



Simply Smart Software

SigmaXL[®] Version 9.0 Workbook

Contact Information:

Technical Support: 1-866-475-2124 (Toll Free in North America) or
1-519-579-5877

Sales: 1-888-SigmaXL (888-744-6295)

E-mail: support@SigmaXL.com Web: www.SigmaXL.com

Table of Contents

SigmaXL® What's New in Version 9.0, System Requirements and Getting Help	1
What's New in Version 9.0.....	2
Activating SigmaXL.....	5
Error Messages	7
SigmaXL® Defaults and Menu Options.....	8
Clear Saved Defaults.....	8
Data Selection Default.....	8
Menu Options (Classical or DMAIC)	9
SigmaXL® System Requirements.....	10
Getting Help and Product Registration	11
Introduction to SigmaXL® Data Format and Tools Summary	13
Introduction	15
The Y=f(X) Model.....	15
Data Types: Continuous Versus Discrete.....	16
Stacked Data Column Format versus Unstacked Multiple Column Format.....	17
Summary of Graphical Tools	18
Summary of Statistical Tools	19
SigmaXL: Measure Phase Tools	21
Part A - Basic Data Manipulation	23
Introduction to Basic Data Manipulation	23
Category Subset	23
Random Subset	24
Numerical Subset	25
Date Subset.....	25
Transpose Data	26
Stack Subgroups Across Rows	26
Stack Columns.....	28
Random Data	29
Box-Cox Transformation.....	29
Standardize Data	29

SigmaXL: Table of Contents

Data Preparation – Remove Blank Rows and Columns	30
Data Preparation – Change Text Data Format to Numeric.....	32
Recall SigmaXL Dialog.....	32
Activate Last Worksheet	32
Part B – Templates & Calculators.....	33
Introduction to Templates & Calculators	33
SigmaXL Templates & Calculators	34
Team/Project Charter	36
SIPOC Diagram	37
Data Measurement Plan	37
Cause & Effect (Fishbone) Diagram.....	38
Cause & Effect (Fishbone) Template	39
Cause & Effect (XY) Matrix Example.....	40
Failure Mode & Effects Analysis (FMEA) Example.....	41
Quality Function Deployment (QFD) Example.....	44
Pugh Concept Selection Matrix Example.....	46
Control Plan	47
Lean – Takt Time Calculator Example.....	48
Lean – Value Analysis/Process Load Balance.....	49
Lean – Value Stream Mapping	50
Basic Graphical Templates – Pareto Example.....	52
Basic Graphical Templates – Histogram Example.....	53
Basic Graphical Templates – Run Chart Example	55
Basic Statistical Templates – Sample Size – Discrete.....	57
Basic Statistical Templates – Sample Size – Continuous.....	59
Basic Statistical Templates – Minimum Sample Size for Robust t-Tests and ANOVA	60
Basic Statistical Templates – 1 Sample Z-Test and Confidence Interval for Mean.....	64
Basic Statistical Templates – 1 Sample t-Test and Confidence Interval for Mean	65
Basic Statistical Templates – 2 Sample t-Test and Confidence Interval (Compare 2 Means).....	66
Basic Statistical Templates – 1 Sample Equivalence Test for Mean.....	67
Basic Statistical Templates – 2 Sample Equivalence Test (Compare 2 Means)	68
Basic Statistical Templates – 1 Sample Chi-Square Test and CI for Standard Deviation	69

Basic Statistical Templates – 2 Sample F-Test and CI (Compare 2 Standard Deviations).....	70
Basic Statistical Templates – 1 Proportion Test and Confidence Interval.....	71
Basic Statistical Templates – 2 Proportions Test and Confidence Interval	73
Basic Statistical Templates – 2 Proportions Equivalence Test	75
Basic Statistical Templates – 1 Poisson Rate Test and Confidence Interval.....	77
Basic Statistical Templates – 2 Poisson Rates Test and Confidence Interval	79
Basic Statistical Templates – 2 Poisson Rates Equivalence Test	81
Basic Statistical Templates – One-Way Chi-Square Goodness-of-Fit Test	83
Basic Statistical Templates – One-Way Chi-Square Goodness-of-Fit Test - Exact.....	84
Probability Distribution Calculators – Normal.....	87
Basic MSA Templates – Type 1 Gage Study.....	88
Basic MSA Templates – Gage Bias and Linearity Study	90
Basic MSA Templates – Gage R&R Study (MSA).....	92
Basic MSA Templates - Gage R&R: Multi-Vari & X-bar R Charts	94
Basic MSA Templates – Attribute Gage R&R (MSA)	95
Basic Process Capability Templates – Process Sigma Level – Discrete Data	97
Basic Process Capability Templates – Process Sigma Level – Continuous Data	98
Basic Process Capability Templates – Process Capability Indices	99
Basic Process Capability Templates – Process Capability & Confidence Intervals	100
Basic Process Capability Templates – Tolerance Interval Calculator (Normal Exact).....	101
Basic Control Chart Templates – Individuals Chart.....	102
Part C – Descriptive/Summary Statistics	105
Descriptive Statistics	105
Descriptive Statistics - Options.....	107
Part D – Histograms	114
Basic Histogram Template.....	114
Single (Basic) Histogram.....	114
Multiple Histograms.....	115
Part E – Dotplots	117
Dotplots	117
Part F – Boxplots	118
Boxplots	118

SigmaXL: Table of Contents

Part G – Normal Probability Plots..... 122
 Normal Probability Plots 122

Part H– Run Charts..... 126
 Basic Run Chart Template 126
 Run Charts..... 126
 Overlay Run Charts..... 129

Part I – Measurement Systems Analysis..... 130
 Type 1 Gage Study Template..... 130
 Gage Bias and Linearity Study Template 130
 Gage R&R Study (MSA) Template..... 130
 Attribute Gage R&R (MSA) Template 130
 Create Gage R&R (Crossed) Worksheet 131
 Analyze Gage R&R (Crossed)..... 133
 Attribute MSA (Binary)..... 137
 Attribute MSA (Ordinal) 148
 Attribute MSA (Nominal)..... 154

Part J – Process Capability..... 156
 Process Capability Templates..... 156
 Histograms and Process Capability..... 157
 Capability Combination Report (Individuals)..... 159
 Capability Combination Report (Subgroups) 161
 Capability Combination Report (Individuals Nonnormal)..... 163
 Distribution Fitting 171
 Box-Cox Transformation..... 175

Part K – Reliability/Weibull Analysis..... 177
 Reliability/Weibull Analysis 177

SigmaXL: Analyze Phase Tools 181

Part A – Stratification with Pareto..... 183
 Basic Pareto Chart Template 184
 Basic (Single) Pareto Charts..... 184
 Advanced (Multiple) Pareto Charts 186

Part B - EZ-Pivot 189

Example of Three X's, No Response Y's..... 189

Example of Three X's and One Y..... 192

Example of 3 X's and 3 Y's 193

Part C – Confidence Intervals 194

 One Sample Confidence Interval Templates 194

 Confidence Intervals – Descriptive Statistics..... 194

Part D – Hypothesis Testing – One Sample Z and t-Test..... 197

 One Sample Z, t and Equivalence Test Templates 197

 One Sample t-Test..... 198

Part E – Power and Sample Size 201

 Power and Sample Size – One Sample t-Test – Customer Data..... 201

 Power and Sample Size – One Sample t-Test – Graphing the Relationships between Power,
 Sample Size, and Difference 205

Part F – One Sample Nonparametric Tests..... 207

 Introduction to Nonparametric Tests..... 207

 Exact Nonparametric Tests 207

 One Sample Sign Test..... 209

 One Sample Wilcoxon Signed Rank Test 210

 One Sample Wilcoxon Signed Rank Test - Exact..... 211

Part G – Two Sample t-Test..... 217

 Two Sample t and Equivalence Test Templates..... 217

 Two Sample t-Test..... 218

 Paired t-Test..... 222

 Unpaired 2 Sample t-Test vs. Paired t-Test 224

 Power & Sample Size for 2 Sample T-Test..... 225

Part H – Two Sample F and Comparison Test..... 227

 Two Sample F-Test Template 227

 Two Sample Comparison Test 228

Part I – Two Sample Nonparametric Test: Mann-Whitney..... 230

 Two Sample Mann-Whitney Test 230

 Two Sample Mann-Whitney Test - Exact..... 231

Part J – One-Way ANOVA & Means Matrix 236

One-Way ANOVA & Means Matrix.....	236
One-Way ANOVA & Means Matrix - Options	239
Power & Sample Size for One-Way ANOVA	245
Part K – Two-Way ANOVA.....	247
Two-Way ANOVA	247
Part L – Tests for Equal Variance & Welch’s ANOVA	251
Bartlett’s Test.....	251
Levene’s Test.....	255
Welch’s ANOVA Test (Assume Unequal Variance)	259
Part M – Nonparametric Multiple Comparison.....	266
Kruskal-Wallis.....	266
Kruskal-Wallis – Exact.....	269
Mood’s Median Test	273
Mood’s Median Test - Exact.....	275
Part N – Nonparametric Runs Test.....	281
Nonparametric Runs Test for Randomness.....	281
Nonparametric Runs Test for Randomness - Examples.....	282
Nonparametric Runs Test for Randomness - Exact	285
Part O – Attribute/Discrete Data Tests.....	290
1 Proportion Test and Confidence Interval Template	290
2 Proportions Test and Confidence Interval Template.....	290
2 Proportions Equivalence Test Template.....	290
One-Way Chi-Square Goodness-of-Fit Template.....	290
One-Way Chi-Square Goodness-of-Fit Template - Exact	291
1 Poisson Rate Test and Confidence Interval Template	291
2 Poisson Rates Test and Confidence Interval Template.....	291
2 Poisson Rates Equivalence Test Template.....	291
2 Proportions Test and Confidence Interval Template Example.....	292
Chi-Square Test – Two-Way Table Data	293
Chi-Square Test – Two-Way Table Data: Advanced Tests and Measures of Association – Nominal Categories	298
Computing Odds Ratio and Confidence Interval for 2x2 Table.....	302

Chi-Square Test – Two-Way Table Data: Advanced Tests and Measures of Association – Ordinal Categories	303
Chi-Square Test (Stacked Column Format Data)	308
Chi-Square Test – Fisher’s Exact	313
Chi-Square Test – Two Way Table Data – Fisher’s Exact	317
References for Fisher’s Exact	321
Power & Sample Size for One Proportion Test.....	322
Power & Sample Size for Two Proportions Test	324
Part P – Analysis of Means (ANOM) Charts	326
Introduction – What is ANOM.....	326
ANOM Normal One-Way.....	330
ANOM Normal Two-Way (with Main Effects and Slice Charts)	332
ANOM Binomial Proportions One-Way.....	340
ANOM Binomial Proportions Two-Way.....	342
ANOM Poisson Rates One-Way.....	344
ANOM Poisson Rates Two-Way.....	346
ANOM Nonparametric Transformed Ranks.....	348
ANOM Variances	352
ANOM Levene Robust Variances.....	354
Part Q – Multi-Vari Charts.....	356
Multi-Vari Charts.....	356
Part R – Scatter Plots.....	360
Scatter Plots	360
Scatter Plot Matrix	362
Part S – Correlation Matrix.....	363
Correlation Matrix.....	363
Reference.....	364
Part T – Multiple Regression	365
Multiple Regression	365
Multiple Regression with Continuous and Categorical Predictors.....	369
Part U – Logistic Regression	371
Binary Logistic Regression	371

Ordinal Logistic Regression	377
SigmaXL: Improve Phase Tools: Design of Experiments (DOE)	381
Part A – Overview of Basic Design of Experiments (DOE) Templates	383
Part B – Three Factor Full Factorial Example Using DOE Template	384
Part C – Design and Analysis of Catapult Full Factorial Experiment	393
Part D – Design and Analysis of Response Surface Experiment – Cake Bake.....	405
Part E – Basic Taguchi DOE Templates	412
Introduction – Taguchi Methods	412
Overview of Taguchi DOE Templates	414
Taguchi L8 (2 Level) Three Factor – Robust Cake Example	418
Taguchi L8 (2 Level) Four Factor – Catapult Example	431
Taguchi L9 (3 Level) Four Factor – Paper Airplane Example	443
SigmaXL: Control Phase Tools: Statistical Process Control (SPC) Charts and Time Series Forecasting	447
Part A - Individuals Charts	449
Individuals Chart Template	449
Individuals Chart	450
Individuals Charts: Advanced Limit Options – Historical Groups	458
Individuals & Moving Range Charts.....	460
Individuals Charts for Nonnormal Data (Box-Cox Transformation)	461
Part B - X-Bar & Range/StDev Charts.....	465
X-Bar & R Charts.....	465
X-Bar & R Charts – Exclude Subgroups	471
X-Bar & S Charts	473
Part C – Attribute Charts (P, NP, C, U)	475
P-Charts.....	475
P-Charts: Advanced Limit Options – Historical Groups.....	476
NP-Charts	478
C Chart Template	479
C-Charts.....	479
U-Charts	481
Part D – P’ & U’ Charts (Laney).....	482

P'-Charts.....	482
U'-Charts	484
Part E – Control Chart Selection Tool	486
Data Types and Definitions	486
Control Chart Selection Tool – Individuals Chart.....	487
Control Chart Selection Tool – X-Bar & R Chart.....	488
Control Chart Selection Tool – P-Chart.....	489
Part F – I-MR-R/S Charts.....	490
I-MR-R Chart.....	490
Part G – Control Chart Templates: Rare Events.....	493
Rare Events T Chart - Introduction	493
Rare Events T Chart - Example	494
Rare Events G Chart - Introduction	499
Rare Events G Chart - Example.....	499
Rare Events Prob G Chart - Introduction	504
Rare Events Prob G Chart - Example	505
Part H – Control Chart Templates: Time-Weighted.....	510
Exponentially Weighted Moving Average (EWMA) Chart - Introduction.....	510
Exponentially Weighted Moving Average (EWMA) Chart - Example	512
Tabular Cumulative Sum (CUSUM) Chart - Introduction	517
Tabular Cumulative Sum (CUSUM) Chart - Example.....	518
Part I – Control Chart Templates: Trend.....	523
Trend Chart - Introduction	523
Trend Chart - Example.....	524
Part J – Control Chart Templates: Average Run Length (ARL) Calculators.....	528
Average Run Length (ARL) - Introduction.....	528
Shewhart ARL.....	529
Attribute P ARL.....	538
Attribute C ARL.....	547
EWMA ARL	552
CUSUM ARL.....	563
Part K – Time Series Forecasting and Control Charts for Autocorrelated Data	573

Introduction – Time Series Forecasting	573
Introduction – Control Charts for Autocorrelated Data	575
Summary of Features in Time Series Forecasting and Control Charts for Autocorrelated Data	576
Time Series Forecasting Menu	577
Run Chart	578
Autocorrelation (ACF/PACF) Plots	594
Cross Correlation (CCF) Plots.....	605
Spectral Density Plot	608
Seasonal Trend Decomposition Plots	614
Seasonal Interaction Plots.....	626
Exponential Smoothing Forecast.....	630
Exponential Smoothing – Multiple Seasonal Decomposition (MSD) Forecast.....	666
Exponential Smoothing Control Chart.....	680
Exponential Smoothing Multiple Seasonal Decomposition (MSD) Control Chart.....	692
ARIMA Forecast.....	700
ARIMA Forecast with Predictors.....	725
ARIMA – Multiple Seasonal Decomposition (MSD) Forecast.....	748
ARIMA Control Chart.....	762
ARIMA Control Chart with Predictors.....	777
ARIMA Multiple Seasonal Decomposition (MSD) Control Chart.....	782
Utilities – Difference Data	790
Utilities – Lag Data	792
Utilities – Interpolate Missing Values	794
SigmaXL Appendix	795
SigmaXL Version 9.0 Feature List Summary	796
Descriptive Statistics Options and Tolerance Interval Calculator (Normal Exact).....	801
Shapiro-Wilk (SW) and Kolmogorov-Smirnov-Lilliefors (KSL) Normality Test	801
Doornik-Hansen (DH) Normality Test.....	801
Confidence and Tolerance Intervals	801
Tolerance Interval Calculator (Normal Exact).....	802
Percentile (Nonparametric) Confidence and Tolerance Intervals.....	803

References for Descriptive Statistics Options and Tolerance Interval Calculator (Normal Exact)	804
Statistical Details for Nonnormal Distributions and Transformations	805
Maximum Likelihood Estimation (MLE)	805
Distributions	806
Transformations	818
Monte Carlo Random Number Generation	822
Pearson Family of Distributions	822
Automatic Best Fit	826
Process Capability Indices (Nonnormal)	826
Control Charts (Nonnormal)	828
Exact and Monte Carlo P-Values for Nonparametric and Contingency Tests	829
Introduction	829
Full Enumeration	829
Exact	830
Monte Carlo	831
Two Tailed (Sided) Tests	831
Sample Sizes for Exact	832
Power	832
Validation of SigmaXL Exact P-Values	833
1 Sample Tests	833
2 Sample Tests	833
K Sample Tests	833
Contingency Tests	834
References for Exact and Monte Carlo P-Values	834
Hypothesis Test Assumptions Report	836
Normality	836
Robustness	836
Outliers (Boxplot Rules)	839
Randomness (Independence)	840
Equal Variance	841
Attribute Measurement Systems Analysis	843

Percent Confidence Intervals (Exact Versus Wilson Score)	843
Kappa	843
Concordance	844
Kendall's Coefficient of Concordance.....	844
Kendall's Correlation Coefficient.....	845
References for Attribute Measurement Systems Analysis	847
Chi-Square Tests and (Contingency) Table Associations	848
Chi-Square Tests for Nominal Categories.....	848
Measures of Association for Nominal Categories.....	849
Tests for Ordinal Categories.....	851
Measures of Association for Ordinal Categories	852
References for Chi-Square Tests and (Contingency) Table Associations	854
Multiple Comparison of Means and Variances (a.k.a. Post-Hoc Tests)	856
One-Way ANOVA	856
Welch's ANOVA (Unequal Variance)	857
Bartlett's Test for Equal Variance.....	857
Levene's Test for Equal Variance.....	857
References for Multiple Comparisons of Means and Variances	858
Analysis of Means (ANOM) Charts	859
ANOM Normal One-Way.....	859
ANOM Normal Two-Way.....	860
ANOM Binomial Proportions One-Way.....	861
ANOM Binomial Proportions Two-Way.....	862
ANOM Poisson Rates One-Way.....	863
ANOM Poisson Rates Two-Way.....	863
ANOM Nonparametric Transformed Ranks.....	864
ANOM Variances	864
ANOM Levene Robust Variances.....	866
References for Analysis of Means (ANOM) Charts	866
Time Series Forecasting and Control Charts for Autocorrelated Data.....	868
Autocorrelation (ACF), Partial Autocorrelation (PACF) and Cross Correlation (CCF)	868
Ljung-Box Test.....	869

Box-Cox Transformation..... 870

Seasonal Frequency 870

Seasonal Trend Decomposition..... 871

Spectral Density and Automatic Detection of Seasonal Frequency..... 871

Seasonally Adjusted Linear Interpolation of Missing Values 872

Information Criteria for Model Comparison..... 872

Forecast Accuracy 873

Exponential Smoothing - ETS..... 874

Autoregressive Integrated Moving Average - ARIMA..... 880

Assessing Forecast Accuracy with Forecast Competitions..... 884

Control Charts for Autocorrelated Data 884

References for Time Series Forecasting and Control Charts for Autocorrelated Data..... 886

About SigmaXL, Inc. 889

SigmaXL® What's New in Version 9.0, System Requirements and Getting Help

What's New in Version 9.0

See [SigmaXL Version 9.0 Feature List Summary](#) for a complete feature list summary.

New features in SigmaXL Version 9.0 include:

- **Powerful and Easy-to-Use Time Series Forecasting and Control Charts for Autocorrelated Data**



- Run Chart
- Autocorrelation Function (ACF)/Partial Autocorrelation (PACF) Plots
- Cross Correlation (CCF) Plots with Pre-Whiten Data option
- Seasonal Trend Decomposition Plots
- Spectral Density Plot with Detection of Seasonal Frequency
- Exponential Smoothing:
 - Forecast with Prediction Intervals
 - Exponential Smoothing models use Rob Hyndman's taxonomy:
 - Additive/Multiplicative Error
 - Additive/Additive Damped Trend
 - Additive/Multiplicative Seasonal
 - This includes all of the classical exponential smoothing models such as Simple/Single/EWMA, Double and Holt-Winters
 - Multiple Seasonal Decomposition (MSD) option
 - Useful for high frequency and/or multiple frequency data, such as Monthly with frequency = 12, Daily with frequency = 7 and Hourly with frequency = 24
- Exponential Smoothing Residuals Control Chart for autocorrelated data:
 - Individuals and Moving Limits (with One-Step Ahead Forecast) Charts
 - Add Data, Show Last 30 Data Points, Enable Scroll options
 - MSD option
- Autoregressive Integrated Moving Average (ARIMA):
 - Forecast with Prediction Intervals
 - ARIMA Forecast with Predictors (Continuous and/or Categorical)
 - MSD option
- ARIMA Residuals Control Chart for autocorrelated data:
 - Individuals and Moving Limits (with One-Step Ahead Forecast) Charts
 - ARIMA Control Chart with Predictors
 - Add Data, Show Last 30 Data Points, Enable Scroll options
 - MSD option

- Utilities:
 - Difference Data
 - Lag Data
 - Interpolate Missing Values (seasonally adjusted linear interpolation)
- Time Series Forecasting Model Features:
 - ARIMA and Exponential Smoothing models are fully automatic or user specified
 - Utilizes modern State Space and Kalman Filter models for accurate parameter estimation
 - ARIMA estimates missing values with Kalman Filter; Exponential Smoothing uses seasonally adjusted linear interpolation
 - Automatic Box-Cox Transformation
 - Automatic seasonal frequency detection
- Model Diagnostics:
 - ACF/PACF Plots
 - Ljung-Box p-values
 - Log-Likelihood, AIC, AICc, BIC, Residual StDev.
 - Residual plots (histogram, normal probability, residual versus fits, residuals versus order)
- Forecast Accuracy:
 - In-Sample (Estimation) one-step-ahead forecast errors (RMSE, MAE, MASE, MAPE)
 - Out-of-Sample (Withhold) one-step-ahead forecast errors
 - Out-of-Sample (Withhold) multi-step-ahead forecast errors
 - Evaluated using the benchmark standard M4 forecast competition data, a total of 100,000 data sets with Yearly, Quarterly, Monthly, Weekly, Daily and Hourly data. Using a hybrid average of automatic Exponential Smoothing and ARIMA, SigmaXL (unofficially) ranked 10th out of 60 in the Overall Weighted Average forecast accuracy score, ahead of three well known commercial forecast software packages.

- **New and Improved Control Charts**
 - New Control Chart Templates
 - Rare Events G and T (Provost-Murray approximation)
 - Rare Events Probability-Based G
 - Trend/Tool Wear
 - Exponentially Weighted Moving Average (EWMA)
 - Tabular Cumulative Sum (CUSUM)
 - Average Run Length (ARL) Calculators
 - Shewhart with Tests for Special Causes
 - Attribute C & P
 - EWMA & CUSUM
 - Markov Chain Approximation - fast and accurate
 - Monte Carlo Simulation - additional Run Length statistics: Standard Deviation and Percentiles
 - Test robustness to non-normality with specified Skewness & Kurtosis
 - Tests for Special Causes now supported for menu-based control charts:
 - Varying Subgroup Sizes (Moving Limits)
 - Historical Groups
 - MR/Range/StDev Charts (Tests 1-4)
 - Anderson-Darling P-Value in Control Chart Process Capability Report

Activating SigmaXL

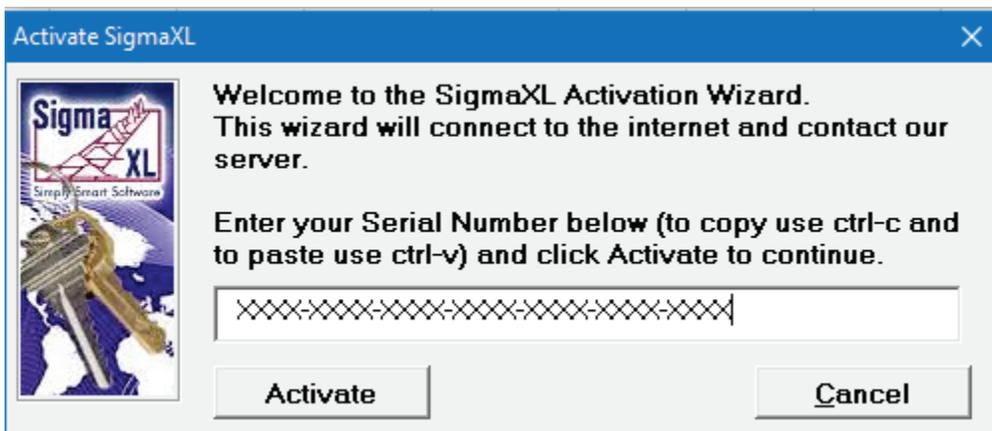
Please proceed with the following steps if you have a valid serial number and your computer is connected to the Internet. If you do not have an Internet connection, please email Support@SigmaXL.com from a different system or call 1-888-SigmaXL (1-888-744-6295) or 1-519-579-5877.

If your trial has timed out and you do not have a serial number but wish to purchase a SigmaXL license, please click **Purchase SigmaXL** in the **Activation Wizard Box** and this will take you to SigmaXL's order page <http://www.sigmaxl.com/Order%20SigmaXL.shtml>. You can also call 1-888-SigmaXL (1-888-744-6295) or 1-519-579-5877 to place an order.

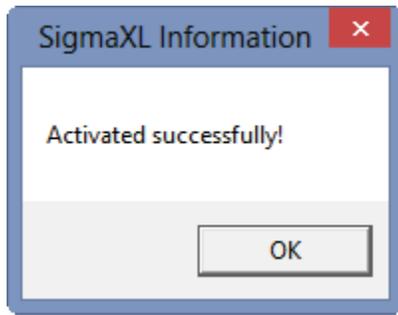
1. In the **Activation Wizard** box select **Activate SigmaXL (Enter a serial number.)**



2. Click **Next**. Enter your serial number. We recommend that you copy and paste the serial number into the serial number field to avoid typos.



3. Click **Next**. The activation process continues and is confirmed as shown:

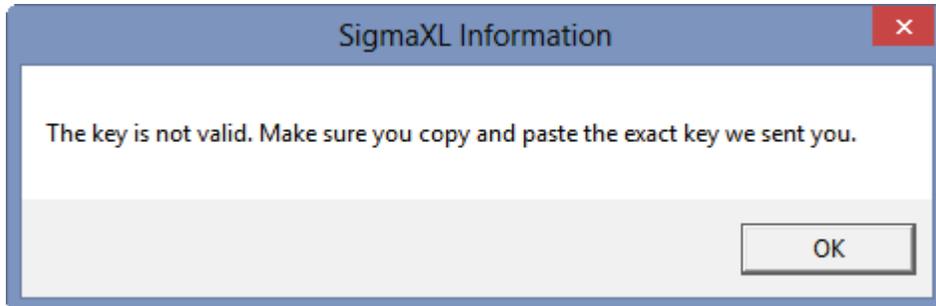


4. Click **OK**. In Excel 2013+ for Windows or Excel 2016+ for Mac, SigmaXL appears as a new Ribbon:



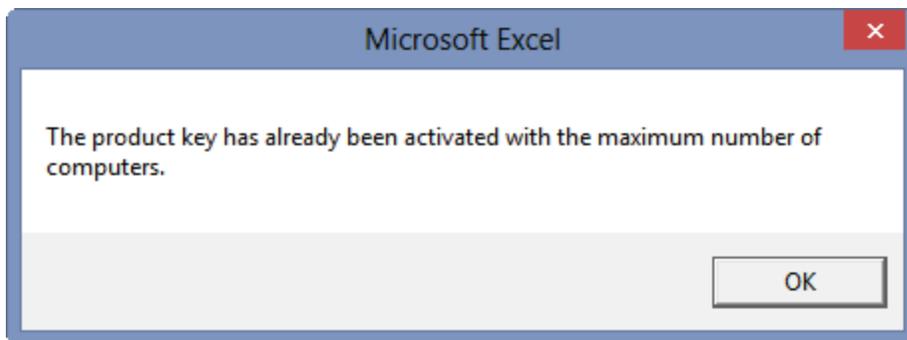
Error Messages

1. Incorrect Serial Number:



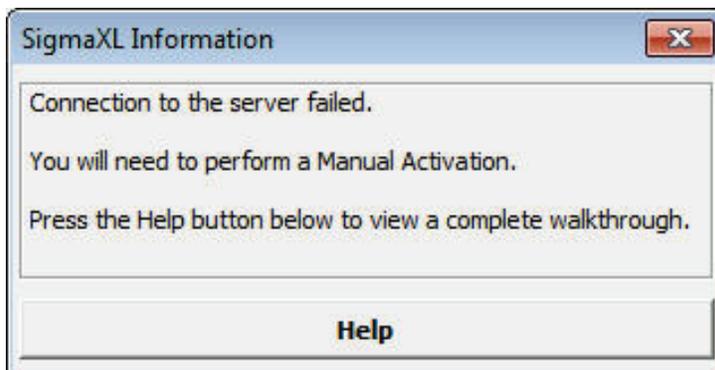
This is due to an incorrect serial number entry in the previous registration window – click **OK** and please reenter your serial number.

2. Serial Number Used in Previous Activation:



The license is currently activated on the maximum allowable number of computers. You will need to deactivate SigmaXL from your old computer while connected to the Internet. To deactivate a SigmaXL license go to **SigmaXL > Help > Deactivate SigmaXL**.

3. Internet Connection Problem:



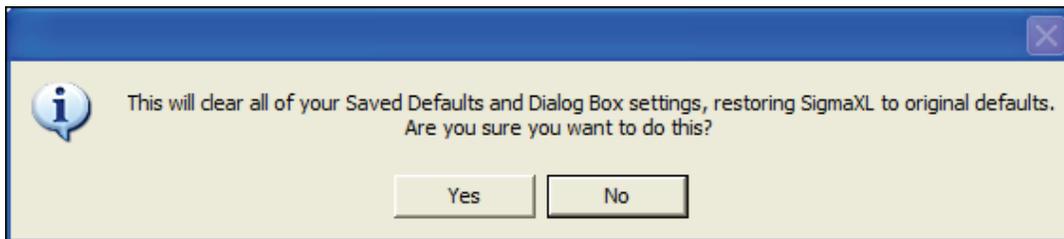
If you do not have an Internet connection, please email Support@SigmaXL.com from a different system or call 1-888-SigmaXL (1-888-744-6295) or 1-519-579-5877.

SigmaXL® Defaults and Menu Options

Clear Saved Defaults

Clear Saved Defaults will reset all saved defaults such as Pareto and Multi-Vari Chart settings, saved control limits, and dialog box settings. All settings are restored to the original installation defaults.

Click **SigmaXL > Help > SigmaXL Defaults > Clear Saved Defaults**. A warning message is given prior to clearing saved defaults.

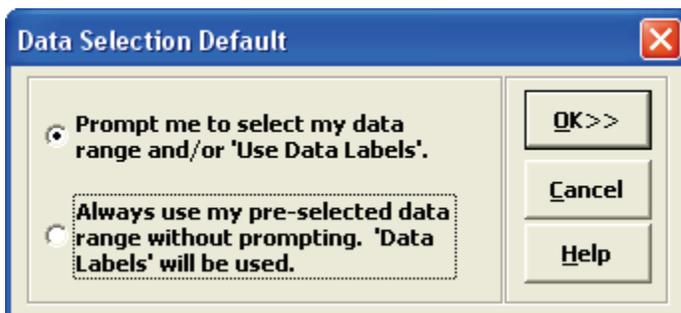


Data Selection Default

The **Data Selection Default** setting is: **Prompt me to select my data range and/or 'Use Data Labels'**.

This can be changed to: **Always use my pre-selected data range without prompting. 'Data Labels' will be used**. This setting saves you from having to click **Next** at the start of every function, but the user is responsible to ensure that the proper data selection is made prior to starting any menu item.

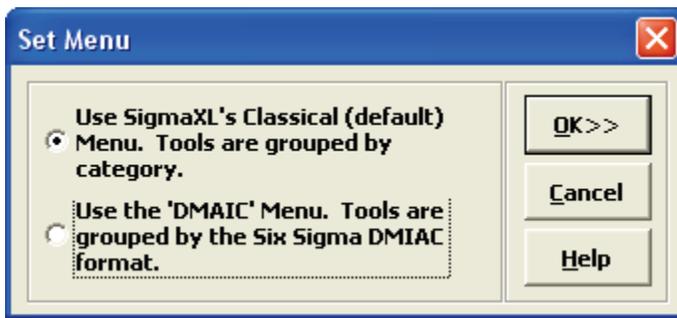
Click **SigmaXL > Help > SigmaXL Defaults > Data Selection Default** to make this change. This will apply permanently unless you revert back to the **Prompt me** setting or click **Clear Saved Defaults** shown above.



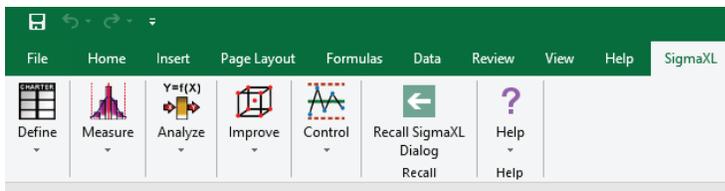
Menu Options (Classical or DMAIC)

The default SigmaXL menu system groups tools by category, but this can be changed to the Six Sigma DMAIC format.

Click **SigmaXL > Help > SigmaXL Defaults > Menu Options – Set SigmaXL's Menu to Classical or DMAIC**. The Set Menu dialog allows you to choose between **Classical** (default) and **DMAIC**:



If you select the **DMAIC** Menu, the **DMAIC** Menu Ribbon appears as shown (Excel 2013 and newer):



All SigmaXL tools are available with this menu format, but they are categorized using the Six Sigma DMAIC phase format. Note that some tools will appear in more than one phase.

This workbook uses the classical (default) menu format, but the chapters are organized as Measure, Analyze, Improve and Control.

SigmaXL® System Requirements

Minimum System Requirements:

Computer and processor: 1.6 gigahertz (GHz) or faster, 2-core.

Memory: 4 GB of RAM or greater.

Hard disk: 1 GB of available hard-disk space.

Display: 1280x768 or higher resolution monitor.

Windows Operating System: Microsoft Windows 10 with current Service Packs, or later operating system.

Microsoft Windows Excel version: Excel 2013 and newer with latest service packs installed.

Mac Operating System: Three most recent versions of macOS.

Mac Excel Version: Excel 2016 or later with latest service packs installed.

Administrative Rights to install software.

Getting Help and Product Registration

To access the help system, please click **SigmaXL > Help > Help**.

Technical support is available by phone at 1-866-475-2124 (toll-free in North America) or 1-519-579-5877 or by e-mail support@sigmaxl.com.

Please note that users obtain free technical support and upgrades for one year from date of purchase. Optional maintenance is available for purchase prior to the anniversary date.

To register by web, simply click **SigmaXL > Help > Register SigmaXL**.

**Introduction to SigmaXL®
Data Format and Tools Summary**

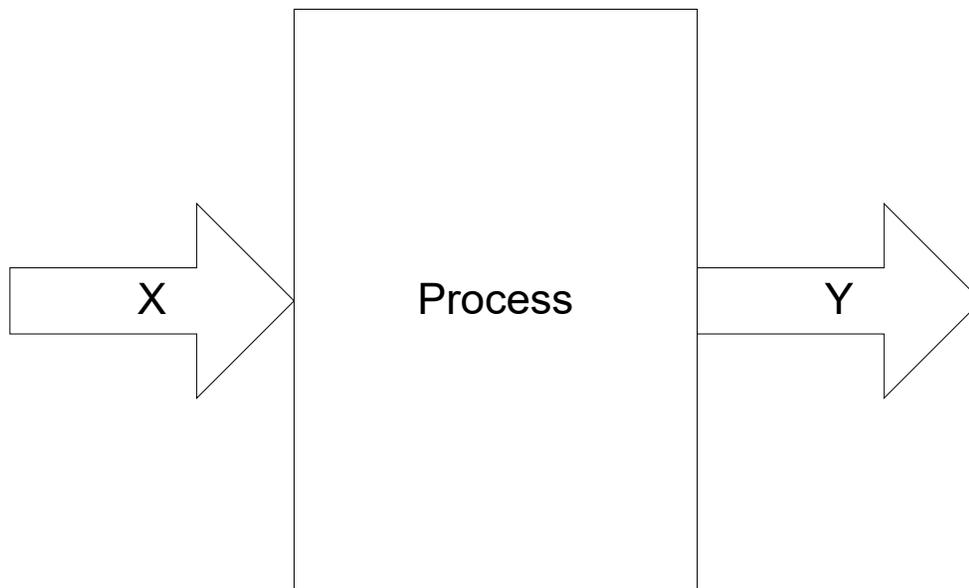
Introduction

SigmaXL is a powerful but easy to use Excel Add-In that will enable you to Measure, Analyze, Improve and Control your service, transactional, and manufacturing processes. This is the ideal cost-effective tool for Six Sigma Green Belts and Black Belts, Quality and Business Professionals, Engineers, and Managers.

SigmaXL will help you in your problem solving and process improvement efforts by enabling you to easily slice and dice your data, quickly separating the “vital few” factors from the “trivial many”. This tool will also help you to identify and validate root causes and sources of variation, which then helps to ensure that you develop permanent corrective actions and/or improvements.

The Y=f(X) Model

SigmaXL utilizes the “Y=f(X)” model in its dialog boxes. Y denotes a key process output metric; X denotes a key process input metric. This process is shown pictorially as:



The mathematical expression $Y = f(X)$ denotes that the variable Y is a function of X. Y can also be viewed as the effect of interest and X is the cause. For example, Y could be customer satisfaction as measured on a survey and X could be location or responsiveness to calls (also measured on a survey). The goal is to figure out which X's from among many possible are the key X's and to what extent do they impact the Y's of interest. Solutions and improvements then focus on those key X's.

Data Types: Continuous Versus Discrete

X and Y metrics can each be continuous or discrete. A continuous measure will have readings on a continuous scale where a mid-point has meaning. For example, in a customer satisfaction survey using a 1 to 5 score, the value 3.5 has meaning. Other examples of continuous measures include cycle time, thickness, and weight. A discrete measure is categorical in nature. If we have Customer Types 1, 2, and 3, customer type 1.5 has no meaning. Other examples of discrete measures include defect counts and number of customer complaints.

It is possible to have various combinations of discrete/continuous X's and discrete/continuous Y's. Some examples are given below:

Examples of Discrete (Category) X and Discrete Y

- X = Customer Type, Y = Number of Complaints
- X = Product Type, Y = Number of Defects
- X = Day Shift vs. Night Shift, Y = Proportion of Defective Units

Examples of Discrete (Category) X and Continuous Y

- X = Customer Type, Y = Customer Satisfaction (1-5)
- X = Before Improvement vs. After Improvement, Y = Customer Satisfaction (1-5)
- X = Location, Y = Order to Delivery Time

Examples of Continuous X and Discrete Y

- X = Responsiveness to Calls (1-5), Y = Number of Complaints
- X = Process Temperature, Y = Number of Defects

Examples of Continuous X and Continuous Y

- X = Responsiveness to Calls (1-5), Y = Customer Satisfaction (1-5)
- X = Amount of Loan (\$), Y = Cycle Time (Loan Application to Approval)

Note that in SigmaXL, a discrete X can be text or numeric, but a continuous X must be numeric. Y's must be numeric. If Y is discrete, count data will be required. If the data of interest is discrete text, it should be referenced as X1 and SigmaXL will automatically search through the text data to obtain a count (applicable for Pareto, Chi-Square and EZ-Pivot tools).

Stacked Data Column Format versus Unstacked Multiple Column Format

SigmaXL can accommodate two data formats: stacked column and unstacked multiple column. The stacked column format has an X column also referred to as the “Group Category” column and a Y column that contains the data of interest. The following is an example of data in stacked column format, with three unique groups of Customer Type:

Customer Type (X)	Overall Satisfaction (Y)
Type 2	3.5
Type 3	3.2
Type 3	3.3
Type 2	4.1
Type 1	3.2
Type 1	2.9
Type 1	1.9
Type 2	3.7
Type 3	4.0
Type 1	2.0
Type 3	2.6
Type 1	3.0
Type 2	4.1
Type 3	3.5
Type 2	5.0
Type 2	4.0
Type 3	4.4
Type 2	4.6
Type 1	2.5

If the data is in unstacked multiple column format, each unique group of X corresponds to a different column. The above data is now shown in unstacked format with customer satisfaction scores for each customer type in separate columns:

Sat_Type 1	Sat_Type 2	Sat_Type 3
3.2	3.5	3.2
2.9	4.1	3.3
1.9	3.7	4
2	4.1	2.6
3	5	3.5
2.5	4	4.4
	4.6	4.2

Summary of Graphical Tools

Tool	What	Type of Data	When to Use	Location in SigmaXL
Pareto Chart	Plots a bar chart of the response in descending order with a cumulative sum line.	Y=Discrete (e.g., Defect Count) or Continuous (e.g., Cost; must be additive) X=Discrete (Category)	To separate the vital few from the trivial many, help specify a problem statement, and prioritize potential root causes. This chart is based on the Pareto principle, which states that typically 80% of the defects in a process or product are caused by only 20% of the possible causes.	<ul style="list-style-type: none"> ➤ SigmaXL > Templates & Calculators > Basic Graphical Templates > Pareto Chart ➤ SigmaXL > Graphical Tools > Basic Pareto Chart (Single) ➤ SigmaXL > Graphical Tools > Advanced Pareto Charts (Multiple)
Pivot Chart	Plots a stacked (or clustered) bar chart from an Excel Pivot Table.	Y=Discrete or Continuous X = Discrete (Category)	To easily 'slice and dice' your data, quickly look at different X factors and their contribution to the total. It is similar to the Pareto Chart, but without the descending bar order.	<ul style="list-style-type: none"> ➤ SigmaXL > Graphical Tools > EZ-Pivot/Pivot Charts
Histogram	Visual display of one variable showing data center, spread, shape and outliers.	Y=Continuous X=Discrete (Category)	<ol style="list-style-type: none"> 1. Summarize large amounts of data 2. To get a 'feel for the data' 3. To compare actual description to customer specifications 	<ul style="list-style-type: none"> ➤ SigmaXL > Templates & Calculators > Basic Graphical Templates > Histogram ➤ SigmaXL > Graphical Tools > Basic Histogram (Single) ➤ SigmaXL > Graphical Tools > Histograms & Descriptive Statistics ➤ SigmaXL > Graphical Tools > Histograms & Process Capability
Dotplots	Visual display of one variable showing data center, spread, shape and outliers.	Y=Continuous X=Discrete (Category)	<ol style="list-style-type: none"> 1. Small sample size ($n < 30$) 2. To get a 'feel for the data' 	<ul style="list-style-type: none"> ➤ SigmaXL > Graphical Tools > Dotplots
Boxplots	Visual display of the summary of Y data grouped by category of X.	Y=Discrete (Category)	Summary display to visualize differences in data center, spread and outliers across categories.	<ul style="list-style-type: none"> ➤ SigmaXL > Graphical Tools > Boxplots
Normal Probability Plot	Plots data in a straight line if the data is normally distributed.	Y=Continuous X=Discrete (Category)	To check for Normality and Outliers.	<ul style="list-style-type: none"> ➤ SigmaXL > Graphical Tools > Normal Probability Plots
Run Charts	Plots observations in time sequence	Y=Continuous or Discrete	To view process performance over time for trends, shifts or cycles. To test for Randomness using the Nonparametric Runs Test	<ul style="list-style-type: none"> ➤ SigmaXL > Templates & Calculators > Basic Graphical Templates > Run Chart ➤ SigmaXL > Graphical Tools > Run Chart ➤ SigmaXL > Graphical Tools > Overlay Run Chart
Multi-Vari Charts	Plots vertical lines with dots to allow comparison of subgroups on one variable.	Y=Continuous X=Discrete (Category)	To visually compare subgroups by individual data points and the mean. To identify major sources of variation (e.g., within a subgroup, between subgroups, or over time).	<ul style="list-style-type: none"> ➤ SigmaXL > Graphical Tools > Multi-Vari Charts
Scatter Plot (Diagram)	Plots a response Y versus a predictor X.	Y=Continuous X=Continuous	To understand the possible relationships between two variables. To identify possible root causes which are related to Y.	<ul style="list-style-type: none"> ➤ SigmaXL > Graphical Tools > Scatter Plot ➤ SigmaXL > Graphical Tools > Scatter Plot Matrix
Analysis of Means (ANOM)	Plots response Y mean for each level of X category with decision limits.	Y=Continuous X=Discrete (Category)	ANOM is a complement to ANOVA showing which group means are significantly different than the grand mean.	<ul style="list-style-type: none"> ➤ SigmaXL > Graphical Tools > Analysis of Means (ANOM)
Control Charts	Plots observations in time sequence against a mean and control limits.	Y=Continuous or Discrete	To monitor the process over time for trends, shifts or cycles in order to control and improve process performance. To identify special causes.	<ul style="list-style-type: none"> ➤ SigmaXL > Control Charts > SigmaXL > Time Series Forecasting > Exponential Smoothing Control Chart ➤ SigmaXL > Time Series Forecasting > ARIMA Control Chart

Summary of Statistical Tools

Tool	What	Type of Data	When to Use	Location in SigmaXL
t-Test	Determine if there is a significant difference between two group means or if the true mean of the data is equal to a standard value.	Y=Continuous X=Discrete (Category)	<ol style="list-style-type: none"> 1. Test if mean = specified value 2. Test if 2 sample means are equal 3. Paired t. to reduce variation when comparing two sample means 4. Multiple pairwise comparisons 	SigmaXL > Templates & Calculators > Basic Statistical Templates > 1 Sample t-Test and Confidence Interval for Mean 2 Sample t-Test and Confidence Interval (Compare 2 Means) SigmaXL > Statistical Tools > 1 Sample t-Test & Confidence Intervals / Paired t-Test 2 Sample t-Test / 2 Sample Comparison Tests One-Way ANOVA & Means Matrix
One-Way ANOVA (Analysis of Variance)	Determine if there is a difference in mean among many groups.	Y=Continuous X=Discrete (Category)	Determine if there is a statistically significant difference in means among the groups.	SigmaXL > Statistical Tools > One-Way ANOVA & Means Matrix (for equal variance) SigmaXL > Statistical Tools > Equal Variance Tests > Welch's ANOVA (for unequal variance) SigmaXL > Statistical Tools > Two-Way ANOVA
Two-Way ANOVA (Analysis of Variance)	Determine if there is a difference in mean among many groups for two factors plus their interaction.	Y=Continuous X1=Discrete (Category) X2=Discrete (Category)	Determine if there is a statistically significant difference in means among the groups.	SigmaXL > Statistical Tools > Two-Way ANOVA
Nonparametric Tests	Determine if there is a difference between two or more group medians or if the median of the data is equal to a standard value.	Y=Continuous X=Discrete (Category)	<ol style="list-style-type: none"> 1. Test if median = specified value; 2. Test if 2 sample medians are equal; 3. Test if there is a difference in medians among the groups: Kruskal-Wallis or Mood's Median 	SigmaXL > Statistical Tools > Nonparametric Tests SigmaXL > Statistical Tools > Nonparametric Tests – Exact (for small sample)
CI for Standard Deviation/ F-test / Bartlett's Test/ Levene's Test	Determine if there is a difference between two or more group variances or determine the confidence interval of a single standard deviation.	Y=Continuous X=Discrete (Category)	<ol style="list-style-type: none"> 1. Determine the confidence interval for a single standard deviation. 2. Test if 2 sample variances (standard deviations) are equal. 3. Determine if there is a statistically significant difference for the variances among the groups. Use Bartlett's test for normal data. Use Levene's test for non-normal data. 	SigmaXL > Templates & Calculators > Basic Statistical Templates > 1 Sample Chi-Square Test and CI for Standard Deviation 2 Sample F-Test and CI (Compare 2 Standard Deviations) SigmaXL > Statistical Tools > Two Sample Comparison Tests SigmaXL > Statistical Tools > Equal Variance Tests > Bartlett / Levene
Proportions Test	Determine if there is a difference between two proportions or determine the confidence interval of a single proportion.	Y=Discrete (Proportion) X=Discrete (Category)	<ol style="list-style-type: none"> 1. Determine the confidence interval for a single proportion. 2. Determine if there is a statistically significant difference for two proportions. 	SigmaXL > Templates & Calculators > Basic Statistical Templates > 1 Proportion Test & Confidence Interval 2 Proportions Test & Confidence Interval
Poisson Rate Test	Determine if there is a difference between two rates or determine the confidence interval of a single rate.	Y=Discrete (Count) X=Discrete (Category)	<ol style="list-style-type: none"> 1. Determine the confidence interval for a single rate. 2. Determine if there is a statistically significant difference for two rates. 	SigmaXL > Templates & Calculators > Basic Statistical Templates > 1 Poisson Rate Test and Confidence Interval 2 Poisson Rates Test and Confidence Interval
One-Way Chi-Square Goodness-of-Fit	Determine if the observed frequencies for one discrete variable are distributed equally	Y=Discrete (Count)	Test the distribution of observed frequency counts against expected (typically uniform discrete).	SigmaXL > Templates & Calculators > Basic Statistical Templates > One-Way Chi-Square Goodness-of-Fit Test One-Way Chi-Square Goodness-of-Fit Test Exact (for small sample)
X ² Chi Square	Determine if there is a difference for observed frequencies of two discrete variables.	Y=Discrete (Count) X=Discrete (Category)	Determine if there is a relationship between two discrete variables.	SigmaXL > Statistical Tools > Chi-Square Tests > Chi-Square Test (for raw discrete data in stacked column format) Chi-Square Test – Two Way Table Data (for pivot or contingency table)
Anderson-Darling Normality Test	Determine if the data is normally distributed.	Y=Continuous	Test if the sample data is normally distributed.	SigmaXL > Statistical Tools > Chi-Square Tests – Exact (for small sample) SigmaXL > Graphical Tools > Histograms & Descriptive Statistics
Correlation	Quantify strength of relationships.	Y=Continuous X=Continuous	Determine if there is evidence of a relationship between Xs and Ys. identify root causes.	SigmaXL > Statistical Tools > Descriptive Statistics SigmaXL > Statistical Tools > Correlation Matrix
Regression (Simple Linear & Multiple)	Summarizes, describes, predicts and quantifies relationships.	Y=Continuous X=Continuous	<ol style="list-style-type: none"> 1. Determine if there is evidence of a relationship between Xs and Ys. 2. Model data to develop a mathematical equation to quantify the relationship. 3. Identify root causes. 4. Make predictions using the model. 	SigmaXL > Statistical Tools > Regression > Multiple Regression
Logistic Regression	Summarizes, describes, predicts and quantifies relationships.	Y=Discrete (Binary or Ordinal) X=Continuous	<ol style="list-style-type: none"> 1. Model data to develop a mathematical equation to quantify the relationship. 3. Identify root causes. 4. Make predictions using the model. 	SigmaXL > Statistical Tools > Regression > Binary Logistic Regression Ordinal Logistic Regression
Design of Experiments (DOE)	Systematic and efficient proactive approach to testing relationships.	Y=Continuous or Discrete X=Continuous or Discrete	To establish cause and effect relationship between Ys and Xs. To identify 'vital few' Xs.	SigmaXL > Design of Experiments

SigmaXL: Measure Phase Tools

Part A - Basic Data Manipulation

Introduction to Basic Data Manipulation

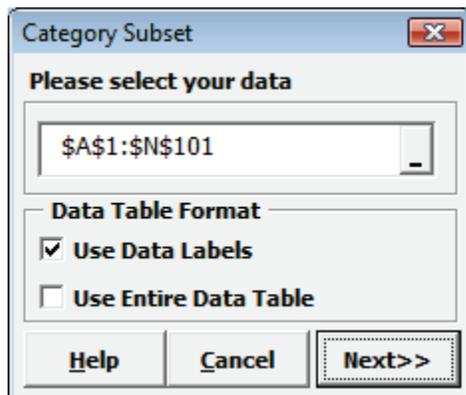
Open **Customer Data.xlsx** (to access, click **SigmaXL > Help > Sample Data** or **Start > Programs > SigmaXL > Sample Data**). This data is in **stacked** column format. This format is highly recommended for use with SigmaXL. Note that all pertinent information is provided in each record (row). Also note that only one row is used for column headings (labels) and there are no blank rows or columns. Each column contains a consistent format of either numeric, text, or date. This is also the data format used by other major statistical software packages.

Customer Record No	Order Date	Customer Type	Avg No. of orders per mo	Avg days Order to delivery time	Loyalty - Likely to Recommend	Overall Satisfaction	Responsive to Calls	Ease of Communications	Staff Knowledge	Size of Customer	Major-Complaint	Product Type
1	1/5/2001	2	24	38	4	3.54	3.02	4.07	3.02	1.65 Small	Return-calls	Consumer
2	1/5/2001	3	36.4	42	4	3.16	3.21	3.11	3.11	3.8 Large	Difficult-to-order	Consumer
3	1/5/2001	2	32.8	44	2	2.42	1.93	2.9	2.9	2.88 Small	Return-calls	Manufacturer
4	1/5/2001	2	47.6	48	3	2.7	1.88	2.52	2.52	4.08 Large	Difficult-to-order	Manufacturer
5	1/5/2001	3	30.6	51	3	3.31	3.75	2.86	2.86	3.88 Small	Not-available	Consumer
6	1/5/2001	2	52.2	55	4	4.12	4.51	3.93	3.93	1.12 Large	Return-calls	Consumer
7	1/5/2001	1	35.8	49	4	3.24	4.06	2.42	2.42	4.64 Large	Return-calls	Manufacturer
8	1/5/2001	2	36.5	39	4	4.47	4.75	4.2	4.2	4.98 Large	Return-calls	Manufacturer
9	1/5/2001	2	39.9	44	4	3.83	3.18	4.48	4.48	3.16 Large	Difficult-to-order	Consumer
10	1/5/2001	1	28	43	3	2.94	2.03	3.85	3.85	4.01 Small	Return-calls	Consumer
11	1/8/2001	2	25.9	44	3	3.24	3.05	4.43	4.43	4.72 Small	Return-calls	Manufacturer
12	1/8/2001	2	23.9	50	3	4.18	3.67	4.69	4.69	4.86 Small	Difficult-to-order	Manufacturer
13	1/8/2001	2	37.9	58	5	4.53	4.29	4.77	4.77	1.9 Large	Return-calls	Consumer

Note that Loyalty, Overall Satisfaction, Responsive to Calls, Ease of Communications, and Staff Knowledge were obtained from surveys. A Likert scale of 1 to 5 was used, with 1 being very dissatisfied, and 5 very satisfied. Survey results were averaged to obtain non-integer results.

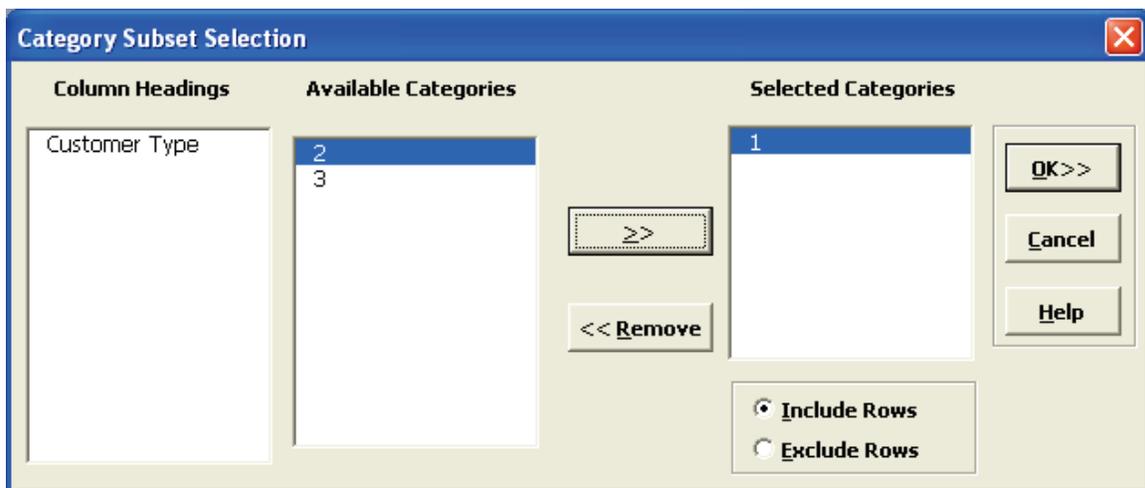
Category Subset

1. Click **SigmaXL > Data Manipulation > Category Subset**.
2. If you are working with a portion of a dataset, specify the appropriate range, otherwise check **Use Entire Data Table**.



3. Click **Next**.

4. Select *Customer Type, 1, >>* as shown:



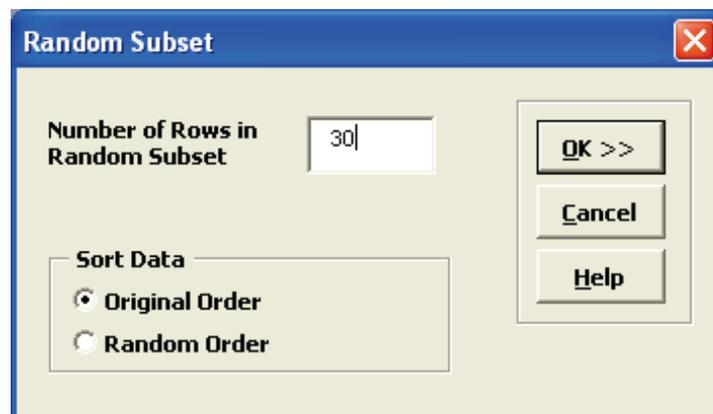
5. Click **OK**.

A new subset worksheet is created containing only Customer Type 1.

Note: We could have chosen more than one Customer Type and had the option to create a subset which included or excluded these Customer Types.

Random Subset

1. Click **Sheet 1** Tab of **Customer Data.xlsx**.
2. Click **SigmaXL > Data Manipulation > Random Subset**.
3. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
4. Enter **Number of Rows in Random Subset** as 30. The default **Sort Data** selection is **Original Order**.



- Click **OK**. A new worksheet is created that contains a random subset of 30 rows. This feature is useful for data collection to ensure a random sample, e.g., given a list of transaction numbers select a random sample of 30 transactions.

Numerical Subset

- Click **Sheet 1** Tab of **Customer Data.xlsx**.
- Click **SigmaXL > Data Manipulation > Numerical Subset**.
- Ensure that entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
- Select *Overall Satisfaction*, **>=**, **Enter Value** as 4.

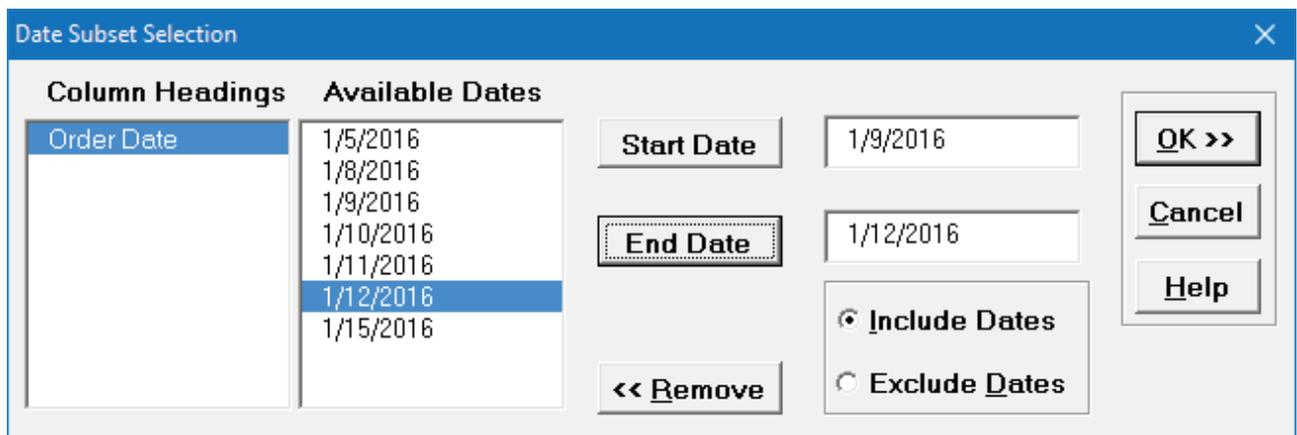
- Click **OK**.

A new subset worksheet is created containing only those rows with Overall Satisfaction ≥ 4 .

Date Subset

- Click **Sheet 1** Tab of **Customer Data.xlsx**.
- Click **SigmaXL > Data Manipulation > Date Subset**.
- Ensure that entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.

4. Select *Order Date*, select 1/9/2016, click **Start Date**, select 1/12/2016, click **End Date**.



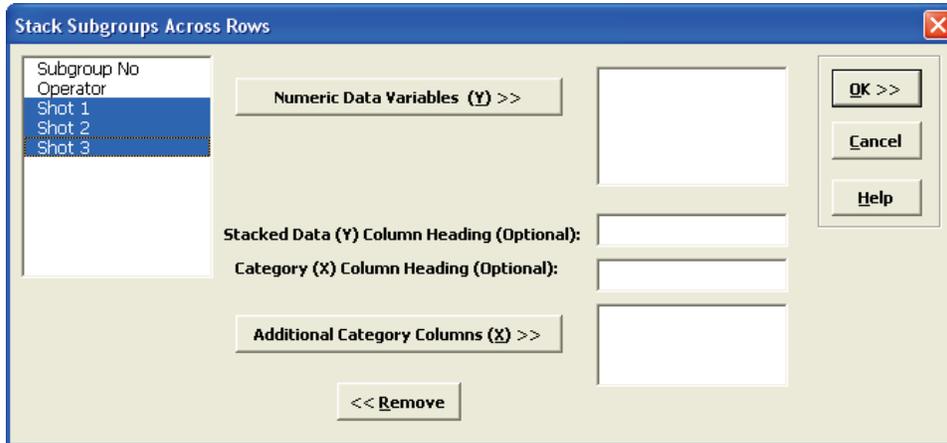
5. Click **OK**. A new subset worksheet is created containing only those rows with Order Date between 1/9/2016 to 1/12/2016.

Transpose Data

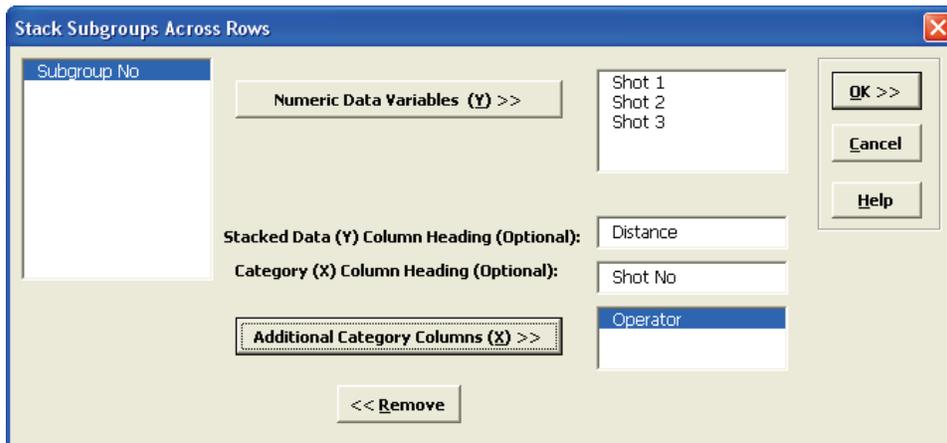
1. Open **Catapult Data Row Format.xlsx**.
2. Manually select the entire data table if the data is not already selected.
3. Click **SigmaXL > Data Manipulation > Transpose Data**.
4. This will transpose rows to columns or columns to rows. It is equivalent to **Copy, Paste Special, Transpose**.

Stack Subgroups Across Rows

1. Now we will stack the subgroups across rows for the transposed data. Ensure that the **Transposed Data Sheet** of **Catapult Data Row Format.xlsx** is active.
2. Click **SigmaXL > Data Manipulation > Stack Subgroups Across Rows**.
3. Check **Use Entire Data Table**, click **Next**.
4. Click on *Shot 1*. Shift Click on *Shot 3* to highlight the three columns of interest.



- Click **Numeric Data Variables (Y) >>** to select *Shot 1 to Shot 3*. Enter *Distance* as the **Stacked Data (Y) Column Heading** and *Shot No* as the **Category (X) Column Heading**. Select *Operator*, click **Additional Category Columns (X) >>**:



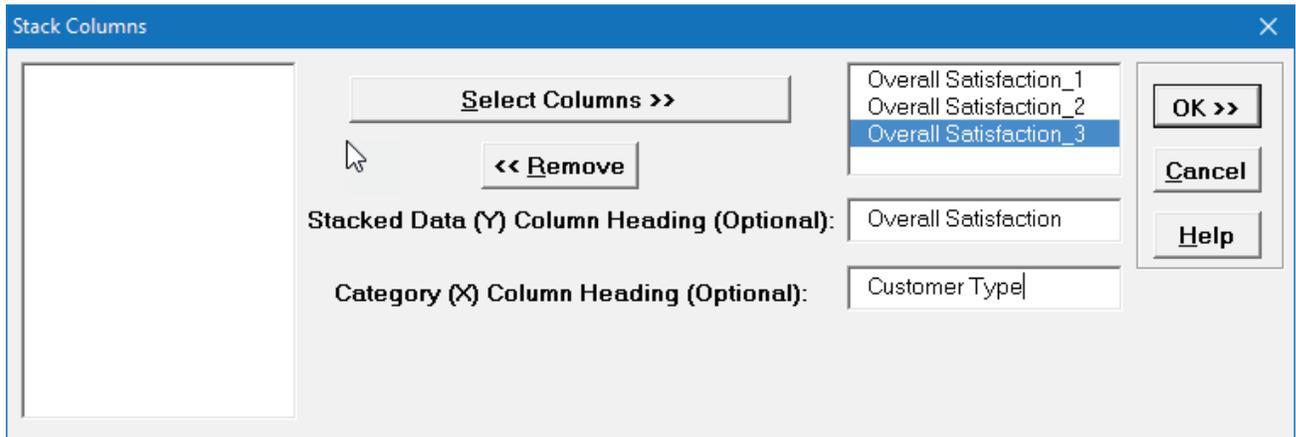
Note that any selected column may be removed by highlighting and double-clicking or clicking the **Remove** button.

- Click **OK**. The resulting stacked data is shown:

Operator	Shot No	Distance
John	Shot 1	105
John	Shot 2	104
John	Shot 3	99
Moe	Shot 1	105
Moe	Shot 2	99
Moe	Shot 3	91
Sally	Shot 1	96
Sally	Shot 2	96
Sally	Shot 3	99
Sue	Shot 1	102
Sue	Shot 2	105
Sue	Shot 3	103
David	Shot 1	100
David	Shot 2	96
David	Shot 3	100

Stack Columns

1. Open **Customer Satisfaction Unstacked.xlsx**.
2. Click **SigmaXL > Data Manipulation > Stack Columns**.
3. Check **Use Entire Data Table**, click **Next**.
4. Shift Click on *Overall Satisfaction_3* to highlight all three column names. Click **Select Columns >>**. Enter the **Stacked Data (Y) Column Heading (Optional)** as *Overall Satisfaction*. Enter the **Category (X) Column Heading (Optional)** as *Customer Type*.



5. Click **OK**. Shown is the resulting stacked column format:

Customer Type	Overall Satisfaction
Overall Satisfaction_1	3.24
Overall Satisfaction_1	2.94
Overall Satisfaction_1	1.86
Overall Satisfaction_1	2.04
Overall Satisfaction_1	2.96
Overall Satisfaction_1	2.53
Overall Satisfaction_1	4.67
Overall Satisfaction_1	4.67
Overall Satisfaction_1	2.57
Overall Satisfaction_1	3.09
Overall Satisfaction_1	3.57
Overall Satisfaction_1	4.25
Overall Satisfaction_1	4.05
Overall Satisfaction_1	3.58
Overall Satisfaction_1	3.82
Overall Satisfaction_1	3.8
Overall Satisfaction_1	2.81
Overall Satisfaction_1	3.99
Overall Satisfaction_1	4.15
Overall Satisfaction_1	3.56
Overall Satisfaction_1	3.26
Overall Satisfaction_1	4.8
Overall Satisfaction_1	1.72
Overall Satisfaction_1	3.01
Overall Satisfaction_1	2.65
Overall Satisfaction_1	3.92
Overall Satisfaction_1	4.24
Overall Satisfaction_1	3.97
Overall Satisfaction_1	4.02
Overall Satisfaction_1	2.56
Overall Satisfaction_1	2.9
Overall Satisfaction_2	3.54
Overall Satisfaction_2	2.42
Overall Satisfaction_2	2.7

6. Data that is in stacked column format can be unstacked using **Data Manipulation > Unstack Columns**.

Random Data

The normal random data generator is used to produce normal random data. Column headings are automatically created with Mean and Standard Deviation values (e.g. 1: Mean = 0; StDev = 1). This utility works with Recall SigmaXL Dialog (**F3**) to append columns to the current Normal Random Data worksheet. An example is shown in **Measure Phase Tools, Part G – Normal Probability Plots**.

Additional random number generators include Uniform (Continuous & Integer), Lognormal, Exponential, Weibull and Triangular. The column headings show the specified parameter values.

Box-Cox Transformation

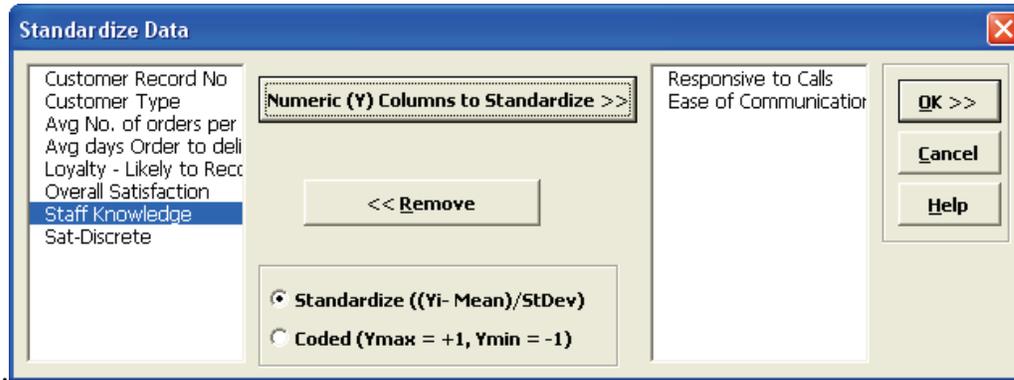
This tool is used to convert nonnormal data to normal by applying a power transformation. Examples of use are given in **Measure Phase Tools, Part J – Process Capability for Nonnormal Data** and **Control Phase Tools, Part A – Individuals Charts for Nonnormal Data**.

Standardize Data

This tool is used to Standardize $((Y_i - \text{Mean})/\text{StDev})$ or Code ($Y_{\text{max}} = +1, Y_{\text{min}} = -1$) your data. This is particularly useful when performing Multiple Regression. Standardized Predictors have better statistical properties. For example, the importance of model coefficients can be determined by the relative size because units are removed. Another statistical benefit is reduced multicollinearity when investigating two-factor interactions.

1. Click **Sheet 1** Tab of **Customer Data.xlsx**.

2. Click **SigmaXL > Data Manipulation > Standardize Data**.
3. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
4. Select *Responsive to Calls* and *Ease of Communications* and click **Numeric (Y) Columns to Standardize >>**



5. Click **OK**. The results are given on the **Standardize** sheet:

Responsive to Calls	Responsive to Calls_Standardize	Ease of Communications	Ease of Communications_Standardize
3.02	-0.740511661	4.07	0.353207938
3.21	-0.573887768	3.11	-0.700161496
1.93	-1.696406629	2.9	-0.93058606
1.88	-1.740255022	2.52	-1.347544794
3.75	-0.100325123	2.86	-0.974476453
4.31	0.390776879	3.93	0.199591563

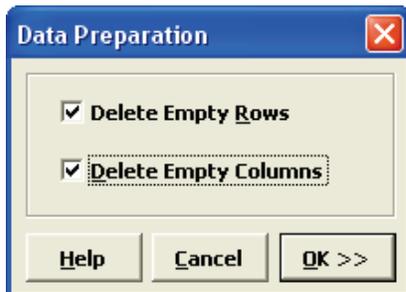
Data Preparation – Remove Blank Rows and Columns

This data preparation utility is provided as a convenient way to prepare data for analysis by deleting any empty rows and/or columns.

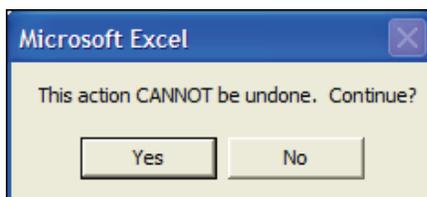
1. Open **Customer Data.xlsx**. Click **Sheet 1** Tab.
2. Insert a new column in B; Click **Column B** heading, **Right-Click > Insert > Columns**.
3. Insert a new row in row 2. Click **Row 2** label, **Right-Click > Insert > Rows** as shown:

	A	B	C	D
1	Customer Record No		Order Date	Customer Type
2				
3	1		1/5/2016	2
4	2		1/5/2016	3
5	3		1/5/2016	2
6	4		1/5/2016	2
7	5		1/5/2016	3
8	6		1/5/2016	2
9	7		1/5/2016	1
10	8		1/5/2016	2
11	9		1/5/2016	2
12	10		1/5/2016	1

- This is now an example of a data set that requires deletion of empty rows and columns. Click **SigmaXL >Data Manipulation >Data Preparation >Remove Blank Rows and Columns**.
- Check **Delete Empty Rows** and **Delete Empty Columns**.



- Click **OK**. A warning message is given prior to the deletion step.



- Click **Yes**. The empty rows and columns are deleted automatically.

	A	B	C
1	Customer Record No	Order Date	Customer Type
2	1	1/5/2001	2
3	2	1/5/2001	3
4	3	1/5/2001	2
5	4	1/5/2001	2
6	5	1/5/2001	3
7	6	1/5/2001	2
8	7	1/5/2001	1
9	8	1/5/2001	2
10	9	1/5/2001	2
11	10	1/5/2001	1

Data Preparation – Change Text Data Format to Numeric

This **Data Preparation** utility will convert data that represents numeric values but are currently in text format. This sometimes occurs when importing data into Excel from another application or text file.

Recall SigmaXL Dialog

Recall SigmaXL Dialog is used to activate the last data worksheet and recall the last dialog, making it very easy to do repetitive analysis. To access, in Excel 2013 and newer, click the **Recall SigmaXL Dialog** menu button as shown below:



Alternatively, you can use the Hot Key **F3** (not available in Excel Mac). This feature can also be accessed by clicking **SigmaXL > Help > Hot Keys > Recall SigmaXL Dialog**.

Note that **Recall SigmaXL Dialog** may not be available for all functions.

Activate Last Worksheet

Activate Last Worksheet is used to activate the last data worksheet without recalling the dialog. To access, press hot key **F4** (not available in Excel Mac). This feature can also be accessed by clicking **SigmaXL > Help > Hot Keys > Activate Last Sheet**.

Part B – Templates & Calculators

Introduction to Templates & Calculators

To use SigmaXL templates, select the appropriate template, enter the inputs and the resulting outputs are produced immediately. If the template does not automatically perform the calculations, click Excel **Formulas > Calculation Options** and select **Automatic** or simply click **Formulas > Calculate Now**. Alternatively, you can click **File > Options > Formulas**, select **Workbook Calculation, Automatic**, and click **OK**.

Some Templates and Calculators are protected worksheets, but this may be modified by clicking **SigmaXL > Help > Unprotect Worksheet**. Alternatively, you can click Excel **File > Info > Unprotect** or **Home > Format > Unprotect Sheet**.

Click **SigmaXL > Templates & Calculators** to access the templates and calculators. Basic Graphical Templates can also be accessed by clicking **SigmaXL > Graphical Tools > Basic Graphical Templates**. Basic Statistical Templates can be accessed by clicking **SigmaXL > Statistical Tools > Basic Statistical Templates**. Basic MSA Templates can be accessed by clicking **SigmaXL > Measurement Systems Analysis > Basic MSA Templates**. Basic Process Capability Templates can be accessed by clicking **SigmaXL > Process Capability > Basic Process Capability Templates**. Basic DOE Templates can also be accessed by clicking **SigmaXL > Design of Experiments > Basic DOE Templates**. Basic Taguchi DOE Templates can also be accessed by clicking **SigmaXL > Design of Experiments > Basic Taguchi DOE Templates**. Basic Control Chart Templates can also be accessed by clicking **SigmaXL > Control Charts > Basic Control Chart Templates**.

SigmaXL Templates & Calculators

- **DMAIC & DFSS Templates:**

- Team/Project Charter
- SIPOC Diagram
- Data Measurement Plan
- Cause & Effect (Fishbone) Diagram and Quick Template
- Cause & Effect (XY) Matrix
- Failure Mode & Effects Analysis (FMEA)
- Quality Function Deployment (QFD)
- Pugh Concept Selection Matrix
- Control Plan

- **Lean Templates:**

- Takt Time Calculator
- Value Analysis/Process Load Balance
- Value Stream Mapping

- **Basic Graphical Templates:**

- Pareto Chart, Histogram and Run Chart

- **Basic Statistical Templates:**

- Sample Size – Discrete and Continuous
- Minimum Sample Size for Robust t-Tests and ANOVA
- 1 Sample Z-Test and Confidence Interval for Mean
- 1 Sample t-Test and Confidence Interval for Mean
- 2 Sample t-Test and Confidence Interval (Compare 2 Means)
- 1 Sample Equivalence Test for Mean
- 2 Sample Equivalence Test (Compare 2 Means)
- 1 Sample Chi-Square Test and CI for Standard Deviation
- 2 Sample F-Test and CI (Compare 2 Standard Deviations)
- 1 Proportion Test and Confidence Interval

- 2 Proportions Test and Confidence Interval
- 2 Proportions Equivalence Test
- 1 Poisson Rate Test and Confidence Interval
- 2 Poisson Rates Test and Confidence Interval
- 2 Poisson Rates Equivalence Test
- One-Way Chi-Square Goodness-of-Fit Test
- One-Way Chi-Square Goodness-of-Fit Test Exact
- **Probability Distribution Calculators:**
 - Normal, Inverse Normal, Lognormal, Exponential, Weibull
 - Binomial, Poisson, Hypergeometric
- **Basic MSA Templates:**
 - Type 1 Gage Study
 - Gage Bias and Linearity Study
 - Gage R&R Study – with Multi-Vari Analysis
 - Attribute Gage R&R (Attribute Agreement Analysis)
- **Basic Process Capability Templates:**
 - Process Sigma Level – Discrete and Continuous
 - Process Capability & Confidence Intervals
 - Tolerance Interval Calculator (Normal Exact)
- **Basic DOE Templates:**
 - 2 to 5 Factors; Main Effects & Interaction Plots
- **Basic Taguchi DOE Templates:**
 - Taguchi L4, L8, L9, L12, L16, L18, L27
- **Control Chart Templates:**
 - Basic: Individuals, C (Count)
 - Rare Events: G, T, Prob G
 - Time Weighted: EWMA, CUSUM
 - Trend
 - Average Run Length (ARL) Calculators: Shewhart, Attribute P & C, EWMA, CUSUM

Team/Project Charter

Click **SigmaXL > Templates & Calculators > DMAIC & DFSS Templates > Team/Project Charter** to access the Team/Project Charter template.

TEAM/PROJECT CHARTER

Project Name:	
Date (Last Revision):	
Prepared By:	
Approved By:	

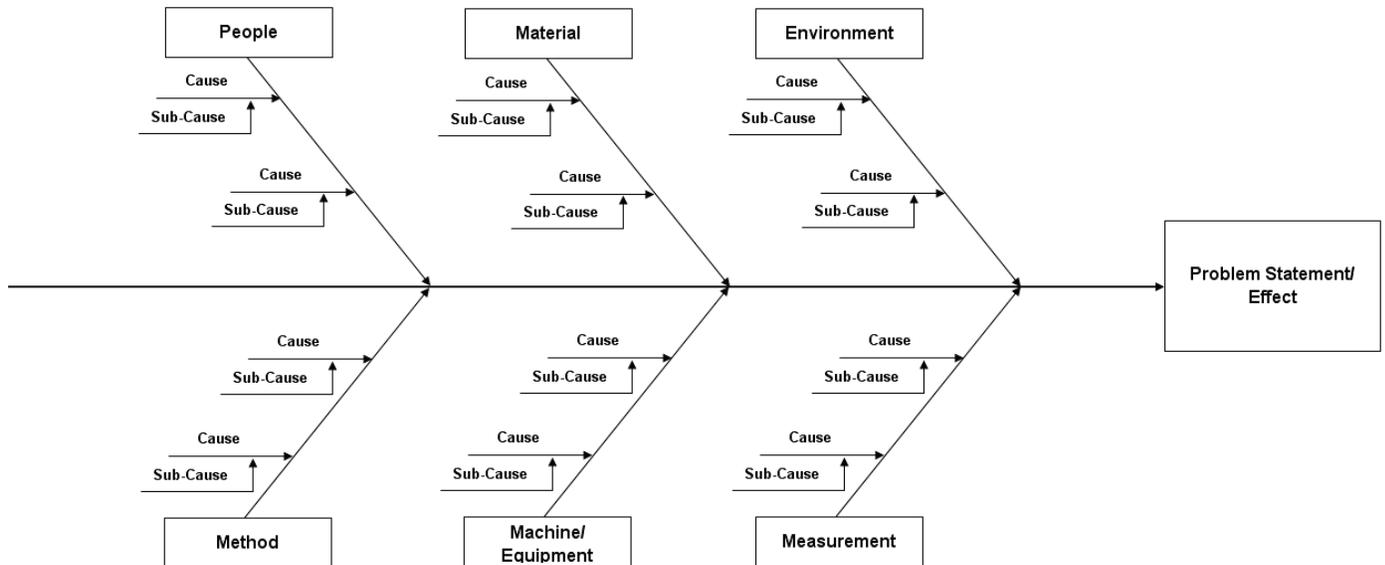
Business Case:				Opportunity Statement (High Level Problem Statement):		
				Defect Definition:		
Goal Statement:				Project Scope:		
				Process Start Point:		
				Process End Point:		
Expected Savings/Benefits:				In Scope:		
				Out of Scope:		
Project Plan:				Team:		
Task/Phase	Start Date	End Date	Actual End	Name:	Role:	Commitment (%)

Cause & Effect (Fishbone) Diagram

Click **SigmaXL > Templates & Calculators > DMAIC & DFSS Templates > Cause & Effect (Fishbone) Diagram** to access the Cause & Effect Diagram form.

CAUSE & EFFECT (FISHBONE) DIAGRAM

Process/Project Name:	
Date:	
Prepared By:	
Notes:	



Notes:

1. Overwrite the text in the Cause, Sub-Cause and Problem Statement/Effect as appropriate.
2. Use copy and paste to create additional causes or sub-causes.
3. The arrows with text are grouped. To ungroup use Excel's Draw > Ungroup tool.

Cause & Effect (Fishbone) Template

Click **SigmaXL > Templates & Calculators > DMAIC & DFSS Templates > Cause & Effect (Fishbone) Template** to access the Cause & Effect Diagram template.

Cause & Effect (Fishbone) Template			
Process/Project Name:			
Date:			
Prepared By:			
Notes:			
Fishbone Diagram			
Problem Statement/Effect:			
Variable	People	Sub Cause 1	Sub Cause 2
X1			
X2			
X3			
X4			
X5			
X6			
X7			
X8			
X9			
X10			
Variable	Method	Sub Cause 1	Sub Cause 2
X1			
X2			
X3			
X4			
X5			
X6			
X7			
X8			
X9			
X10			
Variable	Material	Sub Cause 1	Sub Cause 2
X1			
X2			
X3			
X4			
X5			
X6			
X7			
X8			
X9			
X10			

Simply fill in the template table and click the Fishbone Diagram button to create the diagram.

Cause & Effect (XY) Matrix Example

Click **SigmaXL > Templates & Calculators > DMAIC & DFSS Templates > Cause & Effect (XY) Matrix** to access the Cause & Effect (XY) Matrix template.

The following example is given in **SigmaXL > Help > Template Examples > C&E Matrix**. This is a simple Cause and Effect Matrix example for a Call Center. If prompted, please ensure that macros are enabled in order to allow Pareto Charts to be created.

CAUSE & EFFECT (XY) MATRIX

Process/Project Name:	Call Center Example
Date:	
Performed By:	
Notes:	

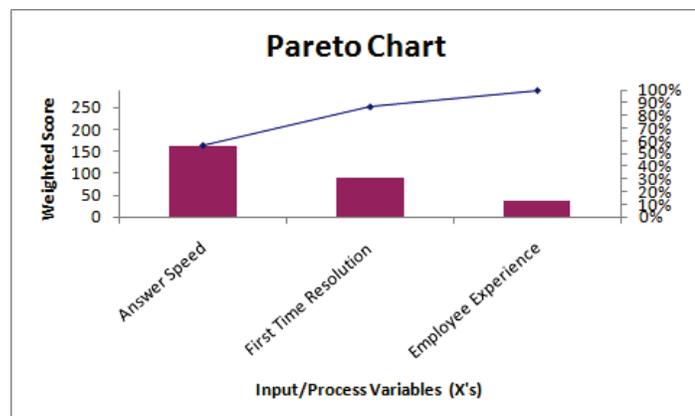
	Call Abandon Rate	Customer Satisfaction	Y3	Y4	Y5	Y6	Y7	Y8	Y9	Y10	Pareto Chart
Output Variables (Y's):											
Importance Score (1-10):	8	10									

Input/Process Variables (X's)	Table of Association Scores (X's to Y's)										Weighted Score
Answer Speed	9	9									162
Employee Experience	1	3									38
First Time Resolution	0	9									90

Notes for use of the Cause and Effect Matrix template, also known as the XY Matrix:

1. Weight the Output Variables (Y's) on a scale of 1 to 10 with 10 indicating most important to the Customer.
2. For Root Cause Analysis, assign the association/effect multiplier score for each X to Y using a scale of 0, 1, 3, 9, where 0 = None, 1 = Weak, 3 = Moderate, and 9 = Strong. Initially this assignment will likely be a team subjective assessment. Data should be collected and the degree of association should be validated with Graphical and Statistical Tools.
3. For Project Selection or Solution Selection, assign the association multiplier score for each X to Y using a scale of 1 to 10, with 10 indicating strong association.

Click the Pareto Chart button to create a Pareto Chart of Weighted Scores by Input/Process Variables:



Failure Mode & Effects Analysis (FMEA) Example

Click **SigmaXL > Templates & Calculators > DMAIC & DFSS Templates > Failure Mode & Effects Analysis (FMEA) template.**

The following example is given in **SigmaXL > Help > Template Examples > FMEA.** This is a simple Failure Mode and Effects example for Stocking Inventory. If prompted, please ensure that macros are enabled in order to allow sort by **RPN**.

Potential Failure Mode & Effects Analysis

		Process/Product:	Stock Inventory							
		FMEA Team:								
		Responsibility:								
		Prepared By:								
Row Number	Process									
	Process Steps or Product Functions	Potential Failure Mode	Potential Effects of Failure	Severity (1-10)	Potential Cause(s) of Failure	Occurrence (1-10)	Current Controls	Detection (1-10)	Risk Priority Number (RPN)	
Sort									Sort	
1	Stock Inventory	Stock in wrong location	Unable to locate Stock	5	Correct location is full	7	Stock checked twice a year	9	315	
2	Stock Inventory	Damaged	Insufficient	7	Supplier Defect	3	Incoming	8	168	
3	Stock Inventory	Damaged	Insufficient product	7	Handling Error	5	Standard Operating	9	315	

If you hover the mouse cursor over the **Severity, Occurrence** or **Detection** heading the recommended scale will appear as a comment:

Severity (1-10)	Score	Severity Guidelines	
		Automotive Industry Action Group AIAG (PFMEA 4th ed.)	Six Sigma
10	Failure to Meet Safety and/or Regulatory Requirements: May endanger operator (machine or assembly) without warning.	Injure a customer or employee	Bad ↓ Good
9	Failure to Meet Safety and/or Regulatory Requirements: May endanger operator (machine or assembly) with warning.	Be illegal	
8	Major Disruption: 100% of product may have to be scrapped. Line shutdown or stop ship.	Render product or service unfit for use	
7	Significant Disruption: A portion of the production run may have to be scrapped. Deviation from primary process including decreased line speed or added manpower.	Cause extreme customer dissatisfaction	
6	Moderate Disruption: 100% of production run may have to be reworked off line and accepted.	Result in partial malfunction	
5	Moderate Disruption: A portion of the production run may have to be reworked off line and accepted.	Cause a loss of performance which is likely to result in a complaint	
4	Moderate Disruption: 100% of production run may have to be reworked in station before it is processed.	Cause minor performance loss	
3	Moderate Disruption: A portion of the production run may have to be reworked in station before it is processed.	Cause a minor nuisance but can be overcome with no performance loss	
2	Minor Disruption: Slight inconvenience to process, operation, or operator.	Be unnoticed and have only minor effect on performance	
1	No Effect: No discernible effect.	Be unnoticed and not affect the performance	

SigmaXL: Measure Phase Tools

The recommended scales for Severity, Occurrence, and Detection are shown below:

Score	Severity Guidelines	
	Automotive Industry Action Group AIAG (PFMEA 4th ed.)	Six Sigma
10	Failure to Meet Safety and/or Regulatory Requirements: May endanger operator (machine or assembly) without warning.	Injure a customer or employee
9	Failure to Meet Safety and/or Regulatory Requirements: May endanger operator (machine or assembly) with warning.	Be illegal
8	Major Disruption: 100% of product may have to be scrapped. Line shutdown or stop ship.	Render product or service unfit for use
7	Significant Disruption: A portion of the production run may have to be scrapped. Deviation from primary process including decreased line speed or added manpower.	Cause extreme customer dissatisfaction
6	Moderate Disruption: 100% of production run may have to be reworked off line and accepted.	Result in partial malfunction
5	Moderate Disruption: A portion of the production run may have to be reworked off line and accepted.	Cause a loss of performance which is likely to result in a complaint
4	Moderate Disruption: 100% of production run may have to be reworked in station before it is processed.	Cause minor performance loss
3	Moderate Disruption: A portion of the production run may have to be reworked in station before it is processed.	Cause a minor nuisance but can be overcome with no performance loss
2	Minor Disruption: Slight inconvenience to process, operation, or operator.	Be unnoticed and have only minor effect on performance
1	No Effect: No discernible effect.	Be unnoticed and not affect the performance

Bad
↓
Good

Score	Occurrence Guidelines	
	Automotive Industry Action Group AIAG (PFMEA 4th ed.)	Six Sigma
10	Very High: ≥ 100 per thousand; ≥ 1 in 10	More than once per day > 30%
9	High: 50 per thousand; 1 in 20	Once every 3-4 days < 30%
8	High: 20 per thousand; 1 in 50	Once every week < 5%
7	High: 10 per thousand; 1 in 100	Once per month < 1%
6	Moderate: 2 per thousand; 1 in 500	Once every 3 months < 0.03%
5	Moderate: 0.5 per thousand; 1 in 2,000	Once every 6 months < 1 per 10,000
4	Moderate: 0.1 per thousand; 1 in 10,000	Once per year < 6 per 100,000
3	Low: 0.01 per thousand; 1 in 100,000	Once every 1-3 years < 6 per million
2	Low: ≤ 0.001 per thousand; 1 in 1,000,000	Once every 3-6 years < 3 per 10 million
1	Very Low: Failure is eliminated through preventive control.	Once every 6-100 years < 2 per billion

Bad
↓
Good

Score	Detection Guidelines	
	Automotive Industry Action Group AIAG (PFMEA 4th ed.)	Six Sigma
10	Almost Impossible. No detection opportunity: No current process control; cannot detect or is not analyzed.	Defect caused by failure is not detectable
9	Very Remote. Not likely to detect at any stage: Failure Mode and/or Error (Cause) is not easily detected (e.g., random audits).	Occasional units are checked for defects
8	Remote. Problem Detection Post Processing: Failure Mode detection post-processing by operator through visual/tactile/audible means.	Units are systematically sampled and inspected
7	Very Low. Problem Detection at Source: Failure Mode detection in-station by operator through visual/tactile/audible means or post-processing through use of attribute gauging (go/no-go, manual torque check/clicker wrench, etc.).	All units are manually inspected
6	Low. Problem Detection Post Processing: Failure Mode detection post-processing by operator through use of variable gauging or in-station by operator through use of attribute gauging (go/no-go, manual torque check/clicker wrench, etc.).	Manual inspection with mistake-proofing modifications
5	Moderate. Problem Detection at Source: Failure Mode or Error (Cause) detection in-station by operator through use of variable gauging or by automated controls in-station that will detect discrepant part and notify operator (light, buzzer, etc.). Gauging performed on setup and first-piece check (for set-up causes only).	Process is monitored (SPC) and manually inspected
4	Moderately High. Problem Detection Post Processing: Failure Mode detection post-processing by automated controls that will detect discrepant part and lock part to prevent further processing.	SPC is used with an immediate reaction to out of control conditions
3	High. Problem Detection at Source: Failure Mode detection in-station by automated controls that will detect discrepant part and automatically lock part in station to prevent further processing.	SPC as above, 100% inspection surrounding out of control conditions
2	Very High. Error Detection and/or Problem Prevention: Error (Cause) detection in-station by automated controls that will detect error and prevent discrepant part from being made.	All units are automatically inspected
1	Almost Certain. Detection not applicable; Error Prevention: Error (Cause) prevention as a result of fixture design, machine design or part design. Discrepant parts cannot be made because item has been error-proofed by process/product design.	Defect is obvious and can be kept from affecting the customer

Bad
↓
Good

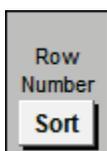
Click **Risk Priority Number (RPN) Sort** to sort the rows by RPN in descending order.



Click **Revised Risk Priority Number Sort** to sort the rows by Revised RPN in descending order.



Click **Row Number Sort** to restore the FMEA worksheet to the original row order (ascending).



Quality Function Deployment (QFD) Example

Click **SigmaXL > Templates & Calculators > DMAIC & DFSS Templates > Quality Function Deployment (QFD) > 10 by 10 QFD** (or 20 by 20, 30 by 30) to access the Quality Function Deployment (QFD) template. Enter data as shown.

Desired Direction of Improvement: ↑ 0 ↑ ↑

Customer Requirements (What):	Customer Importance Rating	Functional Requirements (How):										Satisfaction Ratings:							
		Functional Requirement 1	Functional Requirement 2	Functional Requirement 3	Functional Requirement 4	Functional Requirement 5	Functional Requirement 6	Functional Requirement 7	Functional Requirement 8	Functional Requirement 9	Functional Requirement 10	Our Customer Satisfaction (CS) Rating	Competitor 1 Rating	Competitor 2 Rating	Competitor 3 Rating	Competitor 4 Rating	CS - Max(Competitor Rating)		
Customer Requirement 1	4	9	3		9										3	2			1
Customer Requirement 2	5	9	9	9	9										4	4			0
Customer Requirement 3	2	3	3	1	9										4	3			1
Customer Requirement 4	3	9	1	1	9										1	3			-2
Customer Requirement 5																			0
Customer Requirement 6																			0
Customer Requirement 7																			0
Customer Requirement 8																			0
Customer Requirement 9																			0
Customer Requirement 10																			0
Raw Score:		114	66	50	126	0	0	0	0	0	0	0	0	0					
Rank:		2	3	4	1	5	5	5	5	5	5	5	5	5					

Target:																			
Upper Spec Limit:																			
Lower Spec Limit:																			
Units:																			
Our Current Performance:																			
Benchmark A:																			
Benchmark B:																			
Difficulty:																			
Priority to Improve:																			

Scroll down to view the QFD drop-down lists. The values in yellow highlight may be modified.

Correlations	
Strong Positive	++
Positive	+
None	
Negative	-
Strong Negative	--

Direction of Improvement	
Maximize	↑
Target	0
Minimize	↓

Relationships	
None	
Weak	1
Moderate	3
Strong	9

Customer Importance Rating	
Low	1
	2
Medium	3
	4
High	5

Satisfaction Ratings	
Low	1
	2
Medium	3
	4
High	5

Difficulty	
Very Easy	1
	2
	3
	4
Very Difficult	5

Priority To Improve	
Low Priority	1
	2
Medium Priority	3
	4
High Priority	5

Pugh Concept Selection Matrix Example

Click **SigmaXL > Templates & Calculators > DMAIC & DFSS Templates > Pugh Concept Selection Matrix** to access the Pugh Concept Selection Matrix template. Enter data as shown.

Key Criteria	Weight	Concept A	Concept B	Concept C	Concept D	Concept E	Concept F	Concept G	Concept H	Concept I	Concept J	Current Baseline Datum
Criterion 1	4	+	S	-								S
Criterion 2	5	+	S	S								S
Criterion 3	3	S	+	-								S
Criterion 4	2	+	S	S								S
Criterion 5	5	+	+	S								S
Criterion 6												S
Criterion 7												S
Criterion 8												S
Criterion 9												S
Criterion 10												S
Criterion 11												S
Criterion 12												S
Criterion 13												S
Criterion 14												S
Criterion 15												S
Criterion 16												S
Criterion 17												S
Criterion 18												S
Criterion 19												S
Criterion 20												S

Sum of Positives (+):	4	2	0	0	0	0	0	0	0	0	0	0
Sum of Negatives (-):	0	0	2	0	0	0	0	0	0	0	0	0
Sum of Sames (S):	1	3	3	0	0	0	0	0	0	0	0	0
Positives - Negatives:	4	2	-2	0	0	0	0	0	0	0	0	0

Weighted Sum of Positives (+):	16	8	0	0	0	0	0	0	0	0	0	0
Weighted Sum of Negatives (-):	0	0	7	0	0	0	0	0	0	0	0	0
Weighted Sum of Sames (S):	3	11	12	0	0	0	0	0	0	0	0	0
Weighted Positives - Weighted Negatives:	16	8	-7	0	0	0	0	0	0	0	0	0

Lean – Takt Time Calculator Example

Click **SigmaXL > Templates & Calculators > Lean > Takt Time Calculator** to access the Takt Time calculator. The template gives the following default example:

SigmaXL Lean Templates: Takt Time Calculator		
Daily Customer Demand:	units per day	<input type="text" value="22"/>
Scheduled Work:	hours per shift	<input type="text" value="8"/>
Shifts per Day:		<input type="text" value="2"/>
Lunch:	minutes per shift	<input type="text" value="30"/>
Breaks:	minutes per shift	<input type="text" value="30"/>
Planned Downtime:	minutes per shift	<input type="text" value="30"/>
Staff/Operator Cycle Time:	minutes per unit	<input type="text" value="226"/>
Available Time:	minutes per day	<input type="text" value="780.0"/>
Takt Time:	minutes per unit	<input type="text" value="35.5"/>
Required Number of Staff/Operators:		<input type="text" value="6.4"/>

Lean – Value Analysis/Process Load Balance

Click **SigmaXL > Templates & Calculators > Lean > Value Analysis/Process Load Balance** to access the Value Analysis/Process Load Balance template. The template gives the following default example:

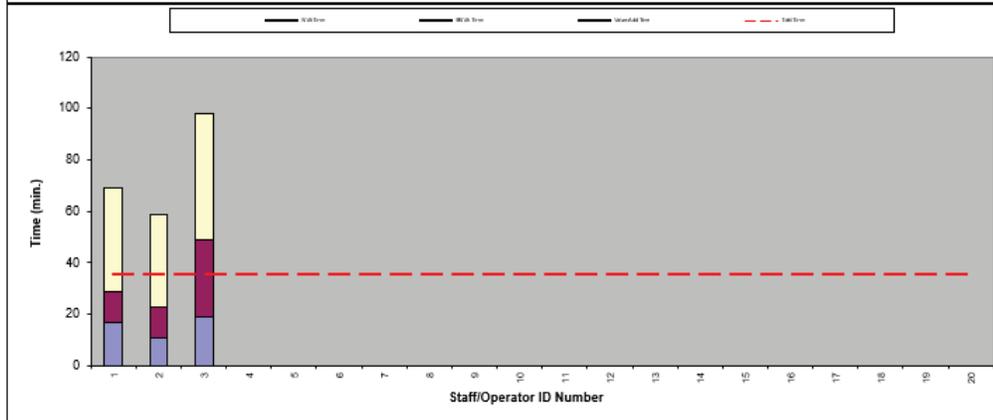
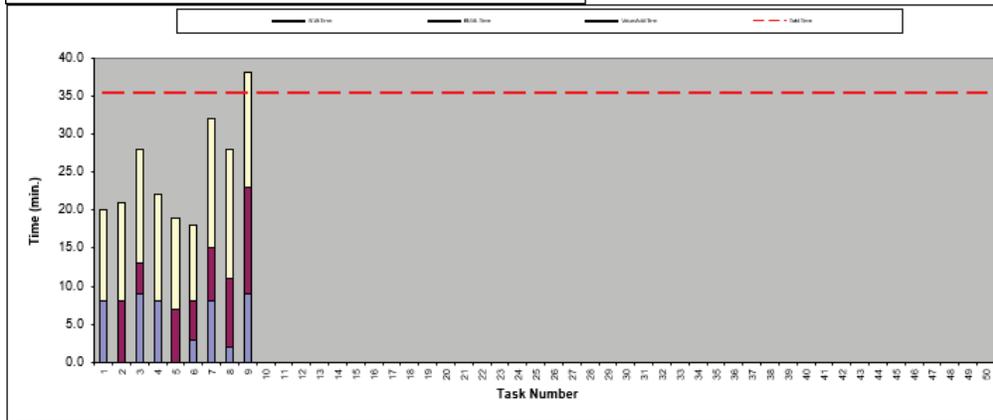
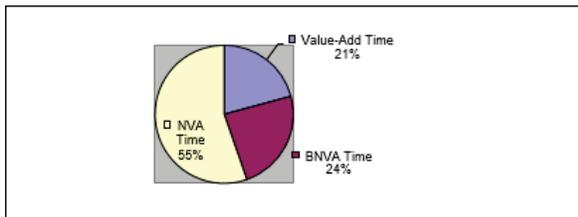
SigmaXL Lean Templates: Value Analysis/Process Load Balance

Process Name:	
Date:	
Prepared By:	
Notes:	

SigmaXL Copyright © 2008

Enter Takt Time:

Task Number	Staff/Operator ID	Task Description	Value-Add Time	BNVA Time	NVA Time	Total Time	Comments
1	1	Task 1	8.0	0.0	12.0	20.0	
2	1	Task 2	0.0	8.0	13.0	21.0	
3	1	Task 3	9.0	4.0	15.0	28.0	
4	2	Task 4	8.0	0.0	14.0	22.0	
5	2	Task 5	0.0	7.0	12.0	19.0	
6	2	Task 6	3.0	5.0	10.0	18.0	
7	3	Task 7	8.0	7.0	17.0	32.0	
8	3	Task 8	2.0	9.0	17.0	28.0	
9	3	Task 9	9.0	14.0	15.0	38.0	
10						0.0	



Lean – Value Stream Mapping

Click **SigmaXL > Templates & Calculators > Lean > Value Stream Mapping** to access the Value Stream Mapping template. The example on the following page is given in **SigmaXL > Help > Template Examples > Value Stream Mapping**.

Process, Material Flow & General Symbols		
<p>Customer/Supplier</p>  <p>Supplier when placed in the upper left of VSM, Customer when in the upper right</p>	<p>External Shipment</p>  <p>External shipment from suppliers or to customers (truck, rail or boat)</p>	<p>Supermarket</p>  <p>Small inventory used when continuous flow is not practical</p>
<p>Process</p>  <p>Process, department, operation or machine, through which material flows</p>	<p>External Shipment (Air)</p>  <p>Air shipment from suppliers or to customers</p>	<p>Safety Stock</p>  <p>Temporary safety stock</p>
<p>Data Box</p>  <p>Typical data includes: Cycle Time (C/T), Operator Cycle Time, Auto Cycle Time, Changover Time (C/O), Batch Size, % Uptime, Available Capacity, First Pass Yield (FPY), Rolled Throughput Yield (RTY), Value Added Time (VA), Scrap Rate, Distance Travelled</p>	<p>Inventory</p>  <p>Inventory count</p>	<p>Material Pull</p>  <p>Physical removal from supermarket</p>
<p>Data Box</p>  <p>Miscellaneous data or other information</p>	<p>Push Arrow</p>  <p>Movement of material that are "pushed" by the production process regardless of the needs of the downstream process</p>	<p>Warehouse</p>  <p>Inventory storage (goal should be to reduce or eliminate)</p>
<p>Work Cell</p>  <p>Multiple processes integrated in a workcell</p>	<p>Material Movement</p>  <p>Finished goods or raw material movement</p>	<p>Kaizen Burst</p>  <p>Highlight area for improvement, to be addressed in a Kaizen Blitz Event</p>
<p>Operator/People</p>  <p>Number of operators or employees required to perform an operation</p>	<p>FIFO Lane</p>  <p>First-In First-Out inventory</p>	<p>VA/NVA Timeline</p>  <p>Highlight value added cycle times versus non-value added wait times</p>

Information Symbols	
<p>PRODUCTION CONTROL</p> <p>Central production scheduling or control</p>	<p>Production Kanban</p> <p>Quantity</p> <p>Triggers production of a pre-defined number of units.</p>
<p>Information Type</p> <p>Weekly Schedule</p> <p>Monthly Forecast</p> <p>Information type includes: Monthly Forecast, Weekly Orders or Schedule, Daily Orders or Schedule</p>	<p>Withdrawal Kanban</p> <p>Quantity</p> <p>Instruction to obtain items from the supermarket</p>
<p>Manual Information Flow</p> <p>Manual information flow from memos, reports, or verbal. Add information type icon above as necessary.</p>	<p>Signal Kanban</p> <p>S</p> <p>Initiate a batch operation</p>
<p>Electronic Information Flow</p> <p>Information flow via computer network. Add information type icon as necessary.</p>	<p>Kanban Post</p> <p>A physical location for kanban cards</p>
<p>MRP/ERP</p> <p>MRP</p> <p>Schedule using MRP/ERP system</p>	<p>Load Leveling</p> <p>XOXO</p> <p>Batch kanbans to load level production</p>
<p>Go See</p> <p>Gather information visually</p>	<p>Sequenced Pull</p> <p>Kanban signal, without using a supermarket</p>

SigmaXL Inc., Copyright © 2011-2020

Portions courtesy and copyright of Strategos, Inc. www.strategosinc.com

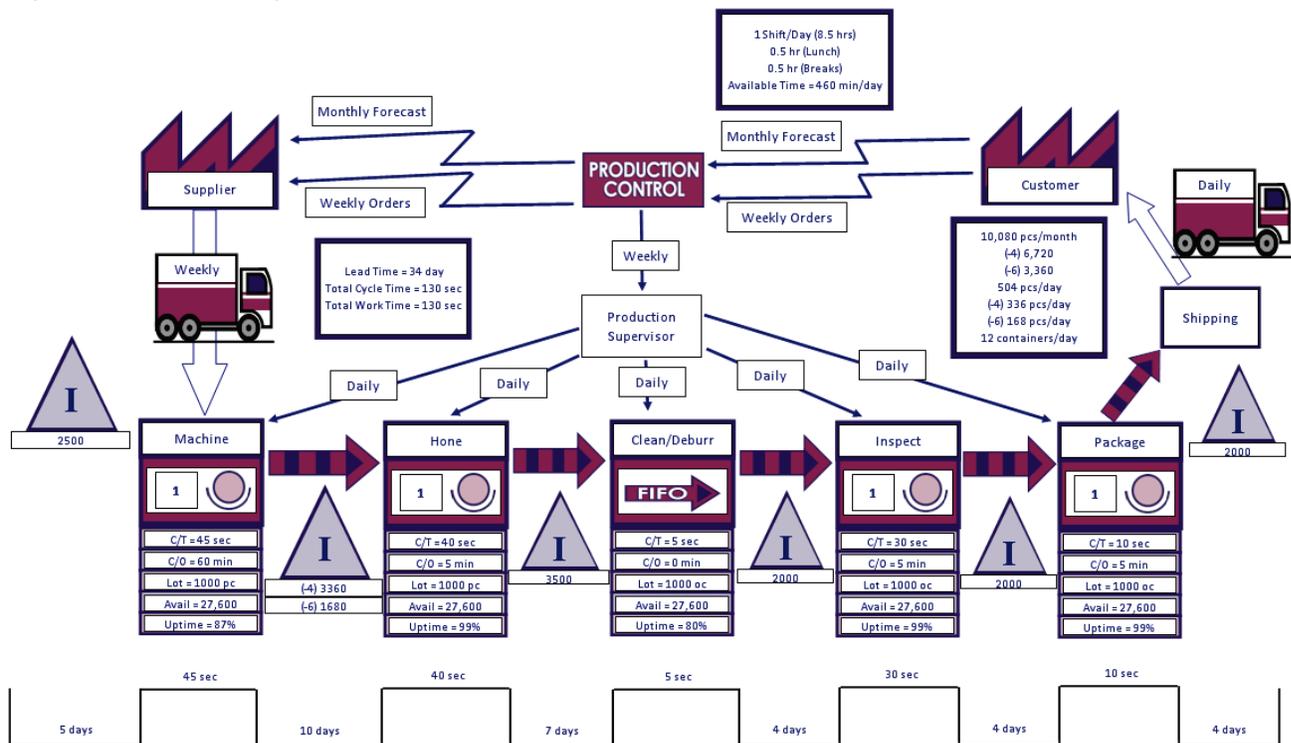
Notes:

1. Copy and Paste the above symbols into a blank worksheet as needed to create the value stream map. Resize if necessary.

2. Recommended resource for Value Stream Mapping: Quarterman Lee and Brad Snyder, "The Strategos Guide to Value Stream & Process Mapping", Enna Products Corp., 2006.

This is a present state value stream map of a manufacturing operation. This example is courtesy and copyright of Strategos, Inc. www.strategosinc.com.

Example Present State Value Stream Map



Basic Graphical Templates – Pareto Example

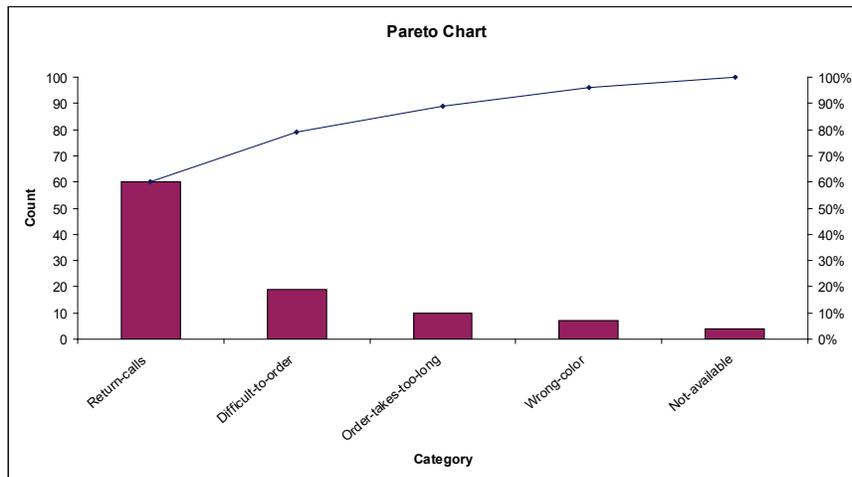
Click **SigmaXL > Templates & Calculators > Basic Graphical Templates > Pareto Chart**. Enter data as shown:

A	B	C	D	E
Category	Count			
Difficult-to-order	19			
Not-available	4			
Order-takes-too-long	10			
Return-calls	60			
Wrong-color	7			



Pareto Chart

Click the **Pareto Chart** button to produce the Pareto Chart:



Notes:

1. This Pareto Chart template should be used with count data like number of errors. You can also create Pareto charts with cost data.
2. You can replace the **Category** and **Count** column headings with any headings that you wish.
3. Enter the Pareto categories in the **Category** column. These can be Name, Location, Error Type or other text information. Pareto categories are required and will appear on the horizontal X-Axis of the Pareto Chart.
4. Enter your count (or cost) data in the **Count** column.
5. Click the **Pareto Chart** button to create a Pareto Chart.

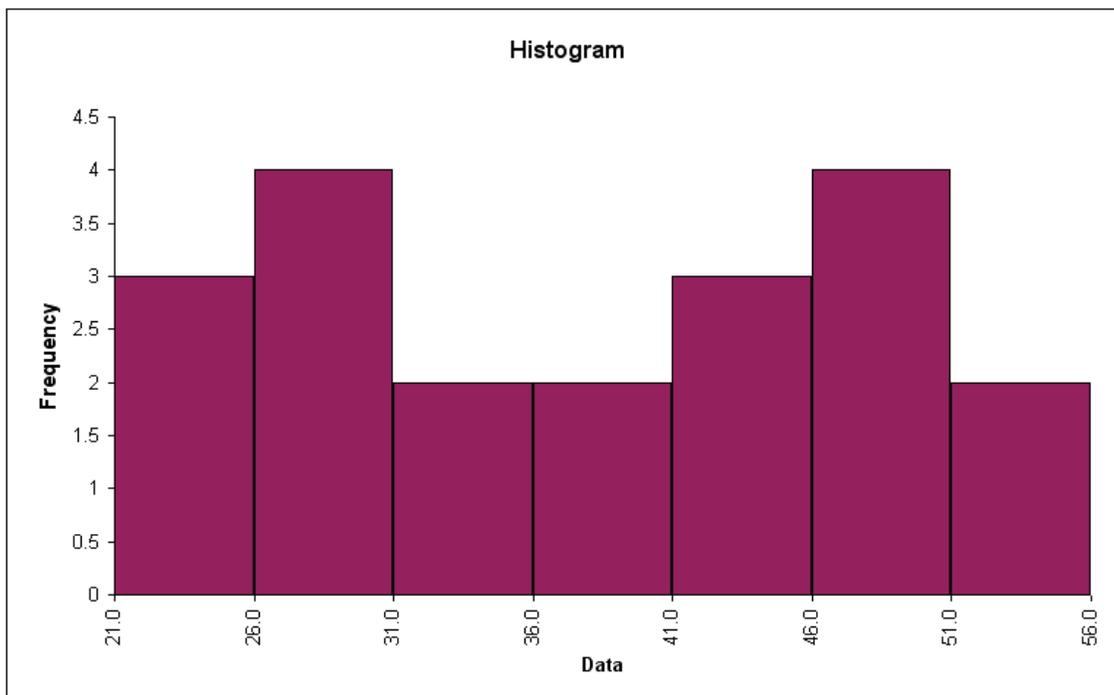
Basic Graphical Templates – Histogram Example

Click **SigmaXL > Templates & Calculators > Basic Graphical Templates > Histogram**. Enter data as shown:

Data
48
47
56
43
43
48
39
50
46
52
29
27
31
28
32
39
34
23
23
21

Histogram

Click the **Histogram** button to produce the Histogram:



Notes:

1. This Histogram template should be used with continuous data like cycle time.
2. You can replace the **Data** column heading with any heading that you wish.
3. Enter your data in the **Data** column.
4. Click the **Histogram** button to create a Histogram.

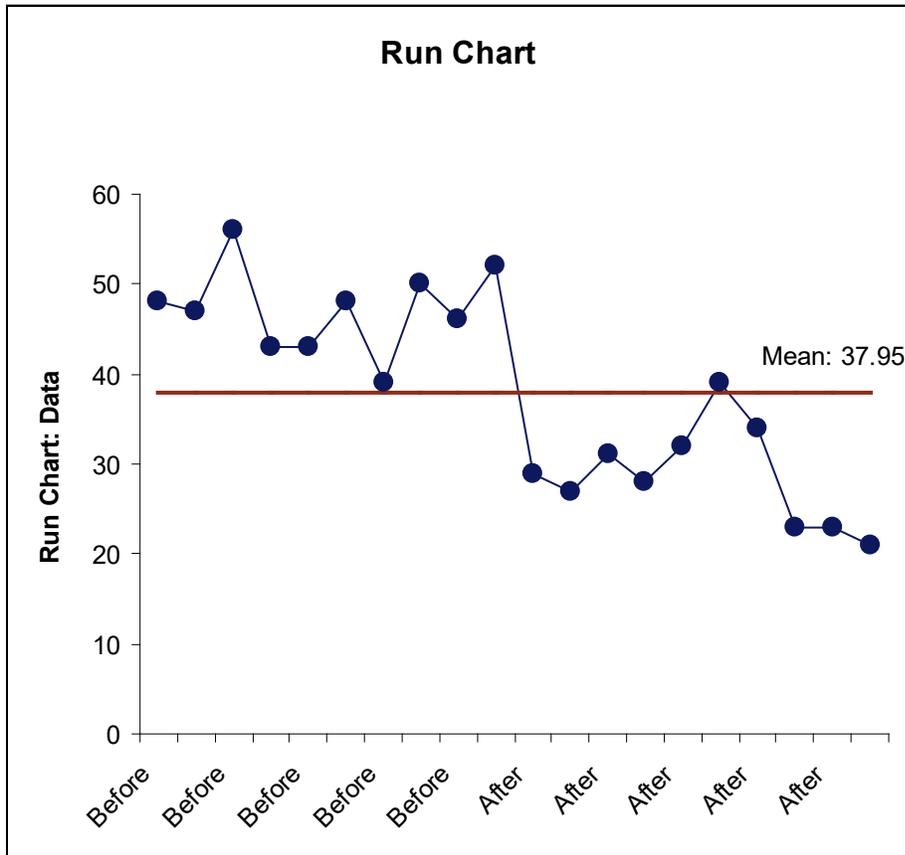
Basic Graphical Templates – Run Chart Example

Click **SigmaXL > Templates & Calculators > Basic Graphical Templates > Run Chart**. Enter the data as shown (copy and paste the histogram data):

Before/After	Data
Before	48
Before	47
Before	56
Before	43
Before	43
Before	48
Before	39
Before	50
Before	46
Before	52
After	29
After	27
After	31
After	28
After	32
After	39
After	34
After	23
After	23
After	21

Run Chart

Click the **Run Chart** button to produce the Run Chart:



Notes:

1. This Run Chart template can be used with a variety of data types: continuous data like cycle time, count data like number of errors, or cost data. The data must be in chronological time-sequence order.
2. You can replace the **X-Axis Label** and **Data** column headings with any headings that you wish.
3. Enter your data in the **Data** column.
4. Enter labels in **X-Axis Label** column. Labels can be Date, Time, Name, or other text information. These labels are optional and will appear on the horizontal X-Axis of the Run Chart.
5. Click the **Run Chart** button to create a Run Chart.

Basic Statistical Templates – Sample Size – Discrete

Click **SigmaXL > Templates & Calculators > Basic Statistical Templates > Sample Size – Discrete** to access the Sample Size – Discrete calculator. The template gives the following default example:

 Sample Size Calculator - Discrete Data		
Sample Data (user inputs):		
Estimate of Proportion	P	0.5
Desired margin of error	delta / half-interval	0.03
Population Size (optional)	N	
Confidence level (enter .95 for 95%)	100*(1-α)%	95.0%
Results:		
Minimum Sample Size	n	1068
	n (adjusted for small N)	
	np check (should be ≥ 5)	534

Notes for use of the Sample Size – Discrete Calculator:

1. P is estimate of proportion for outcome of interest. Use P = 0.5 if unknown.
2. Delta is desired proportion margin of error. Enter as the half-width, i.e. if the desired margin of error is +/- 3%, enter 0.03.
3. Enter population size N to adjust for small populations (N < 10000).
4. np should be ≥ 5 . If necessary, reduce delta to adjust.
5. Power (1 – Beta) is not considered in these calculations. Power and Sample Size may be calculated using **SigmaXL > Statistical Tools > Power & Sample Size Calculators**.

This sample size calculator is based on a confidence interval approach, where you enter the desired half-interval value (delta). Let's say that you desire the proportion margin of error to be +/- 3%. As noted above, you simply enter a delta (half-interval) value of .03.

The challenge with this sample size calculator is that we need an estimate of the population proportion. This is a bit of a chicken and egg situation – which comes first? I want to calculate an appropriate sample size to determine an outcome, but I am being asked to enter an estimate of the population proportion. First, keep in mind that this tool is a planning tool – the true confidence intervals will be determined after you collect your data. Estimating the population proportion can be done if you have good historical data to draw from, for example, in historical customer surveys the percentage of satisfied customers was 80%. In this case, you could use P=0.8. If you do not have *a priori* knowledge, then leave P=0.5 which gives the most conservative value (i.e., largest estimate

of sample size). If you enter a value other than 0.5 it will result in a smaller sample size requirement.

Basic Statistical Templates – Sample Size – Continuous

Click **SigmaXL > Templates & Calculators > Basic Statistical Templates > Sample Size – Continuous** to access the Sample Size – Continuous calculator. The template gives the following default example:

Sample Data (user inputs):		
Estimate of Standard Deviation	S	1
Desired margin of error	delta / half-interval	0.25
Population Size (optional)	N	
Confidence level (enter .95 for 95%)	$100^*(1-\alpha)\%$	95.0%
Results:		
Minimum Sample Size	n	62
	n (adjusted for small N)	

Notes for use of the Sample Size – Continuous Calculator:

1. Delta uses the same units as the standard deviation. Enter as the half-width, i.e., if the desired margin of error is +/- 0.25, enter 0.25.
2. Enter (optional) population size N to adjust for small populations ($N < 1000$).
3. Power ($1 - \text{Beta}$) is not considered in these calculations. Power and Sample Size may be calculated using **SigmaXL > Statistical Tools > Power & Sample Size Calculators**.

Similar to the Sample Size – Discrete calculator, this sample size calculator is based on a confidence interval approach, where you enter the desired half-interval value (delta). If a survey had responses that were on a continuous scale of 1 to 5, and you desired a margin of error on the mean to be +/- .25, then you would use .25 as the delta value. Note that in this continuous case the units are not percentages but level of satisfaction.

Here we need an *a priori* estimate of the population standard deviation. If you have no idea what the standard deviation will be, then you could take a small sample to get a rough estimate of the standard deviation.

Basic Statistical Templates – Minimum Sample Size for Robust t-Tests and ANOVA

It is well known that the central limit theorem enables the t-Test and ANOVA to be fairly robust to the assumption of normality. A question that invariably arises is, “How large does the sample size have to be?” A popular rule of thumb answer for the one sample t-Test is “ $n = 30$.” While this rule of thumb often does work well, the sample size may be too large or too small depending on the degree of nonnormality as measured by the Skewness and Kurtosis. Furthermore, it is not applicable to a One Sided t-Test, 2 Sample t-Test or One-Way ANOVA.

To address this issue, we have developed a unique template that gives a minimum sample size needed for a hypothesis test to be robust.

Click **SigmaXL > Templates & Calculators > Basic Statistical Templates > Minimum Sample Size for Robust t-Tests and ANOVA** to access this template. It includes minimum sample size for robustness for the 1 Sample t-Test, 2 Sample t-Test and the One-Way ANOVA.

The user may specify the alternative hypothesis as “Less Than” (one sided), “Not Equal To” (two sided) or “Greater Than” (one sided). Confidence levels of 90% ($\alpha = 0.1$), 95% ($\alpha = .05$) or 99% ($\alpha = .01$) may also be specified:

 Minimum Sample Size for Robust Hypothesis Testing		
Sample Data (user inputs):		
Hypothesis Test:		1 Sample t-Test
Alternative Hypothesis :	Ha	Not Equal To
Confidence Level:	100*(1- α)%	95%
Skewness:	Skew	1
Kurtosis:	Kurt	-0.48
Results:		
Minimum sample size for each sample/group:	n	30

To use the template, simply select the appropriate **Hypothesis Test**, **Alternative Hypothesis** and **Confidence Level** using the drop-down selection. Enter **Skewness** and **Kurtosis** values as shown in the yellow highlighted cells.

Note that in the example shown, the rule of thumb for a 1 Sample t-Test “ $n = 30$ ” is confirmed with a moderate skew value of 1.

Now change the **Alternative Hypothesis** to “Less Than” and **Confidence Level** to 99% as shown:

Sample Data (user inputs):		
Hypothesis Test:		1 Sample t-Test
Alternative Hypothesis :	Ha	Less Than
Confidence Level:	100*(1- α)%	99%
Skewness:	Skew	1
Kurtosis:	Kurt	-0.48
Results:		
Minimum sample size for each sample/group:	n	752

The minimum sample size required for robustness is now 752!

On the other hand, if you want to perform a standard One-Way ANOVA, enter the values as shown:

Sample Data (user inputs):		
Hypothesis Test:		3 Sample One Way ANOVA
Alternative Hypothesis :	Ha	Not Equal To
Confidence Level:	100*(1- α)%	95%
Skewness:	Skew	1
Kurtosis:	Kurt	-0.48
Results:		
Minimum sample size for each sample/group:	n	3

Now the minimum sample size requirement is only 3. This value applies to each sample or group, so for the 3 Sample ANOVA that would mean each sample has $n = 3$ for a total number of observations = 9.

Note that this calculator is strictly addressing the question of alpha robustness to nonnormality. Power is not considered here.

If the minimum sample size requirements cannot be met, you should use a nonparametric equivalent to the parametric hypothesis test (i.e. One Sample Sign or Wilcoxon, Two Sample Mann-Whitney, Kruskal-Wallis or Mood’s Median: SigmaXL > Statistical Tools > Nonparametric Tests).

Skewness and Kurtosis

Sample Skewness and Kurtosis values can be obtained from SigmaXL’s descriptive statistics:
SigmaXL > Statistical Tools > Descriptive Statistics.

A slight Skew is +/- 0.5, moderate Skew is +/- 1, severe Skew is +/- 2 and extreme Skew is +/- 5. The Skewness range used should be -5 to +5. Values beyond this range are extrapolated so may be inaccurate.

Kurtosis for a normal distribution is 0. Kurtosis must be greater than or equal to:
(Skew ^ 2 - 1.48).

Kurtosis “delta” is Kurt - Skew² and is used in the regression equation. Kurtosis delta range should be -1.48 to +1.48. A kurtosis delta less than -1.48 denotes a bimodal distribution so this is a lower boundary. Values above +1.48 are extrapolated so may be inaccurate.

The calculator assumes that all samples have the same Skewness and Kurtosis.

Monte Carlo Simulation

The data for minimum sample size formulas are derived from extensive Monte Carlo simulations for $n = 2$ to 2000. Observed alpha values were determined empirically from the P-Values of 100,000 replicate hypothesis tests for each n . Nonnormal data with Skew = -5 to +5 and Kurt delta = -1.48 to +1.48 was generated using the Pearson Family function (see SigmaXL DiscoverSim Workbook: Appendix for details).

Minimum sample size for robustness occurs when the simulated observed alpha is within +/- 20% of the specified alpha (Bradley 1980 and Rhiel 1996).

Regression Models

A separate regression model was constructed for each hypothesis test, alternative hypothesis, and confidence level (total of 21 models), using coding as shown.

	Coded for Regression	Notes for Regression Coding
1 Sample t-Test	1	1=1 Sample t; 2=2 Sample t; 3=3 Sample ANOVA
Not Equal To	2	1=LT for +ve Skew or GT for -ve Skew; 2=NE; 3=GT for +ve Skew or LT for -ve Skew
95%	2	1=90% ($\alpha=0.1$); 2=95% ($\alpha=0.05$); 3=99% ($\alpha=0.01$)
1	1	Skew*2
-0.48	-1.48	Kurt_delta = Kurt - Skew*2. Minimum = -1.48.

The model predictor terms are Skewness² and Kurtosis Delta. The response is n minimum. The term coefficient values are stored in the template as shown.

Hypothesis Test:	1	1	1	1	1
Alternative Hypothesis :	1	1	1	2	2
Confidence Level:	1	2	3	1	2
Description:	1 Sample t; Ha = LT +, GT -; $\alpha=0.1$	1 Sample t; Ha = LT +, GT -; $\alpha=0.05$	1 Sample t; Ha = LT +, GT -; $\alpha=0.01$	1 Sample t; Ha = NE; $\alpha=0.1$	1 Sample t; Ha = NE; $\alpha=0.05$

Predictor Term	Coefficient	Coefficient	Coefficient	Coefficient	Coefficient
Constant:	2.682	-2.940	-257.06	-0.658037	-0.864317
Skew^2:	44.257	137.32	1008.2	9.440	22.848
Kurt Delta:	-2.431	-10.157	-0.564855	-1.356	-5.280

Regression model n:	51	149	752	11	30
Selected n:	0	0	0	0	30
Model R-Square:	99.48%	99.34%	99.93%	99.81%	99.85%

Model R-Square values are typically over 99%, with some exceptions (96%) due to small estimated sample sizes.

The model used is the one which matches the selected hypothesis test, alternative hypothesis, and confidence level. The selected n is highlighted in red and is also displayed in the template Results cell. If $n > 2000$, a text display "> 2000" is shown in the Results cell.

References

Boos, D.D. and Hughes-Oliver, J.M. (2000), "How Large Does n Have to be for Z and t Intervals?" *The American Statistician*, 54(2), 121-128.

Bradley, J. V. (1980), "Nonrobustness in Z, t, and F Tests at Large Sample Sizes," *Bulletin of the Psychonomics Society*, 16(5), 333-336.

Rhiel, G. S., and Chaffin, W. W. (1996), "An Investigation of the Large-Sample/Small-Sample Approach to the One-Sample Test for a Mean (Sigma Unknown)," *Journal of Statistics Education*, 4, No. 3 (www.amstat.org/publications/jse).

SigmaXL, Inc., DiscoverSim Version 1.1 Workbook.

Basic Statistical Templates – 1 Sample Z-Test and Confidence Interval for Mean

Click **SigmaXL > Templates & Calculators > Basic Statistical Templates > 1 Sample Z-Test and Confidence Interval for Mean** to access the 1 Sample Z-Test calculator. The template gives the following default example.

Sample Data (user inputs):		
Sample Size	n	30
Sample Mean	x-bar	1
Known Population Standard Deviation	σ	1
Null Hypothesis (hypothesized mean)	H_0 : Mean (μ) =	0.5
Alternative Hypothesis	H_a : Mean (μ)	Not Equal To
Confidence Level (enter .95 for 95%)	100*(1- α)%	95.0%
Results:		
Standard Error Mean		0.1826
DF		29
Z-statistic		2.7386
alpha		0.0500
P-Value (2 sided)		0.0062
Upper Confidence Limit (2 Sided)		1.3578
Lower Confidence Limit (2 Sided)		0.6422
Test Information		
Null Hypothesis H_0 : Mean (μ) = 0.5		Reject
Alternative Hypothesis H_a : Mean (μ) \neq 0.5		Conclude true at 95.0% confidence level
Interpretation of P-Value and Confidence Intervals		
The P-Value is the two-sided (two-tail) probability of obtaining a Z-statistic at least as extreme as 2.7386, given that the null hypothesis is true. Since the P-Value is less than alpha (0.05), we reject the null hypothesis: Mean (μ) = 0.5 and conclude that the alternative hypothesis H_a : Mean (μ) \neq 0.5 is true at the confidence level of 95.0%.		
We are 95% confident that the true population mean lies within the interval (0.6422 to 1.3578).		

Notes:

1. Enter summarized Sample Data, Known Population Standard Deviation, Null Hypothesis and Confidence Level in cells with yellow highlight. Do not modify any other part of this worksheet.
2. Select Alternative Hypothesis using drop-down.
3. This calculator will automatically detect the version of Excel and use the appropriate statistical functions. If Excel 2010 or higher, then the Excel 2010 statistical functions are used.

Basic Statistical Templates – 1 Sample t-Test and Confidence Interval for Mean

Click **SigmaXL > Templates & Calculators > Basic Statistical Templates > 1 Sample t-Test and Confidence Interval for Mean** to access the 1 Sample t-Test calculator. The template gives the following default example.

Sample Data (user inputs):		
Sample Size	n	30
Sample Mean	x-bar	1
Sample Standard Deviation	s	1
Null Hypothesis (hypothesized mean)	$H_0: \text{Mean } (\mu) =$	0.5
Alternative Hypothesis	$H_a: \text{Mean } (\mu)$	Not Equal To
Confidence Level (enter .95 for 95%)	$100*(1-\alpha)\%$	95.0%
Results:		
Standard Error Mean		0.1826
DF		29
t-statistic		2.7386
alpha		0.0500
P-Value (2 sided)		0.0104
Upper Confidence Limit (2 Sided)		1.3734
Lower Confidence Limit (2 Sided)		0.6266
Test Information		
Null Hypothesis $H_0: \text{Mean } (\mu) = 0.5$		Reject
Alternative Hypothesis $H_a: \text{Mean } (\mu) \neq 0.5$		Conclude true at 95.0% confidence level
Interpretation of P-Value and Confidence Intervals		
The P-Value is the two-sided (two-tail) probability of obtaining a t-statistic at least as extreme as 2.7386, given that the null hypothesis is true. Since the P-Value is less than alpha (0.05), we reject the null hypothesis: Mean (μ) = 0.5 and conclude that the alternative hypothesis $H_a: \text{Mean } (\mu) \neq 0.5$ is true at the confidence level of 95.0%.		
We are 95% confident that the true population mean lies within the interval (0.6266 to 1.3734).		

Notes:

1. Enter summarized Sample Data, Null Hypothesis and Confidence Level in cells with yellow highlight. Do not modify any other part of this worksheet.
2. Select Alternative Hypothesis using drop-down.
3. This calculator will automatically detect the version of Excel and use the appropriate statistical functions. If Excel 2010 or higher, then the Excel 2010 statistical functions are used.

Basic Statistical Templates – 1 Sample Equivalence Test for Mean

Click **SigmaXL > Templates & Calculators > Basic Statistical Templates > 1 Sample Equivalence Test for Mean** to access the 1 Sample Equivalence Test (Two One-Sided t-Tests TOST) calculator. The template gives the following default example.

SigmaXL 1 Sample Equivalence Test for Mean (Two One-Sided t-Tests TOST)			
Sample Data (user inputs):			
Sample Size	n		30
Sample Mean	x-bar		0
Sample Standard Deviation	s		1
Target	T		0
Upper Equivalence Limit	UEL		0.5
Lower Equivalence Limit	LEL		-0.5
Confidence Level (enter .95 for 95%)	100*(1- α)%		95.0%
Results:			
Standard Error Mean			0.1826
DF			29
Difference (Mean - Target)			0
alpha			0.0500
t1-statistic (upper)			-2.7386
t2-statistic (lower)			2.7386
P1-Value (upper)			0.0052
P2-Value (lower)			0.0052
Equivalence P-Value (maximum of P1 and P2)			0.0052
Upper Confidence Limit Difference			0.3102
Lower Confidence Limit Difference			-0.3102

Interpretation of Equivalence P-Value and Confidence Intervals
Since the Equivalence P-Value is less than alpha (0.05), conclude that equivalence is true.
Since the 95% confidence interval is within the equivalence interval, conclude that equivalence is true.

Notes:

1. Enter summarized Sample Data, Target, Upper & Lower Equivalence Limits and Confidence Level in cells with yellow highlight. Do not modify any other part of this worksheet.
2. Null hypothesis for P1: Difference \geq UEL; Null hypothesis for P2: Difference \leq LEL. Both null hypotheses must be rejected to conclude that equivalence is true.
3. LEL and UEL establish the zone or region of equivalence and are determined by what size difference is considered practically significant.
4. This calculator will automatically detect the version of Excel and use the appropriate statistical functions. If Excel 2010 or higher, then the Excel 2010 statistical functions are used.

Basic Statistical Templates – 2 Sample Equivalence Test (Compare 2 Means)

Click **SigmaXL > Templates & Calculators > Basic Statistical Templates > 2 Sample Equivalence Test (Compare 2 Means)** to access the 2 Sample Equivalence Test (Two One-Sided t-Tests TOST) calculator. The template gives the following default example.

2 Sample Equivalence Test (Two One-Sided t-Tests TOST)			
Sample Data (user inputs):		Test/Treatment Sample	Reference/Control Sample
Sample Size	n	30	30
Sample Mean	x-bar	0.0000	0.5000
Sample Standard Deviation	s	1.0000	1.0000
Upper Equivalence Limit	UEL	1.0	
Lower Equivalence Limit	LEL	-1.0	
Confidence Level (enter .95 for 95%)	100*(1- α)%	95.0%	
Assume Equal Variances		Yes	
Results:			
Sample Mean Difference		-0.5000	
Standard Error Difference (using pooled variance)		0.2582	
df (equal variance)		58	
alpha		0.0500	
t1-statistic (upper)		-5.8095	
t2-statistic (lower)		1.9365	
P1-Value (upper)		0.0000	
P2-Value (lower)		0.0288	
Equivalence P-Value (maximum of P1 and P2)		0.0288	
Upper Confidence Limit Mean Difference		-0.0684	
Lower Confidence Limit Mean Difference		-0.9316	

Interpretation of Equivalence P-Value and Confidence Intervals
Since the Equivalence P-Value is less than alpha (0.05), conclude that equivalence is true.
Since the 95% confidence interval is within the equivalence interval, conclude that equivalence is true.

Notes:

1. Enter summarized Sample Data, Upper & Lower Equivalence Limits and Confidence Level in cells with yellow highlight. Do not modify any other part of this worksheet.
2. Select "Assume Equal Variances" using drop-down.
3. Null hypothesis for P1: Mean Difference \geq UEL; Null hypothesis for P2: Mean Difference \leq LEL. Both null hypotheses must be rejected to conclude that equivalence is true.
4. LEL and UEL establish the zone or region of equivalence and are determined by what size mean difference is considered practically significant.
5. This calculator will automatically detect the version of Excel and use the appropriate statistical functions. If Excel 2010 or higher, then the Excel 2010 statistical functions are used.

Basic Statistical Templates – 1 Sample Chi-Square Test and CI for Standard Deviation

Click **SigmaXL > Templates & Calculators > Basic Statistical Templates > 1 Sample Chi-Square Test and CI for Standard Deviation** to access the 1 Sample Chi-Square-Test for Standard Deviation calculator. The template gives the following default example.

1 Sample Standard Deviation Chi-Square Test and Confidence Interval		
Sample Data (user inputs):		
Sample Size	n	10
Sample Standard Deviation	s	1
Null Hypothesis (hypothesized StDev)	H ₀ : Sigma (σ) =	0.5
Alternative Hypothesis	H _a : Sigma (σ)	Not Equal To
Confidence Level (enter .95 for 95%)	100*(1-α)%	95.0%
Results:		
Sample Variance		1.0000
DF		9
Chi-Square test statistic		36.0000
alpha		0.0500
P-Value (2 sided)		0.0001
Upper Confidence Limit (2 Sided)		1.8256
Lower Confidence Limit (2 Sided)		0.6878
Test Information		
Null Hypothesis H ₀ : Mean Difference (μ1 - μ2) = 0		Reject
Alternative Hypothesis H _a : Mean Difference (μ1 - μ2) ≠ 0		Conclude true at 95.0% confidence level
Interpretation of P-Value and Confidence Intervals		
The P-Value is the two-sided (two-tail) probability of obtaining a test statistic at least as extreme as 36.0000, given that the null hypothesis is true. Since the P-Value is less than alpha (0.05), we reject the null hypothesis: Sigma (σ) = 0.5 and conclude that the alternative hypothesis H _a : Sigma (σ) ≠ 0.5 is true at the confidence level of 95.0%.		
We are 95% confident that the true population standard deviation lies within the interval (0.6878 to 1.8256).		

Notes:

1. Enter summarized Sample Data, Null Hypothesis and Confidence Level in cells with yellow highlight. Do not modify any other part of this worksheet.
2. Select Alternative Hypothesis using drop-down.
3. This calculator will automatically detect the version of Excel and use the appropriate statistical functions. If Excel 2010 or higher, then the Excel 2010 statistical functions are used.

Basic Statistical Templates – 1 Proportion Test and Confidence Interval

Click **SigmaXL > Templates & Calculators > Basic Statistical Templates > 1 Proportion Test and Confidence Interval** to access the 1 Proportion Test calculator. The template gives the following default example.

1 Proportion Test and Confidence Interval		
Sample Data (user inputs):		
Number of Events	x	1
Sample Size	n	10
Null Hypothesis (hypothesized proportion)	H ₀ : Proportion =	0.5
Alternative Hypothesis	H _a : Proportion	Not Equal To
Confidence Level (enter .95 for 95%)	100*(1-α)%	95.0%
Hypothesis Test Method		Binomial Exact
Confidence Interval Method		Exact (Clopper-Pearson Beta)
Results:		
Sample proportion (x/n)		0.1000
alpha		0.0500
npq (npq should be >= 5 for normal approximation; q = 1-p)		0.9000
Z-statistic (normal)		-2.5298
Binomial exact probability P-Value (2-sided)		0.0215
Upper Confidence Limit (2-sided)		0.4450
Lower Confidence Limit (2-sided)		0.0025

Test Information	
Null Hypothesis H ₀ : Proportion (p) = 0.5	Reject
Alternative Hypothesis H _a : Proportion (p) ≠ 0.5	Conclude true at 95.0% confidence level

Interpretation of P-Value and Confidence Intervals
The P-Value is the two-sided (two-tail) probability of observing a sample proportion at least as extreme as 0.1000, given that the null hypothesis is true. Since the P-Value is less than alpha (0.05), we reject the null hypothesis: Proportion (p) = 0.5 and conclude that the alternative hypothesis: Proportion (p) ≠ 0.5 is true at the confidence level of 95.0%.
We are 95% confident that the true population proportion lies within the interval (0.0025 to 0.4450).

Notes:

1. Enter summarized Sample Data, Null Hypothesis and Confidence Level in cells with yellow highlight. Do not modify any other part of this worksheet.
2. Select Alternative Hypothesis and Methods using drop-down. If x = 0, it is recommended that the Alternative Hypothesis be set to one-sided "Less Than." If x = n, set the Alternative Hypothesis to one-sided "Greater Than."

3. Hypothesis Test Method:

- Binomial Exact is recommended.
- Use Normal Approximation only when n is large and proportion is close to 0.5 (check that $np \geq 5$ and $nq \geq 5$, or the more conservative npq , shown above, ≥ 5). See Rosner (2010).

4. Confidence Interval Method (recommendations based on literature review, see references below):

- Confidence intervals for binomial proportions have an "oscillation" phenomenon where the coverage probability varies with n and p .
- Exact (Clopper-Pearson) is strictly conservative and will guarantee the specified confidence level as a minimum coverage probability, but results in wide intervals. Recommended only for applications requiring strictly conservative intervals.
- Wilson Score and Jeffreys Beta have mean coverage probability matching the specified confidence interval. We recommend Wilson for two-sided confidence intervals and Jeffreys for one-sided.
- Normal Approximation (Wald) is not recommended, but included here to validate hand calculations. Use only when n is large and proportion is close to 0.5 (check that npq , shown above, ≥ 5).

5. This calculator will automatically detect the version of Excel and use the appropriate statistical functions. If Excel 2010 or higher, then the Excel 2010 statistical functions are used.

REFERENCES

1. Agresti, A. and Coull, B.A. (1998), Approximate is Better than "Exact" for Interval Estimation of Binomial Proportions. *The American Statistician*, 52,119–126.
2. Brown, L. D., Cai, T. T. and DasGupta, A. (2001), Interval estimation for a binomial proportion. *Statistical Science*, 16, 101-133. With comments and a rejoinder by the authors.
3. Cai, T. T. (2005), One-sided confidence intervals in discrete distributions. *J. Statist. Plann. Inference*, 131, 63-88.
4. Clopper, C.J., and Pearson, E.S. (1934), The Use of Confidence or Fiducial Limits Illustrated in the Case of the Binomial. *Biometrika*, 26, 404–413.
5. Newcombe, R. (1998a), Two-sided confidence intervals for the single proportion: Comparison of seven methods, *Statistics in Medicine*, 17, 857-872.
6. Rosner, B. (2010), *Fundamentals of Biostatistics*. Seventh Edition. Equation 5.15, page 133. Duxbury Press.
7. Vollset, S.E. (1993). Confidence intervals for a binomial proportion, *Statistics in Medicine*, 12, 809–824.

Basic Statistical Templates – 2 Proportions Test and Confidence Interval

Click **SigmaXL > Templates & Calculators > Basic Statistical Templates > 2 Proportions Test and Confidence Interval** to access the 2 Proportions Test calculator. The template gives the following default example.

2 Proportions Test and Confidence Interval			
Sample Data (user inputs):		Sample 1	Sample 2
Number of Events	x	1	2
Sample Size	n	10	10
Null Hypothesis (hypothesized difference)	$H_0: P_1 - P_2 =$	0	
Alternative Hypothesis	$H_a: P_1 - P_2$	Not Equal To	
Confidence Level (enter .95 for 95%)	$100*(1-\alpha)\%$	95.0%	
Hypothesis Test Method		Fisher's Exact	
Confidence Interval Method		Newcombe-Wilson Score	
Results:			
Sample proportion (x/n)		0.1000	0.2000
Sample proportion difference		-0.1000	
alpha		0.0500	
Minimum expected value (should be ≥ 5 for normal approximation)		1.5000	
Fisher's Exact probability P-Value (2-sided)		1.0000	
Upper Confidence Limit (2-sided)		0.2362	
Lower Confidence Limit (2-sided)		-0.4205	

Test Information	
Null Hypothesis $H_0: P_1 - P_2 = 0$	Fail to Reject
Alternative Hypothesis $H_a: P_1 - P_2 \neq 0$	

Interpretation of P-Value and Confidence Intervals
The P-Value is the two-sided (two-tail) probability of observing a difference in proportions at least as extreme as -0.1000, given that the null hypothesis is true. Since the P-Value is greater than alpha (0.05), we fail to reject the null hypothesis $H_0: P_1 - P_2 = 0$.
We are 95% confident that the true difference in population proportions lies within the interval (-0.4205 to 0.2362).

Notes:

1. Enter summarized Sample Data, Null Hypothesis and Confidence Level in cells with yellow highlight. Do not modify any other part of this worksheet.
2. Select Alternative Hypothesis and Methods using drop-down.

3. Hypothesis Test Method:

- Fisher's Exact is recommended, but is only valid for $H_0: P_1 - P_2 = 0$.
- Use Normal Approximation only when minimum expected cell value, shown above, is ≥ 5 . Unpooled is used if $H_0: P_1 - P_2 \neq 0$. Use pooled if $H_0: P_1 - P_2 = 0$. (See Peltier).

4. Confidence Interval Method:

- Confidence intervals for difference in binomial proportions have an "oscillation" phenomenon where the coverage probability varies with n and p .
- Newcombe-Wilson Score is recommended and has a mean coverage probability that is close to the specified confidence interval. (See Newcombe method 10).
- Newcombe-Wilson Score (CC = Continuity Corrected) is conservative and will typically meet the specified confidence level as a minimum coverage probability, but results in wider intervals. (See Newcombe method 11).
- Jeffreys-Perks is similar to Newcombe-Wilson, and included as an option because it is the preferred interval for some practitioners. (See Radhakrishna; Newcombe method 4).
- Normal Approximation is not recommended, but included here to validate hand calculations. Use only when minimum expected cell value, shown above, is ≥ 5 .
- Exact methods that are strictly conservative (like Clopper-Pearson for the one proportion case) are not included in this template because they are computationally intensive and slow.

5. Due to the complexity of calculations, this template uses vba macros rather than Excel formulas. SigmaXL must be initialized and appear on the menu in order for this template to function.

REFERENCES

1. Beal, S. L. (1987), Asymptotic confidence intervals for the difference between two binomial parameters for use with small samples. *Biometrics*, 43, 941-950.
2. Newcombe, R. G. (1998b), Interval estimation for the difference between independent proportions: Comparison of eleven methods. *Statistics in Medicine*, 17:873-890.
3. Peltier, C., Why Do We Pool for the Two-Proportion z-Test?
http://apcentral.collegeboard.com/apc/members/courses/teachers_corner/49013.html.
4. Radhakrishna S., Murthy B.N., Nair N.G., Jayabal P., Jayasri R. (1992), Confidence intervals in medical research. *Indian J Med Res.*, Jun;96:199-205.

Basic Statistical Templates – 2 Proportions Equivalence Test

Click **SigmaXL > Templates & Calculators > Basic Statistical Templates > 2 Proportions Equivalence Test** to access the 2 Proportions Equivalence Test (Two One-Sided Tests TOST) calculator. The template gives the following default example.

2 Proportions Equivalence Test (Two One-Sided Tests TOST)			
Sample Data (user inputs):		Test/Treatment Sample	Reference/Control Sample
Number of Events	x	10	11
Sample Size	n	100	100
Upper Equivalence Limit	UEL	0.1	
Lower Equivalence Limit	LEL	-0.1	
Confidence Level (enter .95 for 95%)	100*(1- α)%	95.0%	
Confidence Interval Method		Newcombe-Wilson Score	
Results:			
Sample proportion (x/n)		0.1000	0.1100
Sample proportion difference		-0.0100	
alpha		0.0500	
Minimum expected value (should be ≥ 5 for normal approximation)		10.5000	
Z1-statistic (upper, normal unpooled)		-2.5376	
Z2-statistic (lower, normal unpooled)		2.0762	
P1-Value (upper)		0.0056	
P2-Value (lower)		0.0189	
Equivalence P-Value (maximum of P1 and P2)		0.0189	
Upper Confidence Limit Proportion Difference		0.0633	
Lower Confidence Limit Proportion Difference		-0.0835	

Interpretation of Equivalence P-Value and Confidence Intervals
Since the Equivalence P-Value is less than alpha (0.05), conclude that equivalence is true.
Since the 95% confidence interval is within the equivalence interval, conclude that equivalence is true.

Notes:

- Enter summarized Sample Data, Upper & Lower Equivalence Limits and Confidence Level in cells with yellow highlight. Do not modify any other part of this worksheet.
- Select Confidence Interval Method using drop-down.
 - Confidence intervals for difference in binomial proportions have an "oscillation" phenomenon where the coverage probability varies with n and p.
 - Newcombe-Wilson Score is recommended and has a mean coverage probability that is close to the specified confidence interval. (See Newcombe method 10).
 - Newcombe-Wilson Score (CC = Continuity Corrected) is conservative and will typically meet the specified confidence level as a minimum coverage probability, but results in wider intervals. (See Newcombe method 11).
 - Jeffreys-Perks is similar to Newcombe-Wilson, and included as an option because it is the preferred interval for some practitioners. (See Radhakrishna; Newcombe method 4).
 - Normal Approximation will match the P-Value calculations, but should only be used when minimum expected cell value, shown above, is ≥ 5 .

- Exact methods that are strictly conservative (like Clopper-Pearson for the one proportion case) are not included in this template because they are computationally intensive and slow.
3. Null hypothesis for P1: Proportion Difference \geq UEL; Null hypothesis for P2: Proportion Difference \leq LEL. Both null hypotheses must be rejected to conclude that equivalence is true. The P-Values are based on the normal approximation unpooled method.
 4. LEL and UEL establish the zone or region of equivalence and are determined by what size proportion difference is considered practically significant.
 5. Since the confidence interval options use different methods than the hypothesis test (except for Normal option), it is possible that the conclusion from the Equivalence P-Value will be different from that of the confidence interval. In this case, we recommend using just the confidence interval method.
 6. Due to the complexity of calculations, this template uses VBA macros rather than Excel formulas. SigmaXL must be initialized and appear on the menu in order for this template to function.

REFERENCES

1. Beal, S. L. (1987), Asymptotic confidence intervals for the difference between two binomial parameters for use with small samples. *Biometrics*, 43, 941-950.
2. Newcombe, R. G. (1998b), Interval estimation for the difference between independent proportions: Comparison of eleven methods. *Statistics in Medicine*, 17:873–890.
3. Radhakrishna S., Murthy B.N., Nair N.G., Jayabal P., Jayasri R. (1992), Confidence intervals in medical research. *Indian J Med Res.*, Jun;96:199-205.

Basic Statistical Templates – 1 Poisson Rate Test and Confidence Interval

Click **SigmaXL > Templates & Calculators > Basic Statistical Templates > 1 Poisson Rate Test and Confidence Interval** to access the 1 Poisson Rate calculator. The template gives the following default example.

1 Sample Poisson Rate Test and Confidence Interval		
Sample Data (user inputs):		
Number of occurrences	x	1
Sample Size (area of opportunity)	nt	10.0
Null Hypothesis (hypothesized rate)	H ₀ : Rate (λ) =	0.5
Alternative Hypothesis	H _a : Rate (λ)	Not Equal To
Confidence Level (enter .95 for 95%)	100*(1-α)%	95.0%
Hypothesis Test Method		Poisson Exact
Confidence Interval Method		Exact (Garwood Chi-Square)
Results:		
Sample mean rate (x/(nt))		0.1000
alpha		0.0500
Z-statistic (normal)		-1.7889
Poisson exact probability P-Value (2-sided)		0.0809
Upper Confidence Limit (2-sided)		0.5572
Lower Confidence Limit (2-sided)		0.0025
Test Information		
Null Hypothesis H ₀ : Rate (λ) = 0.5		Fail to Reject
Alternative Hypothesis H _a : Rate (λ) ≠ 0.5		
Interpretation of P-Value and Confidence Intervals		
The P-Value is the two-sided (two-tail) probability of observing a sample rate at least as extreme as 0.1000, given that the null hypothesis is true. Since the P-Value is greater than alpha (0.05), we fail to reject the null hypothesis: Rate (λ) = 0.5.		
We are 95% confident that the true population rate lies within the interval (0.0025 to 0.5572).		

Notes:

1. Enter summarized Sample Data, Null Hypothesis and Confidence Level in cells with yellow highlight. Do not modify any other part of this worksheet.
2. Sample Size (area of opportunity) can be a count, time, length or other unit that defines the area of opportunity for an occurrence.
3. Select Alternative Hypothesis and Methods using drop-down. If x = 0, it is recommended that the Alternative Hypothesis be set to one-sided "Less Than."

4. Hypothesis Test Method:

- Poisson Exact is recommended.
- Use Normal Approximation only when the number of occurrences, $x > 10$.

5. Confidence Interval Method (recommendations based on literature review, see references below):

- Confidence intervals for Poisson rates have an "oscillation" phenomenon where the coverage probability varies with x and n .
- Exact (Garwood Chi-Square) is strictly conservative and will guarantee the specified confidence level as a minimum coverage probability, but results in wide intervals. Recommended only for applications requiring strictly conservative intervals.
- Jeffreys has mean coverage probability matching the specified confidence interval and is recommended.
- Normal Approximation (Wald) is not recommended, but included here to validate hand calculations. Use only when the number of occurrences, $x > 10$.

6. This calculator will automatically detect the version of Excel and use the appropriate statistical functions. If Excel 2010 or higher, then the Excel 2010 statistical functions are used.

REFERENCES

1. Garwood, F. (1936). Fiducial limits for the Poisson distribution. *Biometrika* 28:437–442.
2. Swift, M.B. (2009). Comparison of confidence intervals for a Poisson Mean-Further considerations, *Communication in Statistics — Theory and Methods*, 38, 748–759.

Basic Statistical Templates – 2 Poisson Rates Test and Confidence Interval

Click **SigmaXL > Templates & Calculators > Basic Statistical Templates > 2 Poisson Rates Test and Confidence Interval** to access the 2 Poisson Rates calculator. The template gives the following default example.

2 Poisson Rates Test and Confidence Interval			
Sample Data (user inputs):		Sample 1	Sample 2
Number of occurrences	x	1	2
Sample Size (area of opportunity)	nt	10.0	10.0
Null Hypothesis (hypothesized difference)	$H_0: \text{Rate}_1 - \text{Rate}_2 (\lambda_1 - \lambda_2) =$	0	
Alternative Hypothesis	$H_a: \text{Rate}_1 - \text{Rate}_2 (\lambda_1 - \lambda_2)$	Not Equal To	
Confidence Level (enter .95 for 95%)	$100 \cdot (1 - \alpha)\%$	95.0%	
Hypothesis Test Method		Exact (Binomial Method)	
Confidence Interval Method		Jeffreys Hybrid	
Results:			
Sample mean rate (x/(nt))		0.1000	0.2000
Sample rate difference		-0.1000	
alpha		0.0500	
Exact probability P-Value (2-sided)		1.0000	
Upper Confidence Limit (2-sided)		0.3001	
Lower Confidence Limit (2-sided)		-0.5505	

Test Information	
Null Hypothesis $H_0: \lambda_1 - \lambda_2 = 0$	Fail to Reject
Alternative Hypothesis $H_a: \lambda_1 - \lambda_2 \neq 0$	

Interpretation of P-Value and Confidence Intervals
The P-Value is the two-sided (two-tail) probability of observing a difference in rates at least as extreme as -0.1000, given that the null hypothesis is true. Since the P-Value is greater than alpha (0.05), we fail to reject the null hypothesis $H_0: \lambda_1 - \lambda_2 = 0$.
We are 95% confident that the true difference in population rates lies within the interval (-0.5505 to 0.3001).

Notes:

1. Enter summarized Sample Data, Null Hypothesis and Confidence Level in cells with yellow highlight. Do not modify any other part of this worksheet.
2. Sample Size (area of opportunity) can be a count, time, length or other unit that defines the area of opportunity for an occurrence.
3. Select Alternative Hypothesis and Methods using drop-down.

4. Hypothesis Test Method:

- Exact (Binomial Method) is recommended, but is only valid for $H_0: \lambda_1 - \lambda_2 = 0$.
- Use Normal Approximation only when the number of occurrences for each sample is greater than 10. Unpooled is used if $H_0: \lambda_1 - \lambda_2 \neq 0$. Use pooled if $H_0: \lambda_1 - \lambda_2 = 0$.

5. Confidence Interval Method:

- Confidence intervals for difference in Poisson rates have an "oscillation" phenomenon where the coverage probability varies with x and n .
- Jeffreys Hybrid is recommended and has a mean coverage probability that is close to the specified confidence interval. (See Li, 2011).
- Normal Approximation is not recommended, but included here to validate hand calculations. Use only when the number of occurrences for each sample is greater than 10.

6. Due to the complexity of calculations, this template uses vba macros rather than Excel formulas. SigmaXL must be initialized and appear on the menu in order for this template to function.

REFERENCES

1. Li, H.Q., Tang, M.L., Poon, W.Y. and Tang, N.S. (2011), Confidence intervals for difference between two Poisson rates, *Comm. Stat. Simul. Comput.* 40, pp. 1478–1491.
2. Przyborowski, J., Wilenski, H., (1940), Homogeneity of results in testing samples from Poisson series. *Biometrika* 31, 313–323.

Basic Statistical Templates – 2 Poisson Rates Equivalence Test

Click **SigmaXL > Templates & Calculators > Basic Statistical Templates > 2 Poisson Rates Equivalence Test** to access the 2 Poisson Rates Equivalence Test (Two One-Sided Tests TOST) calculator. The template gives the following default example.

2 Poisson Rates Equivalence Test (Two One-Sided Tests TOST)			
Sample Data (user inputs):		Test/Treatment Sample	Reference/Control Sample
Number of occurrences	x	1	2
Sample Size (area of opportunity)	nt	10.0	10.0
Upper Equivalence Limit	UEL	0.5	
Lower Equivalence Limit	LEL	-0.5	
Confidence Level (enter .95 for 95%)	100*(1- α)%	95.0%	
Confidence Interval Method		Jeffreys Hybrid	
Results:			
Sample mean rate (x/(nt))		0.1000	0.2000
Sample rate difference		-0.1000	
alpha		0.0500	
Z1-statistic (upper, normal unpooled)		-3.4641	
Z2-statistic (lower, normal unpooled)		2.3094	
P1-Value (upper)		0.0003	
P2-Value (lower)		0.0105	
Equivalence P-Value (maximum of P1 and P2)		0.0105	
Upper Confidence Limit Rate Difference		0.2239	
Lower Confidence Limit Rate Difference		-0.4630	

Interpretation of Equivalence P-Value and Confidence Intervals
Since the Equivalence P-Value is less than alpha (0.05), conclude that equivalence is true.
Since the 95% confidence interval is within the equivalence interval, conclude that equivalence is true.

Notes:

1. Enter summarized Sample Data, Upper & Lower Equivalence Limits and Confidence Level in cells with yellow highlight. Do not modify any other part of this worksheet.
2. Sample Size (area of opportunity) can be a count, time, length or other unit that defines the area of opportunity for an occurrence.
3. Select Confidence Interval Method using drop-down.
 - Confidence intervals for difference in Poisson rates have an "oscillation" phenomenon where the coverage probability varies with x and n.
 - Jeffreys Hybrid is recommended and has a mean coverage probability that is close to the specified confidence interval. (See Li, 2011).
 - Normal Approximation will match the P-Value calculations, but should only be used when the number of occurrences for each sample is greater than 10.

4. Null hypothesis for P1: Rate Difference \geq UEL; Null hypothesis for P2: Rate Difference \leq LEL. Both null hypotheses must be rejected to conclude that equivalence is true. The P-Values are based on the normal approximation unpooled method.
5. LEL and UEL establish the zone or region of equivalence and are determined by what size rate difference is considered practically significant.
6. Since the Jeffreys Hybrid confidence interval option uses a different method than the normal based hypothesis test, it is possible that the conclusion from the Equivalence P-Value will be different from that of the confidence interval. In this case, we recommend using just the confidence interval method.
7. Due to the complexity of calculations, this template uses VBA macros rather than Excel formulas. SigmaXL must be initialized and appear on the menu in order for this template to function.

REFERENCES

1. Li, H.Q., Tang, M.L., Poon, W.Y. and Tang, N.S. (2011), Confidence intervals for difference between two Poisson rates, *Comm. Stat. Simul. Comput.* 40, pp. 1478–1491.
2. Przyborowski, J., Wilenski, H., (1940), Homogeneity of results in testing samples from Poisson series. *Biometrika* 31, 313–323.

Basic Statistical Templates – One-Way Chi-Square Goodness-of-Fit Test

Click **SigmaXL > Templates & Calculators > Basic Statistical Templates > One-Way Chi-Square Goodness-of-Fit Test** to access the One-Way Chi-Square Goodness-of-Fit Test calculator. The template gives the following default example.

Sample Data (user inputs):			Results:	
Category ID	Observed Count	Specify Historical Count (optional)		
1	1	1	Percentage of cells with expected count < 5	40.0%
2	2	2	Sum of Observed Counts	26
3	3	3	Chi-Square	3.4231
4	10	4	df	4
5	10	5	p-value	0.4897

Show Calculations

Notes:

1. Enter Category ID, Observed Count values and (optional) Historical Counts. Do not modify any other part of this worksheet.
2. If optional Historical Counts are not specified, chi-square is calculated using equal expected proportions.
3. If optional Historical Counts are specified, a value must be entered for each observed count.
4. Chi-Square statistic requires that no more than 20% of cells have an expected count less than 5.

Basic Statistical Templates – One-Way Chi-Square Goodness-of-Fit Test - Exact

Click **SigmaXL > Templates & Calculators > Basic Statistical Templates > One-Way Chi-Square Goodness-of-Fit Test - Exact** to access the One-Way Chi-Square Goodness-of-Fit Test – Exact calculator. The template gives the following default example.

Sample Data (user inputs):		
Category ID	Observed Count	Specify Historical Count (optional)
1	1	1
2	2	2
3	3	3
4	10	4
5	10	5
6		

Results:	
Percentage of cells with expected count < 5	40.0%
Sum of Observed Counts	26
Chi-Square	3.4231
DF	4
Chi-Square P-Value	0.4897
Exact P-Value	0.5013

Show/Hide Calculations

Exact P-Value

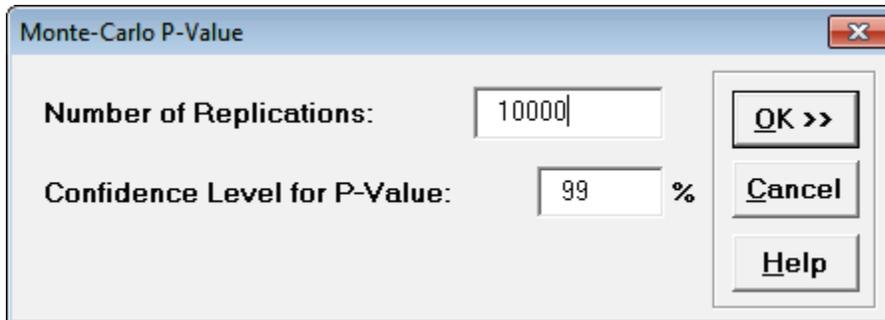
Monte Carlo P-Value

Notes:

1. Enter Category ID, Observed Count values and (optional) Historical Counts. Do not modify any other part of this worksheet.
2. If optional Historical Counts are not specified, chi-square is calculated using equal expected proportions.
3. If optional Historical Counts are specified, a value must be entered for each observed count.
4. The Chi-Square statistic requires that no more than 20% of cells have an expected count less than 5 (and none of the cells have an expected count less than 1). If this assumption is not satisfied, the Chi-Square approximation may be invalid and Exact or Monte Carlo P-Values should be used. In this example 40% of the cells have an expected count less than 5 so Exact should be used.
5. This example shows that the Exact P-Value is 0.5013 and the “large sample” Chi-Square P-Value is 0.4897.
6. Chi-Square Exact solves the permutation problem using enhanced enumeration. For further details refer to the Appendix [Exact and Monte Carlo P-Values for Nonparametric and Contingency Test](#).
7. It is important to note that while exact P-Values are “correct,” they do not increase (or decrease) the power of a small sample test, so they are not a solution to the problem of failure to reject the null due to inadequate sample size.
8. For data that requires more computation time than specified, Monte Carlo P-Values provide an approximate (but unbiased) P-Value that typically matches exact to two decimal places using 10,000 replications. One million replications give a P-Value that is typically accurate to three decimal places. A confidence interval (99% default) is given for the Monte Carlo P-Values.

Monte Carlo Example:

1. The Exact P-Value for this example is solved very quickly, so Monte Carlo is not needed, but we will run it for continuity in the example. Click **Monte Carlo P-Value**. Select **Number of Replications** = 10000 and **Confidence Level for P-Value** = 99%.



Tip: The Monte Carlo 99% confidence interval for P-Value is **not** the same as a confidence interval on the test statistic due to data sampling error. The confidence level for the hypothesis test statistic is still 95%, so **all reported P-Values less than .05 will be highlighted in red** to indicate significance. The 99% Monte Carlo P-Value confidence interval is due to the uncertainty in Monte Carlo sampling, and it becomes smaller as the number of replications increases (irrespective of the data sample size). The Exact P-Value will lie within the stated Monte Carlo confidence interval 99% of the time.

2. Click **OK**. Results:

Results:	
Percentage of cells with expected count < 5	40.0%
Sum of Observed Counts	26
Chi-Square	3.4231
DF	4
Chi-Square P-Value	0.4897
Monte Carlo P-Value	0.5012
Monte Carlo P-Value 99% CI Upper	0.5128
Monte Carlo P-Value 99% CI Lower	0.4896

The Monte Carlo P-Value here is 0.5012 with a 99% confidence interval of 0.4896 to 0.5128. This will be slightly different every time it is run (the Monte Carlo seed value is derived from the system clock). The true Exact P-Value = 0.5013 lies within this confidence interval.

Small Sample Exact Example:

3. Now we will consider a small sample problem. Enter the following values for sample data in the yellow highlight region. Note that the displayed Monte Carlo (or Exact) P-Values are cleared when new data is entered in the template:

Sample Data (user inputs):		
Category ID	Observed Count	Specify Historical Count (optional)
1	7	3
2	1	3
3	1	3
4	1	1
5		
6		
7		
8		

Results:	
Percentage of cells with expected count < 5	100.0%
Sum of Observed Counts	10
Chi-Square	8.0000
DF	3
Chi-Square P-Value	0.0460
Monte Carlo P-Value	
Monte Carlo P-Value 99% CI Upper	
Monte Carlo P-Value 99% CI Lower	

Show/Hide Calculations

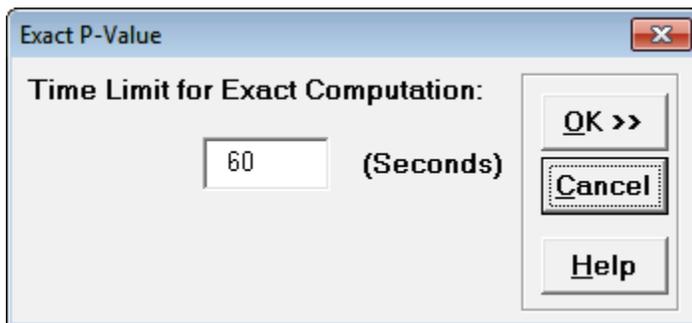
Exact P-Value

Monte Carlo P-Value

This example is adapted from Mehta, C.R. and Patel, N. R., *IBM SPSS Exact Tests*, IBM Corp., page 44.

[ftp://public.dhe.ibm.com/software/analytics/spss/documentation/statistics/21.0/en/client/Manuals/IBM SPSS Exact Tests.pdf](ftp://public.dhe.ibm.com/software/analytics/spss/documentation/statistics/21.0/en/client/Manuals/IBM_SPSS_Exact_Tests.pdf)

- Click **Exact P-Value**. Select **Time Limit for Exact Computation** = 60 seconds.



- Click **OK**. Results:

Results:	
Percentage of cells with expected count < 5	100.0%
Sum of Observed Counts	10
Chi-Square	8.0000
DF	3
Chi-Square P-Value	0.0460
Exact P-Value	0.0523

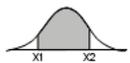
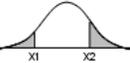
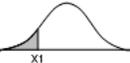
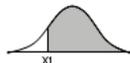
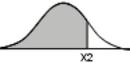
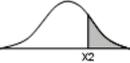
Note that the Exact P-Value is 0.0523 which is a “fail-to-reject” of the null hypothesis (H0), but the “large sample” or “asymptotic” Chi-Square P-Value incorrectly rejected H0 with a P-Value of 0.046. The exact P-Value matches that given in the reference. The error can also go the other way, where a large sample Chi-Square P-Value is a “fail-to-reject” of the null hypothesis and the Exact P-Value is a rejection of H0.

In conclusion, always use the Exact (or Monte Carlo) P-Value when the Chi-Square large sample assumptions are not met.

Probability Distribution Calculators – Normal

Click **SigmaXL > Templates & Calculators > Probability Distribution Calculators > Normal** to access the Normal Distribution Probability calculator. The template gives the following default example:

Normal Distribution Probability Calculator		
Input the following information:		
Mean	μ	0
Standard Deviation	σ	1
Lower Bound (or LSL)	X1	1
Upper Bound (or USL)	X2	2

Select the Probability of Interest from the following probabilities:		
Areas: between X1 and X2		Prob(1 ≤ X ≤ 2) 0.135905122
outside X1 and X2		Prob(X ≤ 1 AND X ≥ 2) 0.864094878
below X1		Prob(X ≤ 1) 0.841344746
above X1		Prob(X ≥ 1) 0.158655254
below X2		Prob(X ≤ 2) 0.977249868
above X2		Prob(X ≥ 2) 0.022750132

Basic MSA Templates – Type 1 Gage Study

Click **SigmaXL > Templates & Calculators > Basic MSA Templates > Gage Bias and Linearity Study** to access the Gage Bias and Linearity Study template.

The following example is given in **SigmaXL > Help > Template Examples > MSA > Type 1 Gage Study**.

Type 1 Gage Study

Gage Name:	Example Type 1 Gage Study	
Date of Study:		
Performed By:		
Notes:		

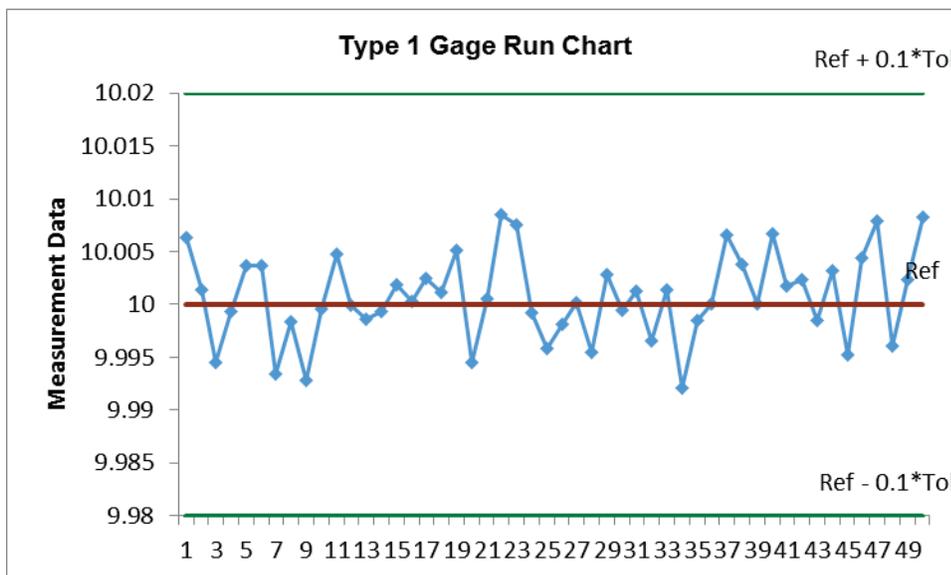
Tolerance (Tol, USL - LSL):	0.2
Part Reference Value (Ref):	10
Gage Resolution:	0.0001
StDev Multiplier:	6.00
Percent of Tolerance for Cg/Cgk:	20

Count	50
Mean	10.00071
StDev	0.004149
StDev*Multiplier (SV)	0.024893
Ref + 0.1*Tol	10.02
Ref - 0.1*Tol	9.98

Bias	0.00071
Bias t-Statistic	1.210106
Bias P-Value	0.2320
Gage Capability Cg	1.61
Gage Capability Cgk	1.55
Gage Repeatability % (100*SV/Tol)	12.45%
Gage Repeatability & Bias % (20/Cgk)	12.90%
Resolution % of Tolerance	0.05%

Trial:	Measurement Data
1	10.0063
2	10.0014
3	9.9945

Type 1 Gage Run Chart



Notes:

1. This calculator assumes that the data are normally distributed.
2. Enter your data in the **Measurement Data** column. Minimum recommended sample size is 25, but 50 is preferred. This template has a maximum sample size of 1000, suitable for automated test systems.
3. Enter **Tolerance** (Upper Specification Limit USL - Lower Specification Limit LSL).
4. Enter **Part Reference Value**. If missing, Gage Capability Cgk is not computed. Bias will assume a 0 value.
5. Enter **Gage Resolution**. If missing, Resolution % of Tolerance is not computed.
6. The default **StDev Multiplier** is 6 (99.73% coverage). Change this to 5.15 if AIAG convention is being used (99% coverage). Dietrich [1] recommends using 4 for approximate 95% coverage. We recommend 6 for typical and critical measurement systems, and 4 for non-critical.
7. The default **Percent of Tolerance** for Cg/Cgk is 20. Stacey [2] suggests using 10 for critical measurement systems.
8. If **Cg** and **Cgk** are ≥ 1.33 , the measurement device is capable. If **Cg** or **Cgk** are < 1.33 , the measurement device is not capable and needs to be improved.
9. If **Resolution % of Tolerance** (RE%) is $\leq 5\%$, the resolution is acceptable. If RE% is $> 5\%$ the resolution is not acceptable.
10. Click the **Type 1 Gage Run Chart** button to create a Run Chart with center line = Ref, upper limit = $\text{Ref} + 0.1 * \text{Tol}$ and lower limit = $\text{Ref} - 0.1 * \text{Tol}$ (assuming 20% of Tolerance).
11. This calculator requires Excel 2010 or higher.

REFERENCES

1. Dietrich, E. (2002), "Measurement System Capability" Reference Manual, http://www.q-das.com/uploads/tx_sbdownloader/Leitfaden_v21_me.pdf.
2. Stacey, J., Mathematics of Measurement Systems Analysis, Rolls-Royce plc., http://www.mei.org.uk/files/pdf/Mathematics_of_Measurement_Systems_Analysis.pdf

Basic MSA Templates – Gage Bias and Linearity Study

Click **SigmaXL > Templates & Calculators > Basic MSA Templates > Gage Bias and Linearity Study** to access the Gage Bias and Linearity Study template.

The following example is given in **SigmaXL > Help > Template Examples > MSA > Gage Bias and Linearity Study**.

Gage Bias and Linearity Study

Gage Name:	Example Gage Bias and Linearity Study - AIAG MSA Reference Manual, 4th Edition, Page 99
Date of Study:	
Performed By:	
Notes:	

Process Variation (SV from Gage R&R or 6 * Historical StDev):	10
---	----

Gage Linearity Report					
-----------------------	--	--	--	--	--

Part:	1	2	3	4	5
Reference Value:	2	4	6	8	10
Trial	Bias Calculations				
1	2.7	5.1	5.8	7.6	9.1
2	2.5	3.9	5.7	7.7	9.3
3	2.4	4.2	5.9	7.8	9.5
4	2.5	5	5.9	7.7	9.3
5	2.7	3.8	6	7.8	9.4
6	2.3	3.9	6.1	7.8	9.5
7	2.5	3.9	6	7.8	9.5
8	2.5	3.9	6.1	7.7	9.5
9	2.4	3.9	6.4	7.8	9.6
10	2.4	4	6.3	7.5	9.2
11	2.6	4.1	6	7.6	9.3
12	2.4	3.8	6.1	7.7	9.4
13					
14					
15					
16					
17					
18					
19					
20					

Part:	1	2	3	4	5
Reference Value:	2	4	6	8	10
Trial	Bias Calculations				
1	0.7	1.1	-0.2	-0.4	-0.9
2	0.5	-0.1	-0.3	-0.3	-0.7
3	0.4	0.2	-0.1	-0.2	-0.5
4	0.5	1	-0.1	-0.3	-0.7
5	0.7	-0.2	0	-0.2	-0.6
6	0.3	-0.1	0.1	-0.2	-0.5
7	0.5	-0.1	0	-0.2	-0.5
8	0.5	-0.1	0.1	-0.3	-0.5
9	0.4	-0.1	0.4	-0.2	-0.4
10	0.4	0	0.3	-0.5	-0.8
11	0.6	0.1	0	-0.4	-0.7
12	0.4	-0.2	0.1	-0.3	-0.6
13					
14					
15					
16					
17					
18					
19					
20					

Bias	0.491666667	0.125	0.025	-0.291666667	-0.616666667
% Bias	4.92%	1.25%	0.25%	2.92%	6.17%
P-Value	0.0000	0.3540	0.6671	0.0000	0.0000

Average Bias	-0.053333333
% Average Bias	0.53%
P-Value	0.0894

	Coeff	SE Coeff	P-Value
Constant	0.736666667	0.072524273	0.0000
Slope	-0.131666667	0.010933445	0.0000

R-Square	71.43%
Adj, R-Square	70.94%
StDev	0.239539789
Linearity	1.316666667
% Linearity	13.17%

Gage Linearity

$y = -0.1317x + 0.7367$

Percent of Process Variation

90

Notes:

1. Enter **Process Variation** (SV from Gage R&R or $6 * \text{Historical StDev}$). If neither are available, use Process Tolerance. % Average Bias will be blank if not specified. Note that % Linearity does not require Process Variation.
2. Enter part **Reference Values**. These should be relative to a traceable standard, but if not available, see AIAG MSA Reference Manual [1], pp. 92-93 for suggestions on how to obtain reference values. Reference values should cover the operating range of the gage.
3. Enter part measurement data in the yellow highlight region below the respective part/reference value. AIAG recommends a minimum of 10 trials per part/reference. **It is important that the parts be selected at random for each trial in order to minimize appraiser "recall" bias.**
4. Click **Gage Linearity Report** button to generate the Linearity Report.
5. If a P-Value is $< .05$, it is highlighted in bold red to indicate significance. If the Average Bias P-Value or Linearity Slope P-Value are $< .05$, the respective bar in Percent of Process Variation will also be red to highlight significance.
6. Statistical significance is the best indicator of a problem with bias or linearity. Some practitioners also use a rule of thumb that a good measurement system will have % Average Bias $< 5\%$ and % Linearity $< 5\%$.
7. The dashed lines on the Linearity scatter plot are the regression fit 95% confidence intervals. If the zero reference line is not within the intervals, the linearity slope is significant, indicating that the measurement system linearity is not acceptable.
8. This calculator requires Excel 2010 or higher. Linearity regression uses vba macros rather than Excel formulas. SigmaXL must be initialized and appear on the menu in order for this template to function.
9. Rows 34-39 and 42 are hidden to simplify the output report. These may be unhidden to view how the calculations are performed.

REFERENCE

1. Automotive Industry Action Group AIAG (2010). Measurement Systems Analysis MSA Reference Manual, 4th Edition

Basic MSA Templates – Gage R&R Study (MSA)

Click **SigmaXL > Templates & Calculators > Basic MSA Templates > Gage R&R Study (MSA)** to access the Gage R&R Study template.

The following example is given in **SigmaXL > Help > Template Examples > MSA > Gage R&R**. If prompted, please ensure that macros are enabled.

Gage Name:	Example Gage R&R
Date of Study:	
Performed By:	
Notes:	

Process Tolerance	
USL:	2
LSL:	1

StDev Multiplier:	6.00
--------------------------	------

Operator A	Part 1	Part 2	Part 3	Part 4	Part 5	Part 6	Part 7	Part 8	Part 9	Part 10
Reading 1	1.34	1.31	1.52	1.44	1.44	1.29	1.52	1.49	1.44	1.52
Reading 2	1.31	1.36	1.52	1.42	1.49	1.23	1.52	1.49	1.47	1.52
Reading 3										

Operator B	Part 1	Part 2	Part 3	Part 4	Part 5	Part 6	Part 7	Part 8	Part 9	Part 10
Reading 1	1.29	1.29	1.55	1.42	1.42	1.21	1.52	1.49	1.39	1.52
Reading 2	1.29	1.26	1.49	1.39	1.39	1.21	1.55	1.47	1.36	1.49
Reading 3										

Operator C	Part 1	Part 2	Part 3	Part 4	Part 5	Part 6	Part 7	Part 8	Part 9	Part 10
Reading 1	1.26	1.44	1.55	1.42	1.42	1.23	1.52	1.49	1.42	1.55
Reading 2	1.29	1.42	1.52	1.42	1.42	1.26	1.55	1.49	1.42	1.55
Reading 3										

Gage R&R Metrics	Variance Component	% Contribution of Variance Component	StDev	StDev * Multiplier	% Total Variation (TV)	% Tolerance
Gage R&R:	0.001289167	11.44	0.035904967	0.215429803	33.83	21.54
Operator (AV Appraiser Variation):	0.000258426	2.29	0.016075631	0.096453789	15.15	9.65
Operator * Part (INT Interaction):	0.000627407	5.57	0.025048102	0.150288611	23.60	15.03
Reproducibility (SQRT(AV^2 + INT^2)):	0.000885833	7.86	0.029762952	0.178577714	28.04	17.86
Repeatability (EV Equipment Variation):	0.000403333	3.58	0.02008316	0.120498963	18.92	12.05
Part Variation (PV):	0.009977222	88.56	0.099886046	0.599316277	94.10	59.93
Total Variation (TV):	0.011266389	100.00	0.106143247	0.636859482	100.00	63.69

Create Stacked Column Format for "Analyze Gage R&R" >>

Click **Create Stacked Column Format for "Analyze Gage R&R" >>** if you wish to analyze the above data using SigmaXL's menu **Gage R&R Analysis**

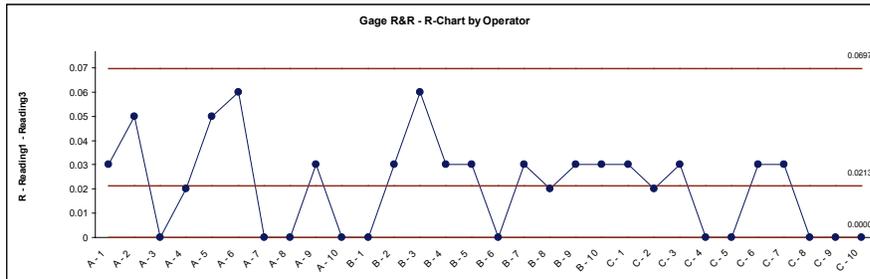
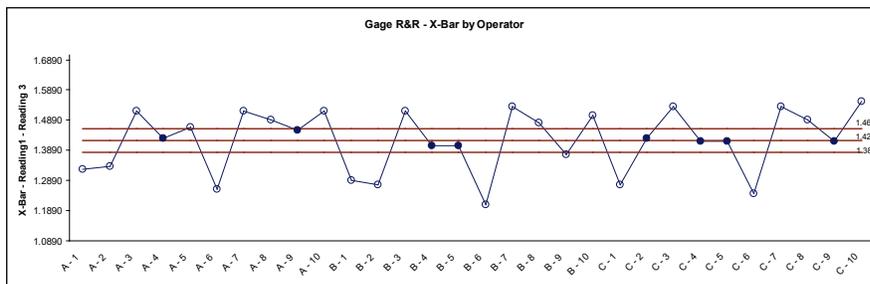
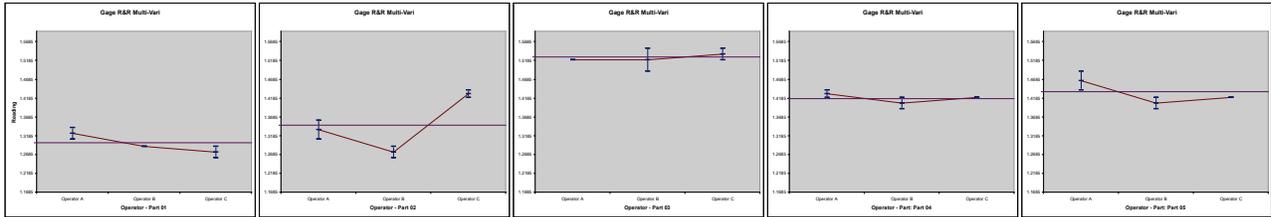
Notes for use of the Gage R&R Template:

1. The Automotive Industry Action Group (AIAG) recommended study includes 10 Parts, 3 Operators and 3 Replicates. The template calculations will work with a minimum of 2 Operators, 2 Parts and 2 Replicates. The data should be balanced with each operator measuring the same number of parts and the same number of replicates. Use **SigmaXL > Measurement Systems Analysis** to specify up to 30 Parts, 10 Operators and 10 Replicates.
2. Enter process Upper Specification Limit (USL) and Lower Specification Limit (LSL) in the Process Tolerance window. This is used to determine the % Tolerance metrics. If the specification is single-sided, leave both entries blank.

3. The default StDev multiplier is 6. Change this to 5.15 if AIAG convention is being used.
4. The cells shaded in light blue highlight the critical metrics Gage R&R % Total Variation (also known as %R&R) and %Tolerance: < 10% indicates a good measurement system; > 30% indicates an unacceptable measurement system.

Basic MSA Templates - Gage R&R: Multi-Vari & X-bar R Charts

Click **SigmaXL > Templates & Calculators > Basic MSA Templates > Gage R&R: Multi-Vari & X-bar R Charts** to access the Multi-Vari & X-bar R Charts template. An example is given in **SigmaXL > Help > Template Examples > MSA > Gage R&R - Multi-Vari** and **Gage R&R – X-Bar R**:



Notes for use of the Gage R&R: Multi-Vari & X-bar R Charts:

1. The Gage R&R Multi-Vari and X-bar & R charts can only be generated if a Gage R&R template has been completed and is the active worksheet. Select the Gage R&R worksheet and click **SigmaXL > Templates & Calculators > Gage R&R: Multi-Vari & X-Bar R Charts** to create the above charts.
2. The Multi-Vari chart shows each Part as a separate graph. Each Operator's response readings are denoted as a vertical line with the top tick corresponding to the Maximum value, bottom tick is the Minimum, and the middle tick is the Mean. The horizontal line across each graph is the overall average for each part.
3. When interpreting the X-bar and R chart for a Gage R&R study, it is desirable that the X-bar chart be out-of-control, and the Range chart be in-control. The control limits are derived from within Operator repeatability.

Basic MSA Templates – Attribute Gage R&R (MSA)

Click **SigmaXL > Templates & Calculators > Basic MSA Templates > Attribute MSA** to access the Attribute MSA Study template. An example is given in **SigmaXL > Help > Template Examples > Attribute MSA**.

MEASUREMENT SYSTEM STUDY FOR ATTRIBUTE DATA (Recommend 3 Appraisers, 2 Repeats, Minimum of 10 Good Parts and 10 Bad Parts)

Product/Unit Name:	Example Attribute MSA
Date of Study:	
Performed By:	
Notes:	

Good Part or Unit:	G
Bad Part or Unit:	NG

Part	True Standard
1	G
2	G
3	NG
4	G
5	NG
6	NG
7	G
8	NG
9	G
10	NG
11	G
12	NG
13	NG
14	G
15	NG
16	G
17	G
18	NG
19	NG
20	G

Appraiser A		Appraiser B		Appraiser C	
Trial # 1	Trial #2	Trial # 1	Trial #2	Trial # 1	Trial #2
G	G	NG	G	G	G
NG	G	G	G	G	G
NG	NG	NG	NG	NG	NG
G	G	G	NG	G	G
NG	NG	NG	NG	NG	NG
NG	NG	NG	NG	NG	NG
G	G	NG	G	G	G
NG	NG	NG	NG	NG	NG
G	G	G	G	G	G
NG	NG	NG	NG	NG	NG
G	G	NG	G	G	G
NG	NG	NG	NG	NG	NG
NG	NG	NG	NG	NG	NG
NG	G	NG	G	G	G
NG	NG	NG	G	NG	NG
G	NG	G	G	G	NG
G	G	G	G	G	G
NG	NG	NG	G	NG	NG
NG	NG	NG	NG	NG	NG
G	G	NG	G	G	G

Create Stacked Column format to Analyze with "Attribute MSA (Binary)" >>

Attribute MSA Analysis:

Within Appraiser Agreement:	# Inspected	# Matched	Percent	95% LC	95% UC	Cohen's Kappa
A	20	17	85.00%	62.11%	96.79%	0.6939
B	20	12	60.00%	36.05%	80.88%	0.2381
C	20	19	95.00%	75.13%	99.87%	0.9000

Type I = False Reject (Appraiser Rejected Good Part)
 Type II = False Accept (Appraiser Accepted Bad Part)
 Mixed = Assessments across trials are not identical

Each Appraiser vs Standard Agreement:	# Inspected	# Matched	Percent	95% LC	95% UC	Type I Errors	Type I Error %	Type II Errors	Type II Error %	Mixed Error	Mixed Error %	Cohen's Kappa
A	20	17	85.00%	62.11%	96.79%	0	0.00%	0	0.00%	3	15.00%	0.8500
B	20	12	60.00%	36.05%	80.88%	0	0.00%	0	0.00%	8	40.00%	0.6000
C	20	19	95.00%	75.13%	99.87%	0	0.00%	0	0.00%	1	5.00%	0.9500

Between Appraiser Agreement:	# Inspected	# Matched	Percent	95% LC	95% UC
	20	10	50.00%	27.20%	72.80%

All Appraisers vs Standard Agreement:	# Inspected	# Matched	Percent	95% LC	95% UC	Cohen's Kappa
	20	10	50.00%	27.20%	72.80%	0.8000

Click **Create Stacked Column Format to Analyze with "Attribute MSA (Binary)" >>** if you wish to analyze the above data using SigmaXL's menu Attribute MSA Analysis tool.

Notes for use of the Attribute Gage R&R (MSA) Template:

1. Attribute Gage R&R is also known as Attribute Agreement Analysis.

SigmaXL: Measure Phase Tools

2. Recommend for study: 3 Appraisers, 2 to 3 Replicates, Minimum of 10 Good Parts and 10 Bad Parts. The data should be balanced with each appraiser evaluating the same number of parts and the same number of replicates.
3. Specify the Good Part or Unit as G or other appropriate text (P, Y, etc.). Specify the Bad Part or Unit as NG or other appropriate text (F, N, etc.). Be careful to avoid typing or spelling errors when entering the results. A space accidentally inserted after a character will be treated as a different value leading to incorrect results.

Basic Process Capability Templates – Process Sigma Level – Discrete Data

Click **SigmaXL > Templates & Calculators > Basic Process Capability Templates > Process Sigma Level – Discrete** to access the Process Sigma Level – Discrete calculator. The template gives the following default example.

Sample Data (user inputs):		
Number of units	n	500
Total number of defects observed	d	5
Number of defect opportunities per unit	o	1
Sigma Shift (typically +1.5 for long term data)		1.5
Results:		
Defects per Unit	dpu	0.01
Defects per Million Opportunities	dpmo	10,000.0
Defects per Opportunity	dpo%	1.00%
Yield	yield%	99.00%
Process Sigma Level	sigma	3.826

Notes for use of the Process Sigma Calculator for Discrete Data:

1. Total number of defects should include defects made and later fixed.
2. Sample size should be large enough to observe 5 defects.

Basic Process Capability Templates – Process Sigma Level – Continuous Data

Click **SigmaXL > Templates & Calculators > Basic Process Capability Templates > Process Sigma Level – Continuous** to access the Process Sigma Level – Continuous calculator. The template gives the following default example.

 Process Sigma Level Calculator - Continuous Data (Assumes that data are normally distributed)		
Sample Data (user inputs):		
Mean	x-bar	0
Standard Deviation	s	1
Upper Specification Limit	USL	3
Lower Specification Limit	LSL	-3
Sigma Shift (typically +1.5 for long term data)		1.5
Results:		
Expected dpm > USL		1349.9
Expected % > USL		0.13%
Expected dpm < LSL		1349.9
Expected % < LSL		0.13%
Expected dpm (overall)		2699.8
Expected yield (overall) %		99.73%
Process Sigma Level		4.282

Note: This calculator assumes that the Mean and Standard Deviation are computed from data that are normally distributed.

Basic Process Capability Templates – Process Capability Indices

Click **SigmaXL > Templates & Calculators > Basic Process Capability Templates > Process Capability** to access the Process Capability Indices calculator. The template gives the following default example.

 Process Capability Indices - Continuous Data (Assumes that data are normally distributed)		
Sample Data (user inputs):		
Mean	x-bar	0
Standard Deviation	s	1
Upper specification limit	USL	3
Lower specification limit	LSL	-3
Results:		
	Cp, Pp	1.00
	Cpu, Ppu	1.00
	Cpl, Ppl	1.00
	Cpk, Ppk	1.00

Notes for use of the Process Capability Indices Calculator:

1. This calculator assumes that the Mean and Standard Deviation are computed from data that are normally distributed.
2. Reports Cp, Cpk if entered S is Within or Short Term (using a control chart).
3. Reports Pp, Ppk if entered S is Overall or Long Term.

Basic Process Capability Templates – Process Capability & Confidence Intervals

Click **SigmaXL > Templates & Calculators > Basic Process Capability Templates > Process Capability & Confidence Intervals** to access the Process Capability Indices & Confidence Intervals calculator. The template gives the following default example.

 Process Capability Indices and Confidence Intervals - Continuous Data (Assumes that data are normally distributed)					
Sample Data (user inputs):					
Mean	x-bar	1	Upper Specification Limit	USL	3
Standard Deviation	s	1	Lower Specification Limit	LSL	-3
Sample Size	n	30	Confidence Level (enter .95 for 95%)		95.0%
Sigma Shift (typically +1.5 for long term data)		1.5			
Process Sigma Level Results:			Process Capability Results:		
Expected dpm > USL		22,750.1	Cp, Pp		1.00
Expected % > USL		2.28%	Lower Limit Cp, Pp		0.74
Expected dpm < LSL		31.7	Upper Limit Cp, Pp		1.26
Expected % < LSL		0.00%	Cpu, Ppu		0.67
Expected dpm (overall)		22,781.8	Cpl, Ppl		1.33
Expected yield (overall)		97.72%	Cpk, Ppk		0.67
			Lower Limit Cpk, Ppk		0.46
Process Sigma Level		3.499	Upper Limit Cpk, Ppk		0.87

Notes for use of the Process Capability & Confidence Intervals Calculator:

1. This calculator assumes that the Mean and Standard Deviation are computed from data that are normally distributed.
2. Reports Cp, Cpk if entered S is Within or Short Term (using a control chart).
3. Reports Pp, Ppk if entered S is Overall or Long Term.

Basic Process Capability Templates – Tolerance Interval Calculator (Normal Exact)

Click **SigmaXL > Templates & Calculators > Basic Process Capability Templates > Tolerance Interval Calculator (Normal Exact)** to access the Tolerance Interval Calculator (Normal Exact). The template gives the following default example.

 Tolerance Interval Calculator (Normal Exact)			
Sample Data (user inputs):			
Sample Size	n		30
Sample Mean	x-bar		1
Sample Standard Deviation	s		1
Population Coverage (enter .99 for 99%)	100*(p)%		99.0%
Confidence Level (enter .95 for 95%)	100*(1- α)%		95.0%
Tolerance Interval Type			Two-Sided
Results:			
	Upper Tolerance Limit		4.3546
	Lower Tolerance Limit		-2.3546

Notes:

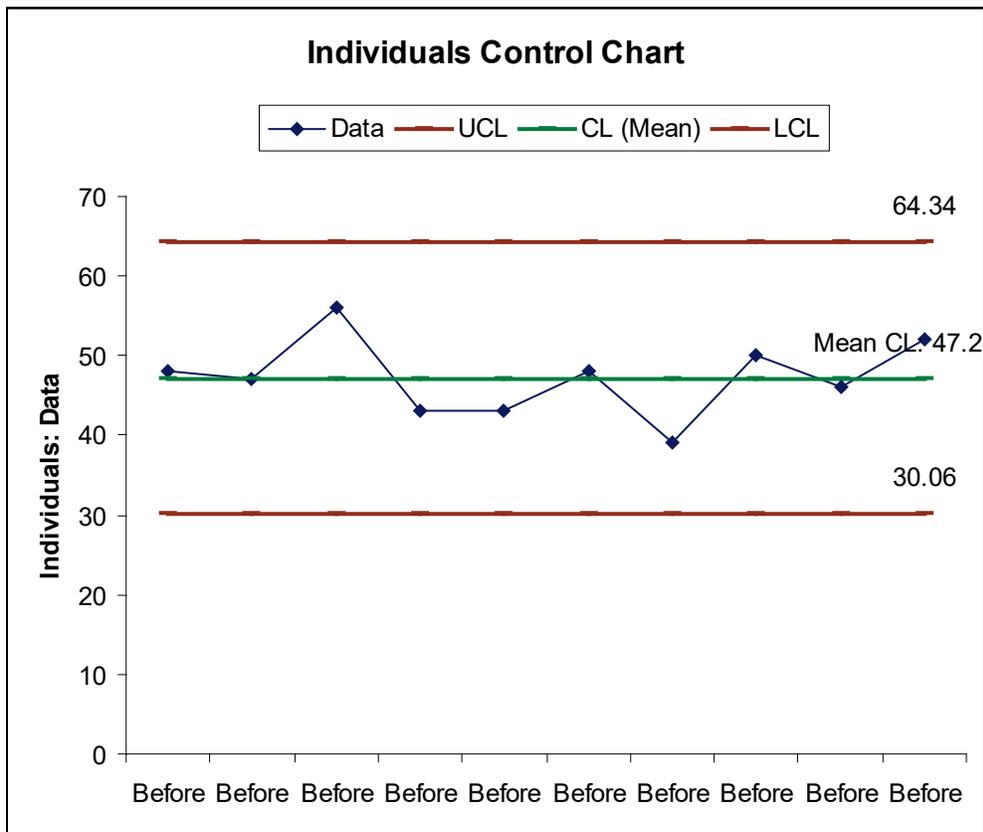
1. Enter summarized Sample Data, Population Coverage and Confidence Level in cells with yellow highlight.
2. Select Tolerance Interval Type using drop-down.
3. This calculator assumes that the summary statistics come from data that is normally distributed. Tolerance intervals are exact, one-sided uses the noncentral t distribution, two-sided uses the noncentral chi-square distribution. See Appendix [Tolerance Interval Calculator \(Normal Exact\)](#) for statistical details.
4. Due to the complexity of calculations, this template uses vba macros rather than Excel formulas. SigmaXL must be initialized and appear on the menu in order for this template to function.

Basic Control Chart Templates – Individuals Chart

Click **SigmaXL > Templates & Calculators > Control Chart Templates > Basic > Individuals**. Enter the data as shown (or copy and paste the “Before” data from the histogram data above):

Before/After	Data
Before	48
Before	47
Before	56
Before	43
Before	43
Before	48
Before	39
Before	50
Before	46
Before	52

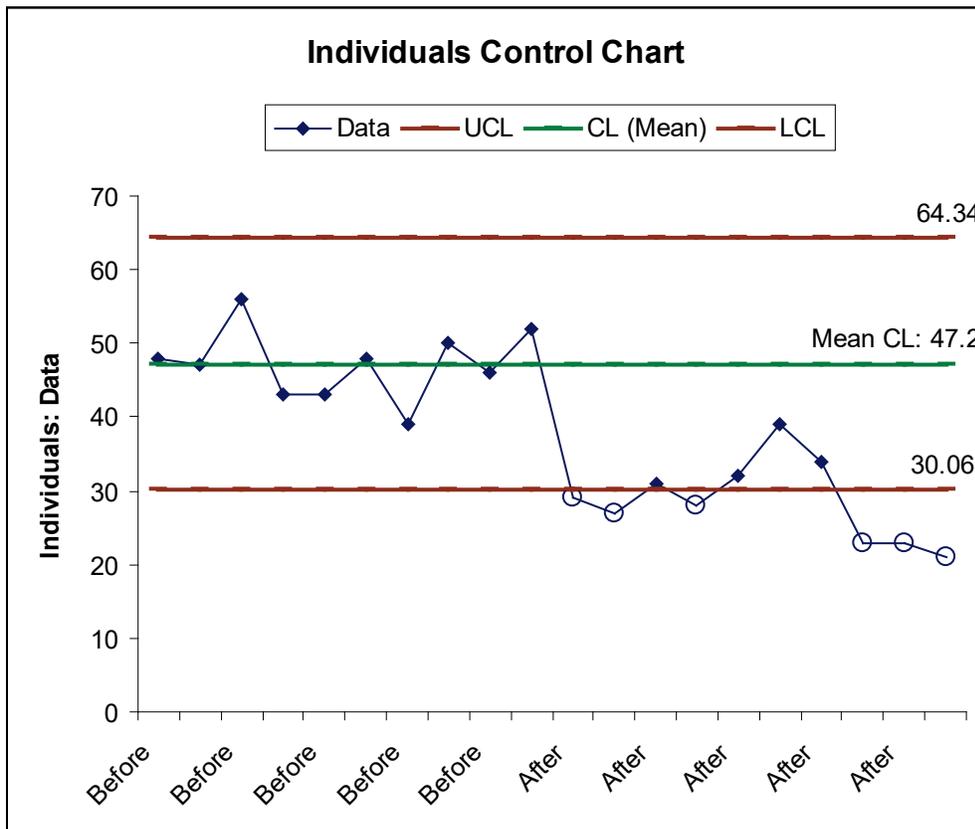
Click the Individuals Control Chart button to produce the Individuals Chart (note that in actual application, one should use a minimum of 20 to 30 data points to compute the initial control limits).



Now enter the data as shown (or copy and paste the “After” data from the histogram data used earlier):

Before/After	Data
Before	48
Before	47
Before	56
Before	43
Before	43
Before	48
Before	39
Before	50
Before	46
Before	52
After	29
After	27
After	31
After	28
After	32
After	39
After	34
After	23
After	23
After	21

Click the Add Data button to add the “After” data to the Individuals chart.



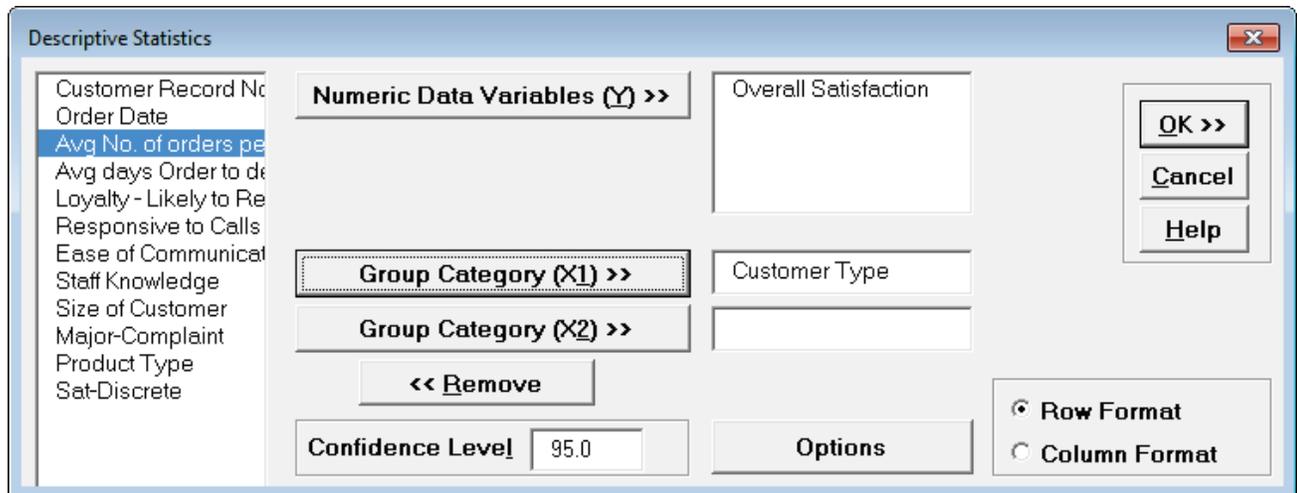
Notes:

1. This Individuals Control Chart template should be used with continuous data like cycle time. The data must be in chronological time-sequence order.
2. You can replace the **X-Axis Label** and **Data** column headings with any headings that you wish.
3. Enter your data in the **Data** column.
4. Enter labels in **X-Axis Label** column. Labels can be Date, Time, Name, or other text information. These labels are optional and will appear on the horizontal X-Axis of the Individuals Control Chart.
5. Click the **Individuals Control Chart** button to create an Individuals Control Chart.
6. After the control chart has been created and additional new data entered into the **Data** column, click the **Add Data** button to add the data to the existing chart.

Part C – Descriptive/Summary Statistics

Descriptive Statistics

1. Open **Customer Data.xlsx**. (To access, click **SigmaXL > Help > Sample Data** or **Start > Programs > SigmaXL > Sample Data**). Click **Sheet 1** Tab.
2. Click **SigmaXL > Statistical Tools > Descriptive Statistics**.
3. Check **Use Entire Data Table**, click **Next**.
4. Select **Overall Satisfaction**, click **Numeric Data Variables (Y) >>**. Select **Customer Type**, click **Group Category (X1) >>** as shown:



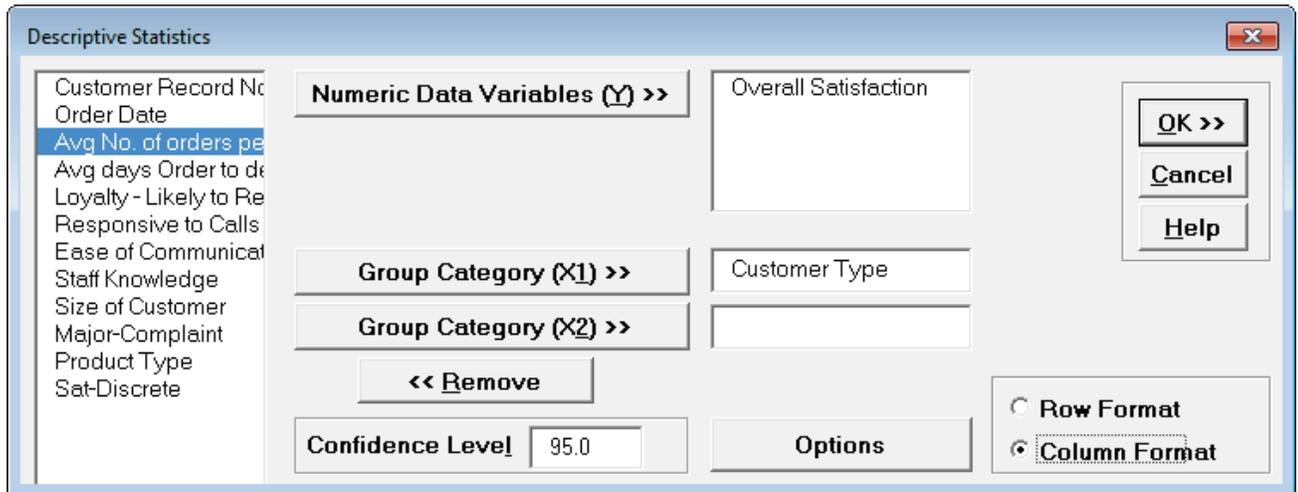
5. Click **OK**.
6. Descriptive Statistics are given for Customer Satisfaction grouped by Customer Type:

Overall Satisfaction by Customer Type			
Descriptive Statistics	Customer Type = 1	Customer Type = 2	Customer Type = 3
Count	31	42	27
Mean	3.394	4.205	3.641
Stdev	0.824680	0.621200	0.670478
Range	3.080	2.560	2.740
Minimum	1.720	2.420	2.190
25th Percentile (Q1)	2.810	3.828	3.240
50th Percentile (Median)	3.560	4.340	3.510
75th Percentile (Q3)	4.020	4.725	4.170
Maximum	4.800	4.980	4.930
95.0% CI Mean	3.091 to 3.696	4.012 to 4.399	3.376 to 3.906
95.0% CI Sigma	0.659012 to 1.102328	0.511126 to 0.792132	0.528013 to 0.918845
Anderson-Darling Normality Test	0.312776	0.826259	0.389291
p-value (A-D Test)	0.5306	0.0302	0.3600
Skewness	-0.235169	-0.967994	0.139571
p-value (Skewness)	0.5557	0.0121	0.7411
Kurtosis	-0.67169	0.679609	-0.313701
p-value (Kurtosis)	0.3705	0.2865	0.8435

Which Customer Type has the highest mean satisfaction score? Clearly Type 2. However, we have to be careful concluding that there is a significant difference in satisfaction solely by looking at the Means. In the Analyze Phase, we will run tests of hypothesis to validate that Type 2 Customers are, in fact, significantly more satisfied.

Tip: Click on **Column B**, click **View | Window | Split, Freeze Panes**. This freezes **Column A** and allows you to scroll across the Descriptive Statistics for each level of the Group Category. This is particularly beneficial when there are a large number of columns.

- Click **Recall SigmaXL Dialog** menu or press **F3** to recall last dialog. Change the format selected to **Column Format** as shown:



- Click **OK**. Descriptive Statistics are given for Customer Satisfaction broken out by Customer Type in Column Format:

Overall Satisfaction by Customer Type																	
Descriptive Statistic	Count	Mean	Stdev	Range	Minimum	25th Percentile (Q1)	50th Percentile (Median)	75th Percentile (Q3)	Maximum	95.0% CI Mean	95.0% CI Sigma	Anderson-Darling Normality Test	p-value (A-D Test)	Skewness	p-value (Skewness)	Kurtosis	p-value (Kurtosis)
Customer Type = 1	31	3.394	0.524683	3.083	1.725	2.810	3.565	4.025	4.805	3.091 to 3.698	0.659012 to 1.102328	0.312776	0.5308	-0.235169	0.5557	-0.671688	0.3705
Customer Type = 2	42	4.205	0.621266	2.568	2.420	3.528	4.340	4.725	4.985	4.012 to 4.395	0.511126 to 0.792132	0.626259	0.0302	-0.967994	0.0121	0.679608	0.2885
Customer Type = 3	27	3.641	0.670478	2.740	2.190	3.240	3.510	4.170	4.930	3.376 to 3.906	0.528013 to 0.916845	0.389291	0.3600	0.139571	0.7411	-0.313701	0.8435

Tip: This column format is useful to create subsequent graphs on the summary statistics.

Descriptive Statistics - Options

1. Click **Recall SigmaXL Dialog** menu or press **F3** to recall last dialog. Click **Options**. Check **Select All** and change Percentile Confidence Intervals to **Percentile** to display all Percentile values in the report.

The screenshot shows the 'Descriptive Statistics' dialog box. On the left, a list of variables includes 'Customer Record No', 'Order Date', 'Avg No. of orders per' (highlighted), 'Avg days Order to de', 'Loyalty - Likely to Re', 'Responsive to Calls', 'Ease of Communicat', 'Staff Knowledge', 'Size of Customer', 'Major-Complaint', 'Product Type', and 'Sat-Discrete'. The 'Numeric Data Variables (Y) >>' button is active, and 'Overall Satisfaction' is in the variable list. The 'Group Category (X1) >>' button is active, and 'Customer Type' is in the variable list. The 'Group Category (X2) >>' button is greyed out. The 'Confidence Level' is set to 95.0. The 'Options' button is active. The 'Row Format' radio button is selected, and the 'Column Format' radio button is greyed out. The 'Select All' checkbox is checked, and the 'Percentile Confidence Intervals' radio button is selected. The 'Interpolated' radio button is selected under 'Percentile Confidence Intervals', and the 'Interpolated' radio button is selected under 'Percentile Tolerance Intervals'.

Tip: Select only those options that are of interest in order to minimize the size of the report. Here we are selecting all options for demonstration purposes. Note that when any option is checked, **Row Format** is automatically selected, **Column Format** and **Group Category (X2)** are greyed out. These display options are limited due to the amount of information displayed in the extended report.

2. Click **OK**. Extended Descriptive Statistics are given for Customer Satisfaction grouped by Customer Type:

Overall Satisfaction by Customer Type			
Descriptive Statistics	Customer Type = 1	Customer Type = 2	Customer Type = 3
Count	31	42	27
Mean	3.394	4.205	3.641
Stdev	0.824680	0.621200	0.670478
Range	3.080	2.560	2.740
Minimum	1.720	2.420	2.190
25th Percentile (Q1)	2.810	3.828	3.240
50th Percentile (Median)	3.560	4.340	3.510
75th Percentile (Q3)	4.020	4.725	4.170
Maximum	4.800	4.980	4.930
95.0% CI Mean	3.0911 to 3.696	4.0117 to 4.3988	3.3759 to 3.9063
95.0% CI Sigma	0.65901 to 1.1023	0.51113 to 0.79213	0.52801 to 0.91884
Anderson-Darling Normality Test	0.312776	0.826259	0.389291
P-Value (A-D Test)	0.5306	0.0302	0.3600
Skewness	-0.235169	-0.967994	0.139571
P-Value (Skewness)	0.5557	0.0121	0.7411
Kurtosis	-0.671690	0.679609	-0.313701
P-Value (Kurtosis)	0.3705	0.2865	0.8435
Additional Descriptive Statistics			
5% Trimmed Mean	3.413	4.251	3.648
Standard Error of Mean	0.148117	0.095853182	0.129034
Variance	0.680097	0.385889	0.449541
Coefficient of Variation	24.301	14.772	18.414
StDev (Within, Short Term)	0.678487	0.521969	0.683988
Additional Normality Tests			
Shapiro-Wilk/KSL	0.969383	0.923744	0.971664
P-Value (Shapiro-Wilk/KSL)	0.5022	0.0080	0.6463
Doornik-Hansen	0.813602	8.554	0.127399
P-Value (Doornik-Hansen)	0.6658	0.0139	0.9383

3. The **Additional Descriptive Statistics** are:

- 5% Trimmed Mean. The highest 5% and lowest 5% are excluded and mean calculated with the rest of the data. This gives a robust alternative to the Median as a measure of central tendency in the presence of outliers.
- Standard Error of Mean ($StDev/\sqrt{N}$)
- Variance ($StDev^2$)
- Coefficient of Variation ($100 * StDev/Mean$)
- Short Term StDev ($MR\text{-}bar/d2$)

4. The **Additional Normality Tests** are:

- Shapiro-Wilk ($n \leq 5000$) and Kolmogorov-Smirnov-Lilliefors (KSL, $n > 5000$)
 - This is a popular alternative to Anderson Darling.
- Doornik-Hansen (DH)
 - Univariate omnibus test based on Skewness and Kurtosis. (Note, the bivariate DH test is used in [Correlation Matrix](#) to test bivariate normality).
 - Best for data with ties, i.e. “chunky” data. Anderson-Darling, Shapiro-Wilk and KSL are severely affected by ties in the data and will trigger a low P-Value even if the data are normal.
 - See Appendix [Doornik-Hansen \(DH\) Normality Test](#) for further details and references.

Overall Satisfaction by Customer Type			
Descriptive Statistics	Customer Type = 1	Customer Type = 2	Customer Type = 3
Percentile Report			
0.135	1.72	2.42	2.19
0.5	1.72	2.42	2.19
1	1.72	2.42	2.19
2.5	1.72	2.441	2.19
5	1.804	2.778	2.37
10	2.138	3.285	2.832
15	2.554	3.518	3.056
20	2.602	3.65	3.144
25	2.81	3.8275	3.24
30	2.924	3.981	3.272
35	2.97	4.091	3.298
40	3.074	4.132	3.33
45	3.248	4.3035	3.422
50	3.56	4.34	3.51
55	3.576	4.4195	3.566
60	3.804	4.464	3.644
65	3.9	4.5275	3.988
70	3.978	4.558	4.074
75	4.02	4.725	4.17
80	4.11	4.786	4.294
85	4.242	4.881	4.428
90	4.586	4.91	4.688
95	4.722	4.9655	4.878
97.5	4.8	4.97925	4.93
99	4.8	4.98	4.93
99.5	4.8	4.98	4.93
99.865	4.8	4.98	4.93
75 - 25 (50%, Interquartile Range IQR)	1.21	0.8975	0.93
90 - 10 (80%, Interdecile Range IDR)	2.448	1.625	1.856
95 - 5 (90%, Span)	2.918	2.1875	2.508
97.5 - 2.5 (95%, +/- 1.96 Sigma Equivalent)	3.08	2.53825	2.74
99 - 1 (98%)	3.08	2.56	2.74
99.5 - 0.5 (99%)	3.08	2.56	2.74
99.865 - 0.135 (99.73%, +/- 3 Sigma Equivalent)	3.08	2.56	2.74

5. The **Percentile Report** gives 27 values from 0.135 to 99.865.
6. The **Percentile Ranges** are:
 - 75 - 25 (50%, Interquartile Range IQR)
 - 90 - 10 (80%, Interdecile Range IDR)
 - 95 - 5 (90%, Span)
 - 97.5 - 2.5 (95%, +/- 1.96 Sigma Equivalent)
 - 99 - 1 (98%)
 - 99.5 - 0.5 (99%)
 - 99.865 - 0.135 (99.73%, +/- 3 Sigma Equivalent)

Overall Satisfaction by Customer Type			
Descriptive Statistics	Customer Type = 1	Customer Type = 2	Customer Type = 3
Percentile Confidence Intervals (Interpolated 95.0%)			
0.135	Min. sample size = 2731	Min. sample size = 2731	Min. sample size = 2731
0.5	Min. sample size = 736	Min. sample size = 736	Min. sample size = 736
1	Min. sample size = 368	Min. sample size = 368	Min. sample size = 368
2.5	Min. sample size = 146	Min. sample size = 146	Min. sample size = 146
5	Min. sample size = 72	Min. sample size = 72	Min. sample size = 72
10	Min. sample size = 36	2.5826 to 3.6589	Min. sample size = 36
15	1.8393 to 2.9095	3.1206 to 3.8676	2.4389 to 3.2707
20	2.0087 to 2.9711	3.2747 to 4.0232	2.7108 to 3.3014
25	2.4809 to 3.1075	3.4847 to 4.1163	2.9143 to 3.3899
30	2.5606 to 3.2586	3.5891 to 4.3005	3.0606 to 3.4556
35	2.6047 to 3.5658	3.7308 to 4.3359	3.1354 to 3.5275
40	2.7749 to 3.6297	3.8487 to 4.4153	3.2052 to 3.6277
45	2.906 to 3.817	3.9942 to 4.4601	3.2554 to 3.862
50	2.9535 to 3.9362	4.0946 to 4.5184	3.2891 to 4.0227
55	3.0221 to 3.987	4.1397 to 4.5458	3.3036 to 4.1239
60	3.2061 to 4.0266	4.3049 to 4.6959	3.3845 to 4.1874
65	3.387 to 4.1066	4.3441 to 4.7459	3.4401 to 4.3391
70	3.5707 to 4.2345	4.4284 to 4.8516	3.5236 to 4.4274
75	3.7743 to 4.2921	4.4737 to 4.8928	3.626 to 4.6107
80	3.8978 to 4.67	4.53 to 4.91	3.9334 to 4.7587
85	3.9853 to 4.6892	4.6717 to 4.9457	4.0779 to 4.8581
90	Min. sample size = 36	4.77 to 4.9742	Min. sample size = 36
95	Min. sample size = 72	Min. sample size = 72	Min. sample size = 72
97.5	Min. sample size = 146	Min. sample size = 146	Min. sample size = 146
99	Min. sample size = 368	Min. sample size = 368	Min. sample size = 368
99.5	Min. sample size = 736	Min. sample size = 736	Min. sample size = 736
99.865	Min. sample size = 2731	Min. sample size = 2731	Min. sample size = 2731
Percentile Tolerance Intervals (Interpolated 95.0%)			
50%	2.56 to 4.1129	3.62 to 4.8262	3.04 to 4.2764
80%	1.7453 to 4.67	3.108 to 4.97	2.4552 to 4.93
90%	Min. sample size = 46	Min. sample size = 46	Min. sample size = 46
95%	Min. sample size = 93	Min. sample size = 93	Min. sample size = 93
98%	Min. sample size = 236	Min. sample size = 236	Min. sample size = 236
99%	Min. sample size = 473	Min. sample size = 473	Min. sample size = 473
99.73%	Min. sample size = 1756	Min. sample size = 1756	Min. sample size = 1756

7. The **Percentile Confidence Intervals** give 27 values from 0.135 to 99.865.
8. The **Quartile Confidence Intervals** give 3 values: 25, 50 and 75.
9. The **Percentile Tolerance Intervals** are 50%, 80%, 90%, 95%, 98%, 99%, and 99.73%.
10. Confidence Intervals and Tolerance Intervals can be exact or interpolated. If exact, the actual exact confidence level will be a value greater than or equal to specified, due to percentile values being discrete in nature. The actual exact level will also be reported in this case. If interpolated, the result will be an interpolated estimate of the specified confidence level (typically 95.0%) and is the recommended setting. See Appendix [Percentile \(Nonparametric\) Confidence and Tolerance Intervals](#) for further details.

Tip: The Tolerance Intervals given here are nonparametric, so the data does not have to be normal. However, if you have normal data, you can use the Tolerance Interval Calculator for Normal data, which will allow smaller sample sizes to be used. See [Basic Process Capability Templates – Tolerance Interval Calculator \(Normal Exact\)](#).

11. If the Confidence Interval or Tolerance Interval cannot be computed due to inadequate sample size, a minimum sample size is reported.

Overall Satisfaction by Customer Type			
Descriptive Statistics	Customer Type = 1	Customer Type = 2	Customer Type = 3
Outlier and Randomness Tests			
Outliers (Boxplot Rules)	No outliers found.	Potential (1.5*IQR) outlier lower count = 1.	No outliers found.
Grubbs Outlier Test	Grubbs' Test P-Value = 1.000. Fail to reject null hypothesis: "There are no outliers in the data set." Note that Grubbs' Test assumes normality and tests only if the maximum or minimum is an outlier.	Grubbs' Test P-Value = 0.105. Fail to reject null hypothesis: "There are no outliers in the data set." Note that Grubbs' Test assumes normality and tests only if the maximum or minimum is an outlier.	Grubbs' Test P-Value = 0.654. Fail to reject null hypothesis: "There are no outliers in the data set." Note that Grubbs' Test assumes normality and tests only if the maximum or minimum is an outlier.
Randomness Runs Test	Nonparametric Runs Test (Exact) P-Value = 0.066. Fail to reject null hypothesis: "data are random," so conclude that the assumption of randomness (independence) is not violated.	Nonparametric Runs Test (Exact) P-Value = 1.000. Fail to reject null hypothesis: "data are random," so conclude that the assumption of randomness (independence) is not violated.	Nonparametric Runs Test (Exact) P-Value = 1.000. Fail to reject null hypothesis: "data are random," so conclude that the assumption of randomness (independence) is not violated.

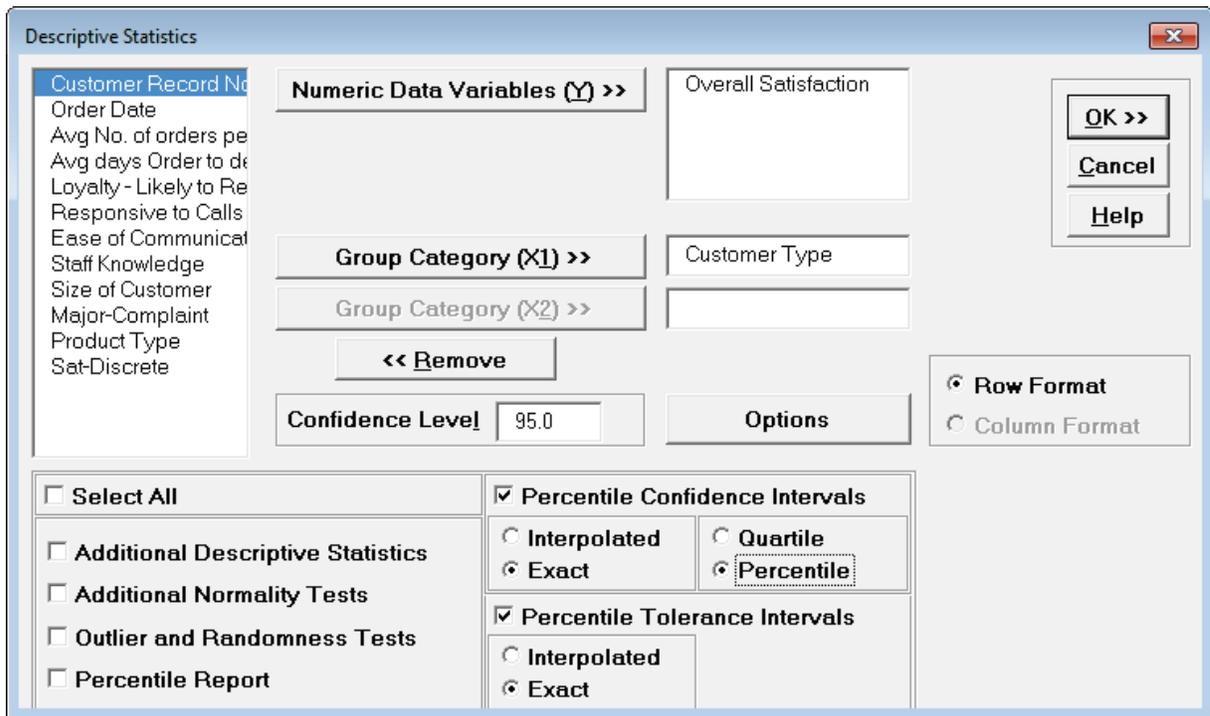
12. The **Outlier (Boxplot Rules)** Tests are: Potential 1.5(IQR), Likely 2.2(IQR), Extreme 3.0(IQR)

13. **Grubbs Outlier Test** is more powerful at detecting a single outlier as maximum or minimum but assumes that the remainder of the data are normally distributed.

14. The **Randomness Runs Test** is a nonparametric exact runs test.

15. The Outlier and Randomness Tests use the same Green, Yellow, Red highlight given in the automatic assumptions report that are included in t-tests and ANOVA.

16. Click **Recall SigmaXL Dialog** menu or press **F3** to recall last dialog. Click **Options**. Uncheck **Select All** to clear the selections and check **Percentile Confidence Intervals**, select **Exact** and **Percentile**. Check **Percentile Tolerance Intervals**, and select **Exact** as shown:



17. Click **OK**. The Percentile Confidence Intervals and Tolerance Intervals are displayed:

Overall Satisfaction by Customer Type			
Descriptive Statistics	Customer Type = 1	Customer Type = 2	Customer Type = 3
Percentile Confidence Intervals (Exact 95.0% Minimum)			
0.135	Min. sample size = 2731	Min. sample size = 2731	Min. sample size = 2731
0.5	Min. sample size = 736	Min. sample size = 736	Min. sample size = 736
1	Min. sample size = 368	Min. sample size = 368	Min. sample size = 368
2.5	Min. sample size = 146	Min. sample size = 146	Min. sample size = 146
5	Min. sample size = 72	Min. sample size = 72	Min. sample size = 72
10	Min. sample size = 36	2.42 to 3.67 (96.7%)	Min. sample size = 36
15	1.72 to 2.94 (98.1%)	2.7 to 3.9 (97.3%)	2.19 to 3.29 (97.4%)
20	1.86 to 3.01 (97.9%)	3.24 to 4.09 (96.6%)	2.64 to 3.31 (97%)
25	2.04 to 3.24 (98%)	3.39 to 4.12 (96.9%)	2.88 to 3.41 (95.8%)
30	2.56 to 3.26 (95.2%)	3.54 to 4.31 (97.2%)	3.04 to 3.51 (96.6%)
35	2.57 to 3.57 (96.4%)	3.67 to 4.35 (96.6%)	3.12 to 3.53 (95.9%)
40	2.65 to 3.8 (97.4%)	3.83 to 4.43 (96.1%)	3.16 to 3.65 (97.1%)
45	2.9 to 3.82 (95.5%)	3.99 to 4.47 (95.7%)	3.24 to 3.98 (96.8%)
50	2.94 to 3.97 (97.1%)	4.09 to 4.53 (95.6%)	3.26 to 4.11 (98.1%)
55	3.01 to 3.99 (95.5%)	4.12 to 4.55 (95.7%)	3.3 to 4.17 (96.8%)
60	3.09 to 4.05 (97.4%)	4.3 to 4.72 (96.1%)	3.31 to 4.21 (97.1%)
65	3.26 to 4.15 (96.4%)	4.33 to 4.75 (96.6%)	3.43 to 4.42 (95.9%)
70	3.57 to 4.24 (95.2%)	4.4 to 4.87 (97.2%)	3.51 to 4.43 (96.6%)
75	3.58 to 4.67 (98%)	4.47 to 4.91 (96.9%)	3.62 to 4.66 (95.8%)
80	3.82 to 4.67 (97.9%)	4.53 to 4.91 (96.6%)	3.65 to 4.8 (97%)
85	3.97 to 4.8 (98.1%)	4.63 to 4.97 (97.3%)	4.02 to 4.93 (97.4%)
90	Min. sample size = 36	4.75 to 4.98 (96.7%)	Min. sample size = 36
95	Min. sample size = 72	Min. sample size = 72	Min. sample size = 72
97.5	Min. sample size = 146	Min. sample size = 146	Min. sample size = 146
99	Min. sample size = 368	Min. sample size = 368	Min. sample size = 368
99.5	Min. sample size = 736	Min. sample size = 736	Min. sample size = 736
99.865	Min. sample size = 2731	Min. sample size = 2731	Min. sample size = 2731
Percentile Tolerance Intervals (Exact 95.0% Minimum)			
50%	2.56 to 4.15 (96.5%)	3.62 to 4.84 (95.6%)	3.04 to 4.42 (97.4%)
80%	1.72 to 4.67 (96.3%)	2.7 to 4.97 (97.9%)	2.19 to 4.93 (98.1%)
90%	Min. sample size = 46	Min. sample size = 46	Min. sample size = 46
95%	Min. sample size = 93	Min. sample size = 93	Min. sample size = 93
98%	Min. sample size = 236	Min. sample size = 236	Min. sample size = 236
99%	Min. sample size = 473	Min. sample size = 473	Min. sample size = 473
99.73%	Min. sample size = 1756	Min. sample size = 1756	Min. sample size = 1756

18. The specified 95% is a guaranteed minimum. The exact confidence level is given with each reported interval.

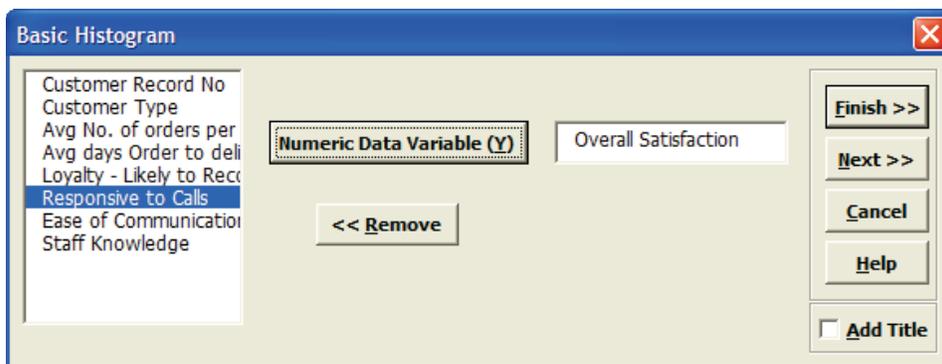
Part D – Histograms

Basic Histogram Template

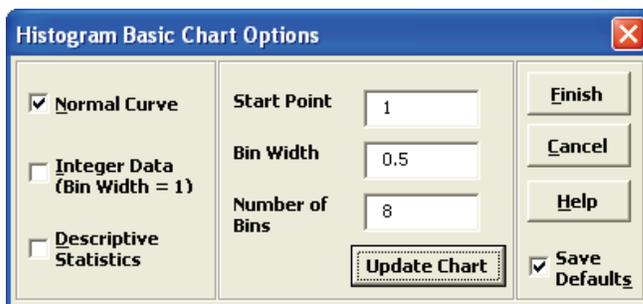
Click **SigmaXL > Templates and Calculators > Basic Graphical Templates > Histogram** or **SigmaXL > Graphical Tools > Basic Graphical Templates > Histogram**. See **Part B – Templates and Calculators** for a Histogram Template example.

Single (Basic) Histogram

1. Click **Sheet 1** Tab of **Customer Data.xlsx** (or press **F4** to activate last worksheet).
2. Click **SigmaXL > Graphical Tools > Basic Histogram**.
3. Ensure that entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
4. Select *Overall Satisfaction*, click **Numeric Data Variable (Y)** >> as shown:

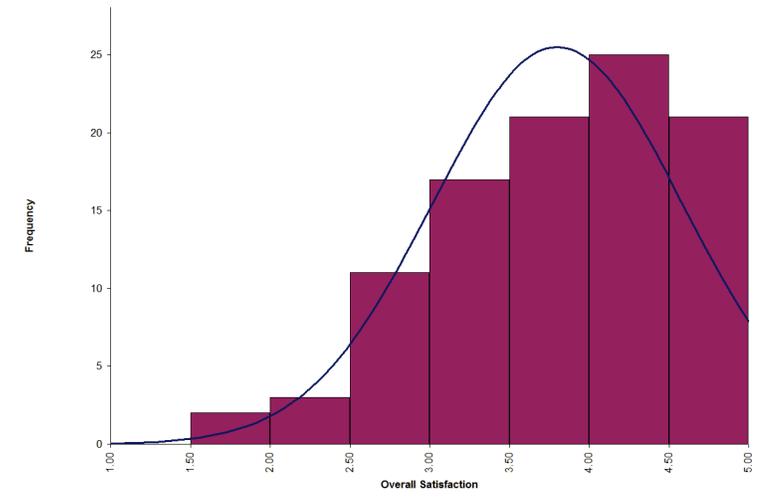


5. Click **Next**. Ensure that **Normal Curve** is checked. Set **Start Point** = 1. Change the **Bin Width** to 0.5, and the **Number of Bins** to 8. Click **Update Chart** to view the histogram. (If the survey satisfaction data was pure integer format, we would have checked the **Integer Data** option).



Note: For SigmaXL Mac version, click **SigmaXL > Graphical Tools > Basic Histogram Options** to set options. The Histogram Basic Chart is displayed after you click **Finish**.

6. Click **Finish**. A histogram of Overall Customer Satisfaction is produced.

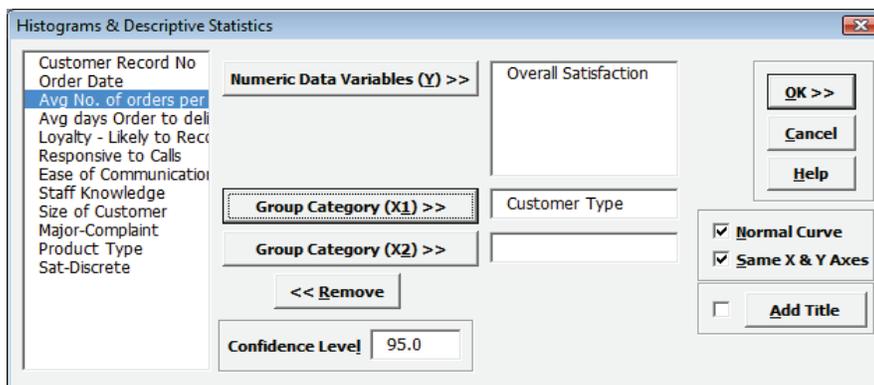


7. Note that bin one is 1 to < 1.5, bin 2 is 1.5 to < 2, etc.

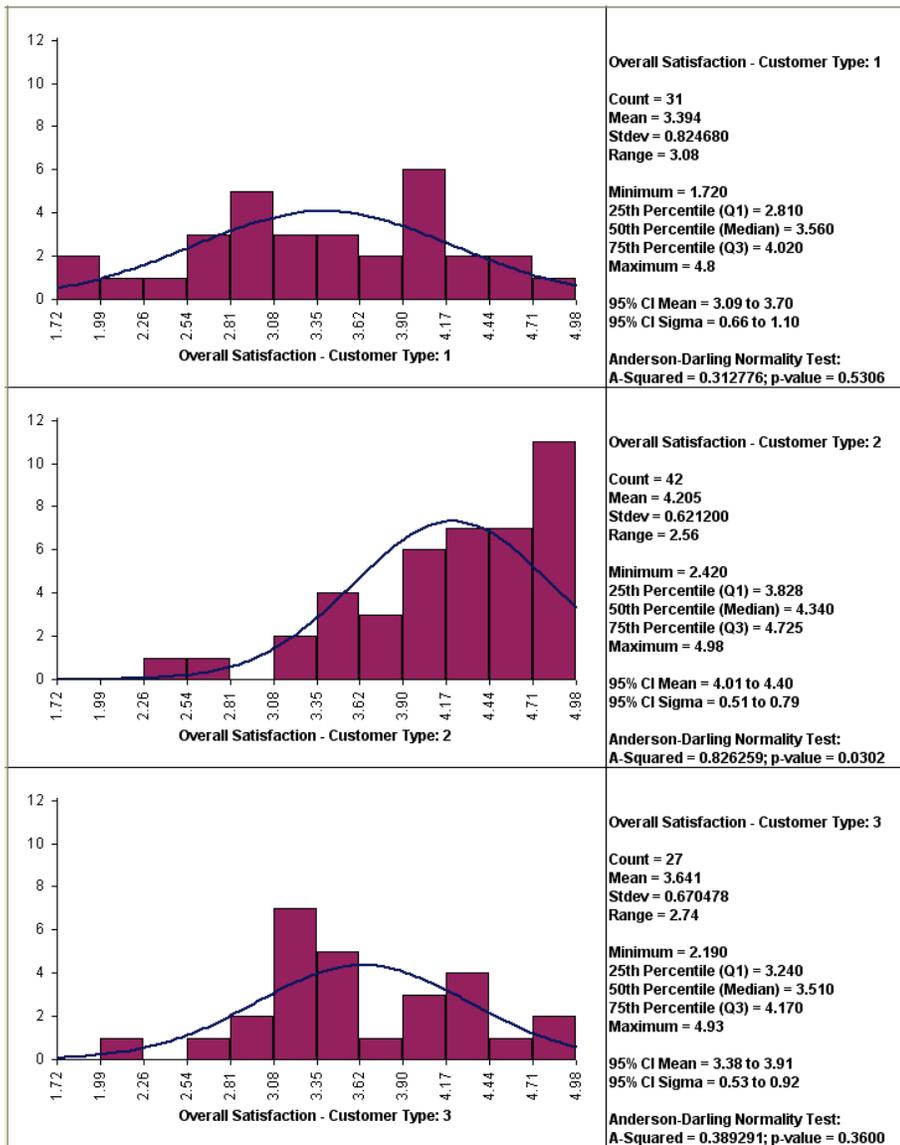
Tip: Any graph produced by SigmaXL can be Copied/Pasted into Word. It can also be enlarged by clicking on the graph and dragging the corner. The number of decimal places displayed can be modified by clicking on the Axis Label and selecting the **Number** tab to adjust. The text label alignment can also be modified by: Select **Axis**, Right Click, **Format Axis**.

Multiple Histograms

1. Click **Sheet 1** Tab of **Customer Data.xlsx** (or press **F4** to activate last worksheet).
2. Click **SigmaXL > Graphical Tools > Histograms & Descriptive Statistics**.
3. Ensure that entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
4. Select *Overall Satisfaction*, click **Numeric Data Variables (Y) >>**, select *Customer Type*, click **Group Category (X1) >>** as shown:



- Click **OK**. Multiple Histograms and Descriptive Statistics of Customer Satisfaction By Customer Type are produced:



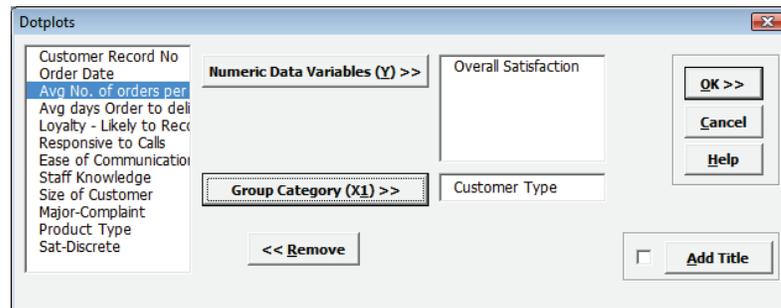
Clearly Customer Type 2 shows a higher level of overall satisfaction, with the data skewed left. Note that Customer Type 1 and 3 have data that is normally distributed, but this is not desirable when the response is a satisfaction score!

- Note that bin one is 1.72 to < 1.99, bin 2 is 1.99 to < 2.26, etc. The number of decimals displayed can be changed by double-clicking on the X axis, click Number tab, and adjust decimal places.

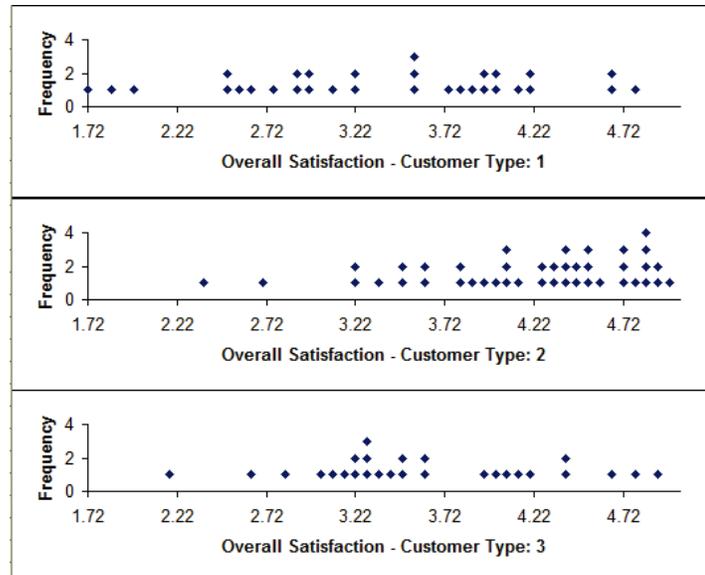
Part E – Dotplots

Dotplots

1. Click Sheet 1 Tab of **Customer Data.xlsx** (or press **F4** to activate last worksheet). Click **SigmaXL > Graphical Tools > Dotplots**.
2. Ensure that entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
3. Select *Overall Satisfaction*, click **Numeric Data Variable (Y) >>**, select *Customer Type*, click **Group Category (X1) >>**:



4. Click **OK**. Multiple dotplots of Customer Satisfaction By Customer Type are produced:

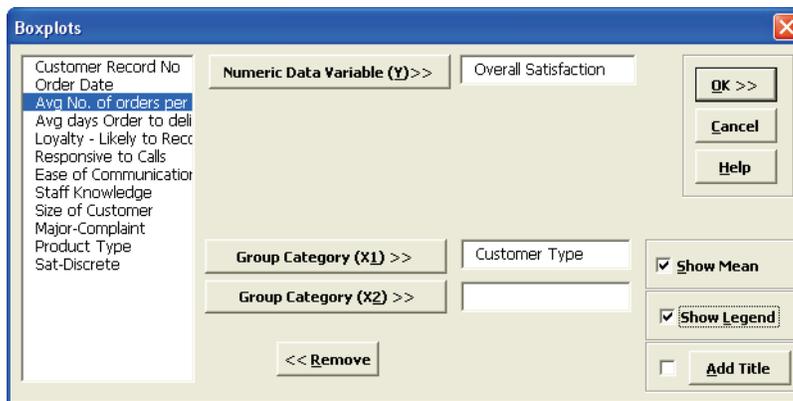


5. Since all data points are shown, dotplots are a useful alternative to Histograms, particularly when the group sample sizes are small ($n < 30$).

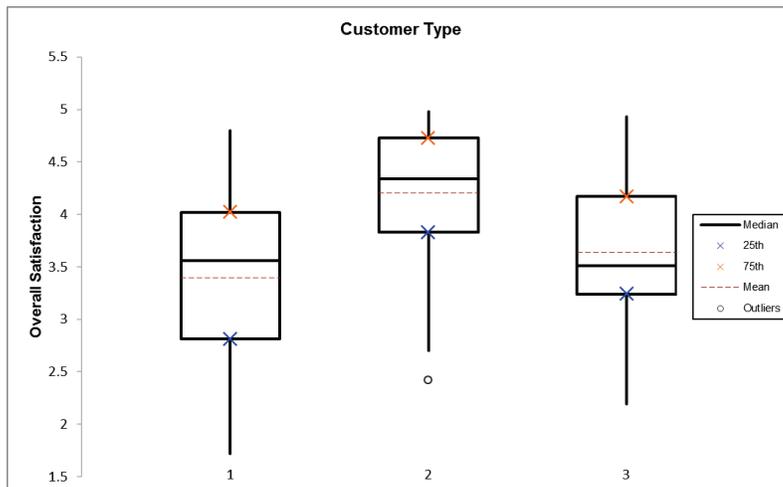
Part F – Boxplots

Boxplots

1. Click **Sheet 1** Tab of **Customer Data.xlsx** (or press **F4** to activate last worksheet).
2. Click **SigmaXL > Graphical Tools > Boxplots**.
3. Ensure that entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
4. Select *Overall Satisfaction*, click **Numeric Data Variable (Y) >>**, select *Customer Type*, click **Group Category (X1) >>**, check **Show Mean**, check **Show Legend**:

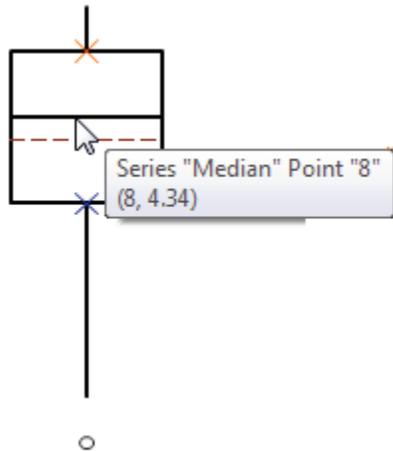


5. Click **OK**. A boxplot of Customer Satisfaction By Customer Type is produced:



6. The legend indicates that the solid center line is the median. The dashed red line shows the sample mean. The top of the box is the 75th percentile (Q3). The bottom of the box is the 25th percentile (Q1). The height of the box is called the Inter-Quartile Range (IQR) and is a robust measure of spread or sample variability. The data point highlighted for Customer Type 2 is a potential outlier ($< Q1 - 1.5 * IQR$ or $> Q3 + 1.5 * IQR$). Note that extreme outliers are highlighted with a solid dot ($< Q1 - 3 * IQR$ or $> Q3 + 3 * IQR$).

Tip: If you hover your mouse cursor in the middle of any of the Boxplot lines as shown, you will see the balloon help indicating what the line is and its numeric Y value.



- Now we will modify the Y axis scale, showing 1 as minimum and 5 as maximum (given that the response data comes from a survey with 1-5 scale). To do this right click on the Y axis and select “Format Axis”, modify the minimum value and maximum value. Change **Horizontal axis crosses** > **Axis value** to 1 as shown:

Format Axis ▼ ×

AXIS OPTIONS ▼ | TEXT OPTIONS

4 AXIS OPTIONS

Bounds

Minimum

Maximum

Units

Major

Minor

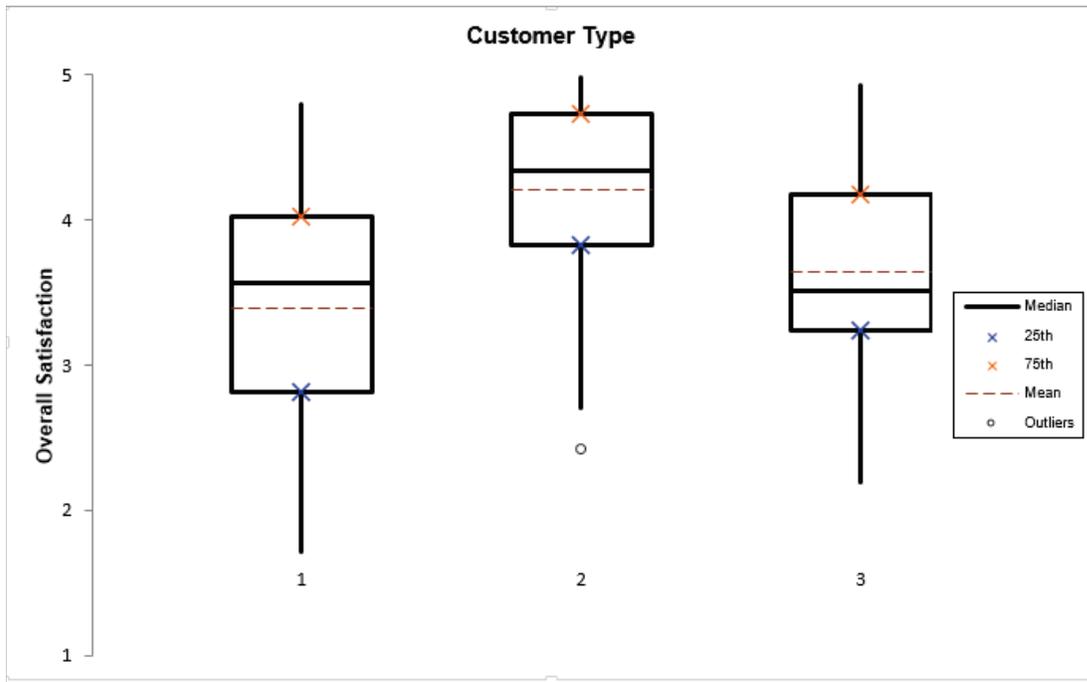
Horizontal axis crosses

Automatic
 Axis value
 Maximum axis value

Display units ▼

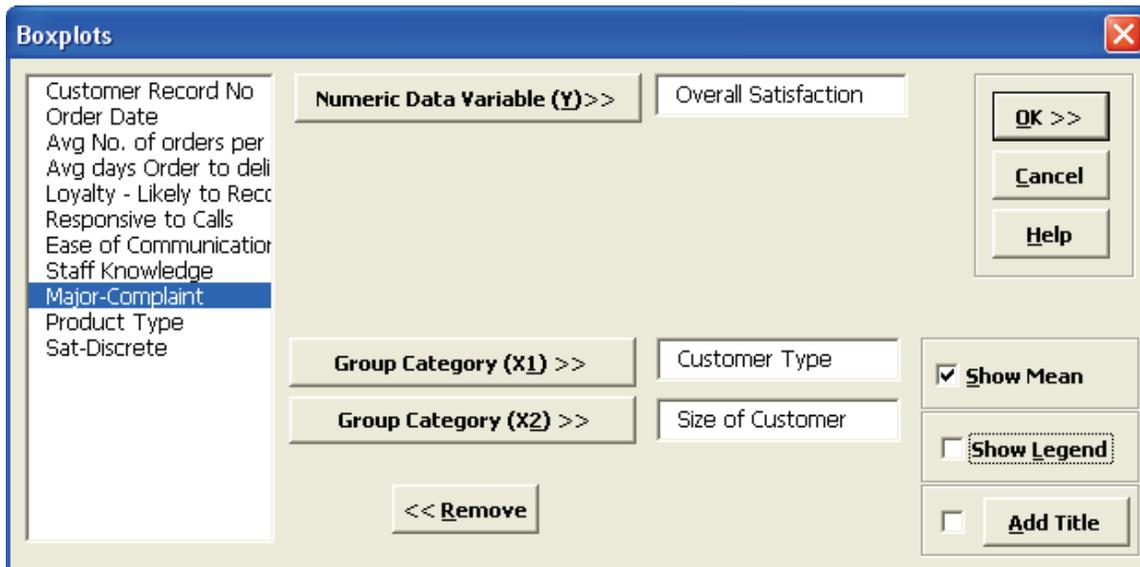
Show display units label on chart
 Logarithmic scale Base
 Values in reverse order

8. The Boxplot axis is modified as shown below:

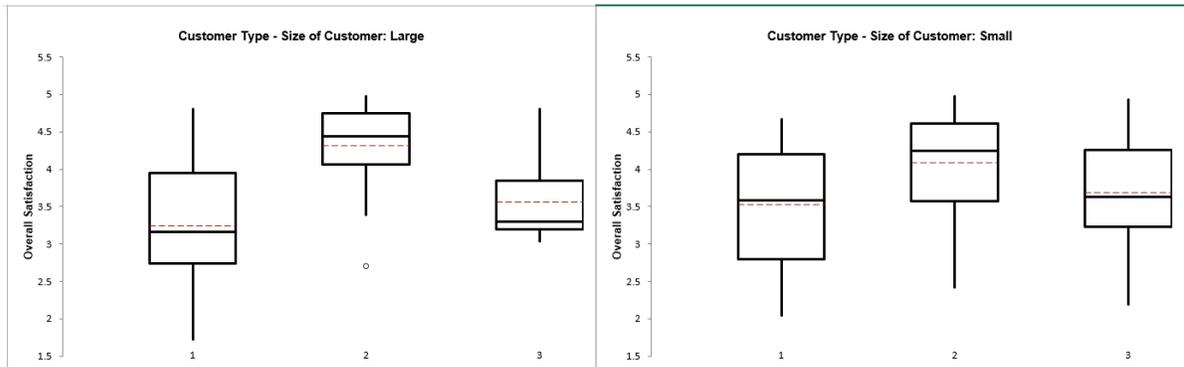


9. Click **Recall SigmaXL Dialog** menu or press **F3** to recall last dialog.

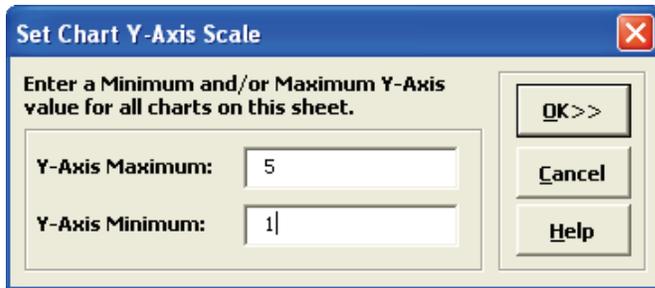
10. Select *Overall Satisfaction*, click **Numeric Data Variable (Y) >>**; select *Customer Type*, click **Group Category (X1) >>**; select *Size of Customer*, click **Group Category (X2) >>**; check **Show Mean**; uncheck **Show Legend**:



11. Click **OK**. Boxplots of Customer Satisfaction By Customer Type and Size are produced:



12. In order to adjust the Y-axis scale for both charts, click **SigmaXL Chart Tools > Set Chart Y-Axis Max/Min**.



13. Click **OK**. The Y-axis scale maximum and minimum are now modified for both charts.

Part G – Normal Probability Plots

Normal Probability Plots

1. Create 100 random normal values as follows: Click **SigmaXL > Data Manipulation > Random Data > Normal**. Specify 1 Column, 100 Rows, Mean of 100 and Standard Deviation of 25 as shown below:

Normal Random Number Generator

Start Cell: \$A\$1

Select Distribution: Normal

Number of Columns: 1

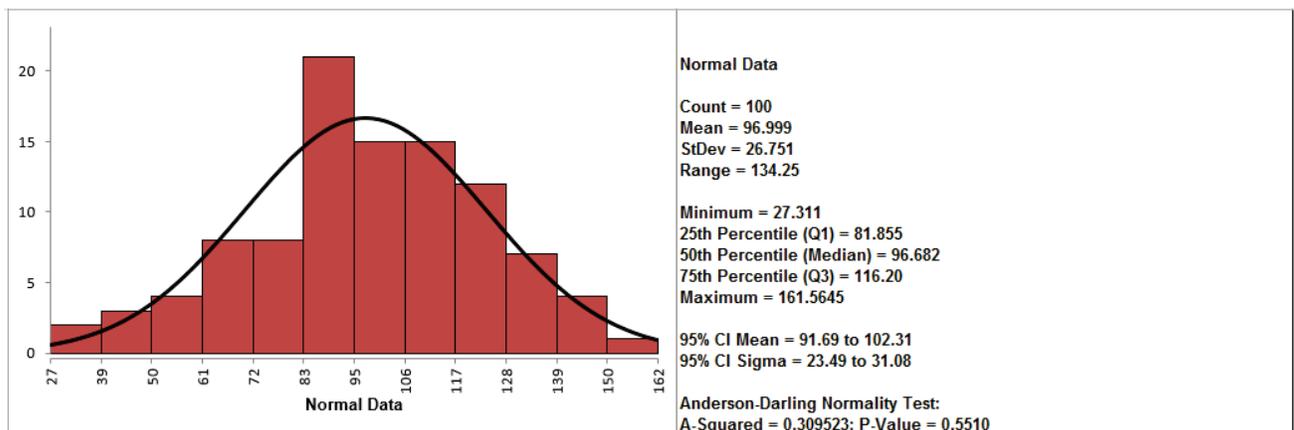
Number of Rows: 100

Mean: 100

Standard Deviation: 25

OK >> Cancel Help

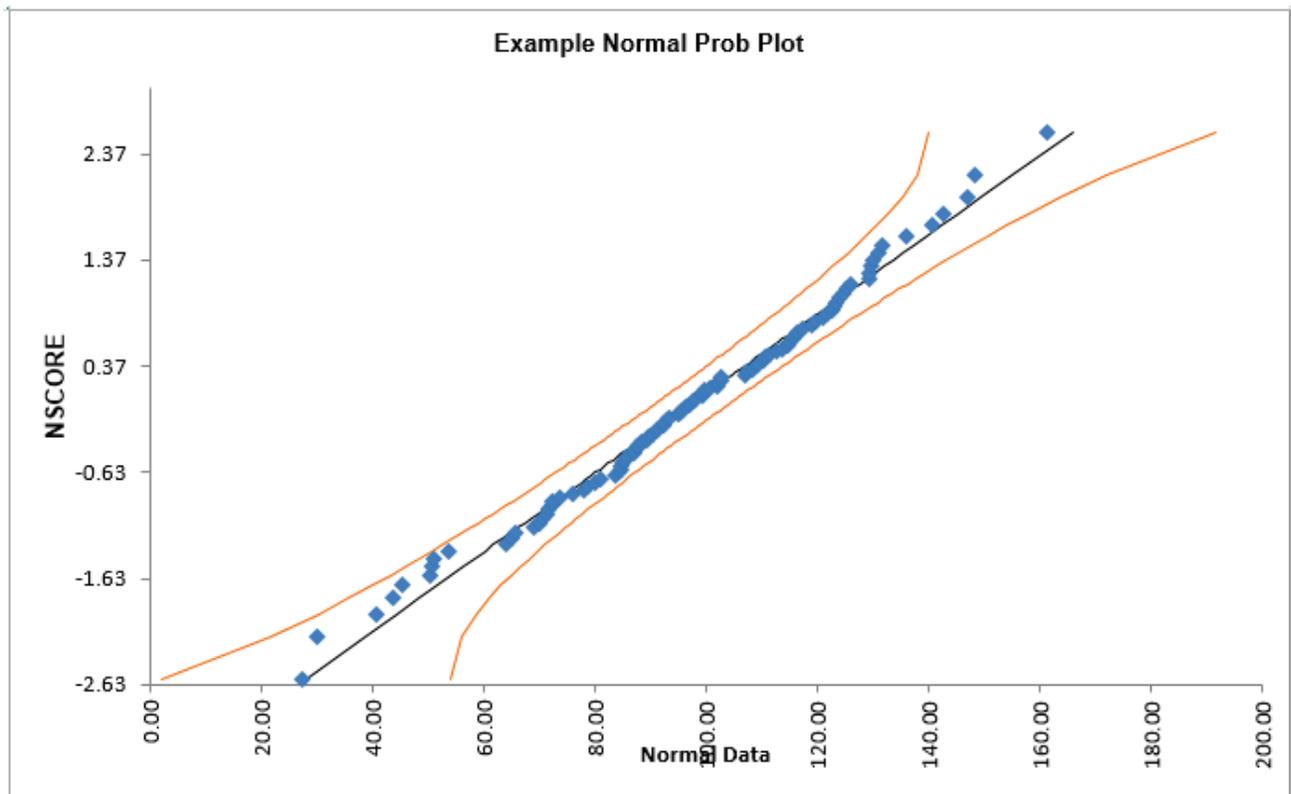
2. Click **OK**. Change Column heading to *Normal Data*.
3. Create a Histogram & Descriptive Statistics for this data. Your data will be slightly different due to the random number generation:



If the P-Value of the Anderson-Darling Normality test is greater than or equal to .05, the data is considered to be normal (interpretation of P-Values will be discussed further in Analyze).

4. Create a normal probability plot of this data: Click **SigmaXL Random Data (1) Sheet**, Click **SigmaXL > Graphical Tools > Normal Probability Plots**.
5. Ensure that entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
6. Select *Normal Data*, click **Numeric Data Variable (Y) >>**. Check **Add Title**. Enter *Example Normal Prob Plot*. Click **OK**.

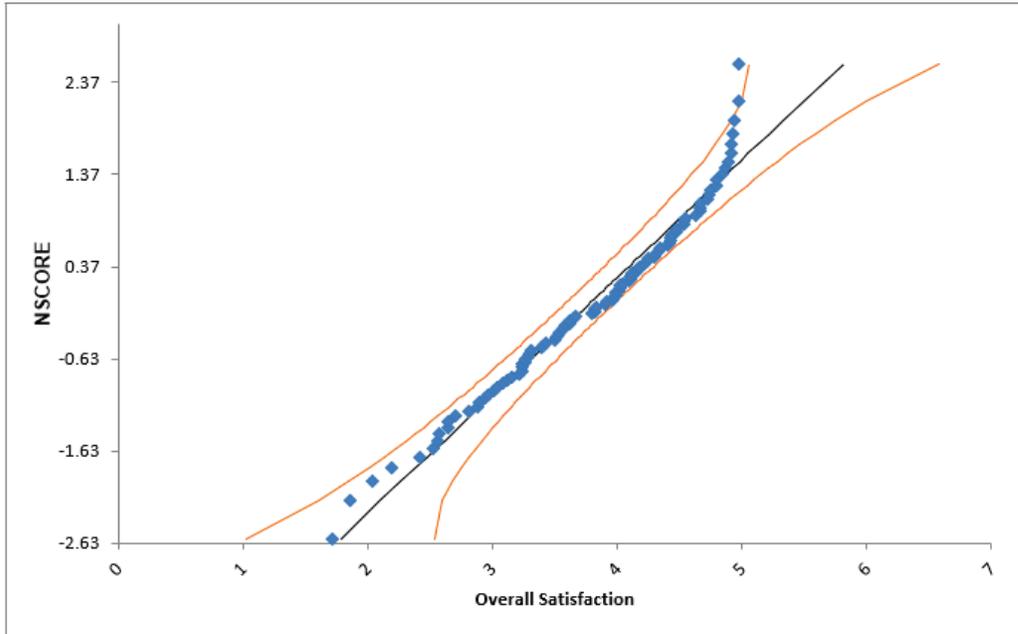
- Click **OK**. A Normal Probability Plot of simulated random data is produced (again, your plot will be slightly different due to the random number generation):



The data points follow the straight line fairly well, indicating that the data is normally distributed. Note that the data will not likely fall in a perfectly straight line. The eminent statistician George Box uses a “Fat Pencil” test where the data, if covered by a fat pencil, can be considered normal! We can also see that the data is normal since the points fall within the normal probability plot 95% confidence intervals (confidence intervals will be discussed further in Analyze).

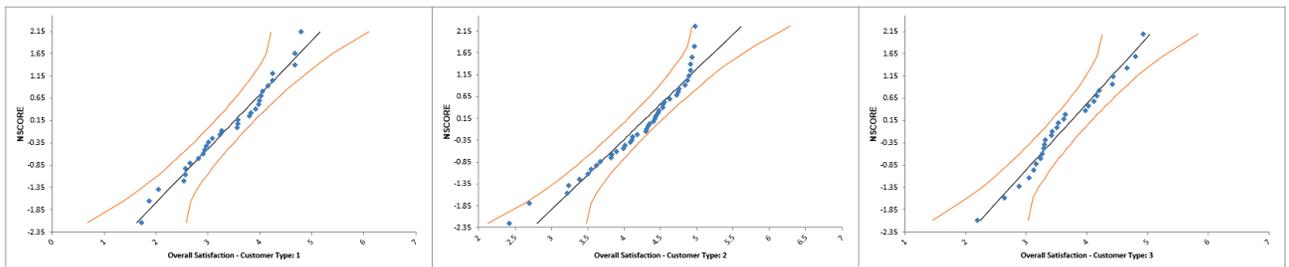
- Click **Sheet 1** Tab of **Customer Data.xlsx**.
- Click **SigmaXL > Graphical Tools > Normal Probability Plots**.
- Ensure that entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.

11. Select *Overall Satisfaction*; click **Numeric Data Variable (Y) >>**. Click **OK**. A Normal Probability Plot of Customer Satisfaction data is produced:



Is this data normally distributed? See earlier histogram and descriptive statistics of Customer Satisfaction data.

12. Now we would like to stratify the customer satisfaction score by customer type and look at the normal probability plots.
13. Click **Sheet 1** of **Customer Data.xlsx**. Click **SigmaXL > Graphical Tools > Normal Probability Plots**. Ensure that **Entire Table** is selected, click **Next**. (Alternatively, press **F3** or click **Recall SigmaXL Dialog** to recall last dialog).
14. Select *Overall Satisfaction*, click **Numeric Data Variable (Y) >>**; select *Customer Type* as **Group Category (X) >>**. Click **OK**. Normal Probability Plots of Overall Satisfaction by Customer Type are produced:



Reviewing these normal probability plots, along with the previously created histograms and descriptive statistics, we see that the satisfaction data for customer type 2 is not normal, and skewed left, which is desirable for satisfaction data! Note that although the customer type 2 data falls within the 95% confidence intervals, the Anderson Darling test from descriptive statistics

shows $p < .05$ indicating nonnormal data. Smaller sample sizes tend to result in wider confidence intervals, but we still see that the curvature for customer type 2 is quite strong.

Tip: Use the Normal Probability Plot (NPP) to distinguish reasons for nonnormality. If the data fails the Anderson-Darling (AD) test (with $p < 0.05$) and forms a curve on the NPP, it is inherently nonnormal or skewed. Calculations such as Sigma Level, Pp, Cp, Ppk and Cpk assume normality and will therefore be affected. Consider transforming the data using LN(Y) or SQRT(Y) or using the Box-Cox Transformation tool (**SigmaXL > Data Manipulation > Box-Cox Transformation**) to make the data normal. Of course, whatever transformation you apply to your data, you must also apply to your specification limits. See also the Process Capability for nonnormal data tools.

If the data fails the AD normality test, but the bulk of the data forms a straight line and there are some outliers, the outliers are driving the nonnormality. Do not attempt to transform this data! Determine the root cause for the outliers and take corrective action on those root causes.

Part H- Run Charts

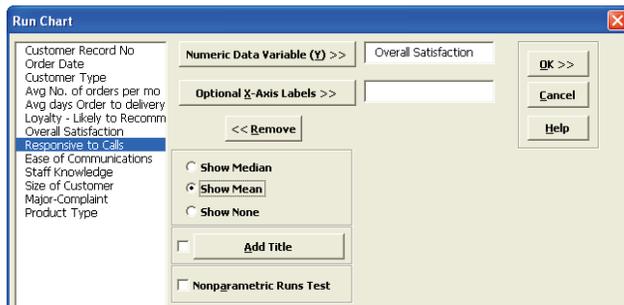
Run charts, also known as trend charts and time series plots, add the dimension of time to the graphical tools. They allow us to see trends and outliers in the data. Run Charts are a precursor to control charts, which add calculated control limits. Note that Run Charts should be used only on unsorted data, in its original chronological sequence.

Basic Run Chart Template

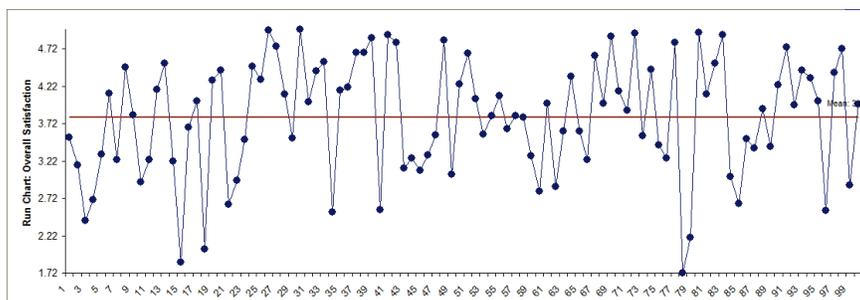
Click **SigmaXL > Templates and Calculators > Basic Graphical Templates > Run Chart** or **SigmaXL > Graphical Tools > Basic Graphical Templates > Run Chart**. See Part B – Templates and Calculators for a Run Chart Template example.

Run Charts

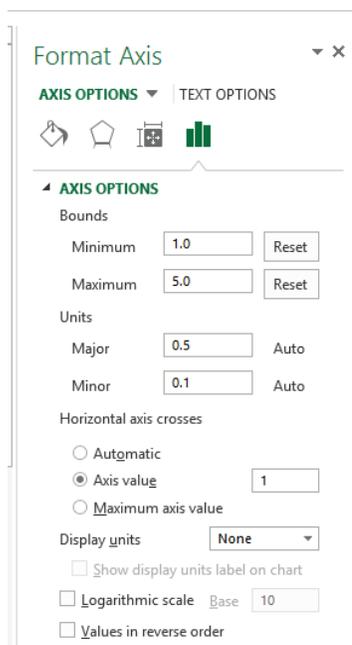
1. Click **Sheet 1** Tab of **Customer Data.xlsx** (or press **F4** to activate last worksheet). Click **SigmaXL > Graphical Tools > Run Chart**. Ensure that entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
2. Select *Overall Satisfaction*, click **Numeric Data Variable (Y) >>**. Select **Show Mean**. Uncheck **Nonparametric Runs Test** (to be discussed later in Part N of Analyze Phase).



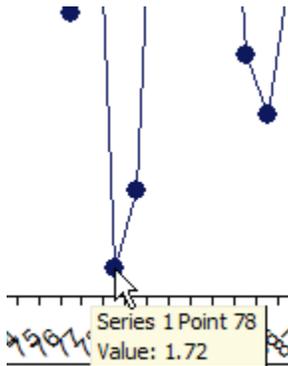
3. Click **OK**. A Run Chart of Overall Satisfaction with Mean center line is produced.



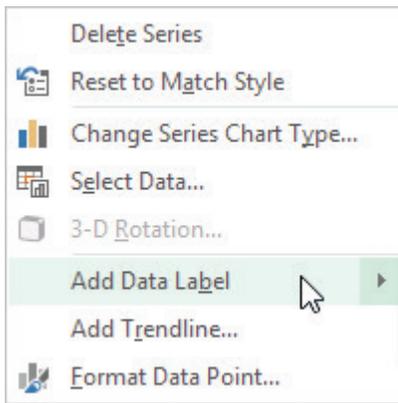
4. Select the Y axis, Right Click, **Format Axis**, to activate the **Format Axis** dialog. Change **Minimum** to 1, **Maximum** to 5, **Horizontal axis crosses** > **Axis value** to 1:



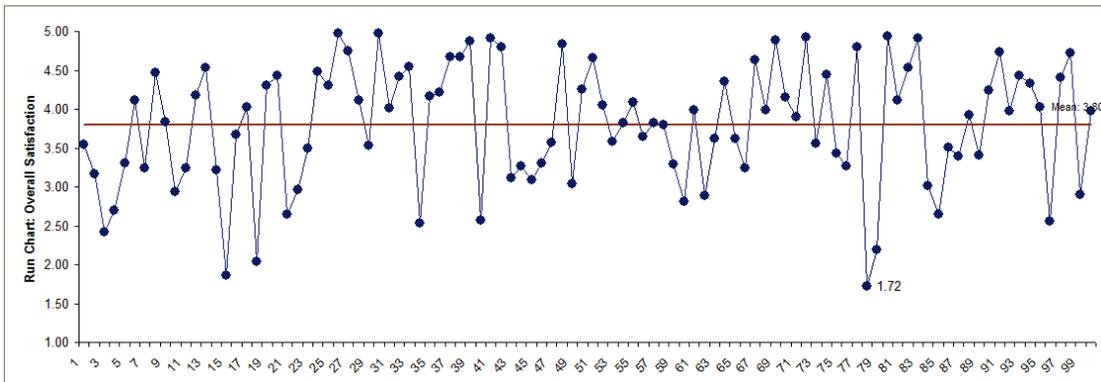
5. The chart is automatically updated.
6. Are there any obvious trends? Some possible cycling, but nothing clearly stands out. It may be interesting to look more closely at a specific data point. Any data point value can be identified by simply moving the cursor over it:



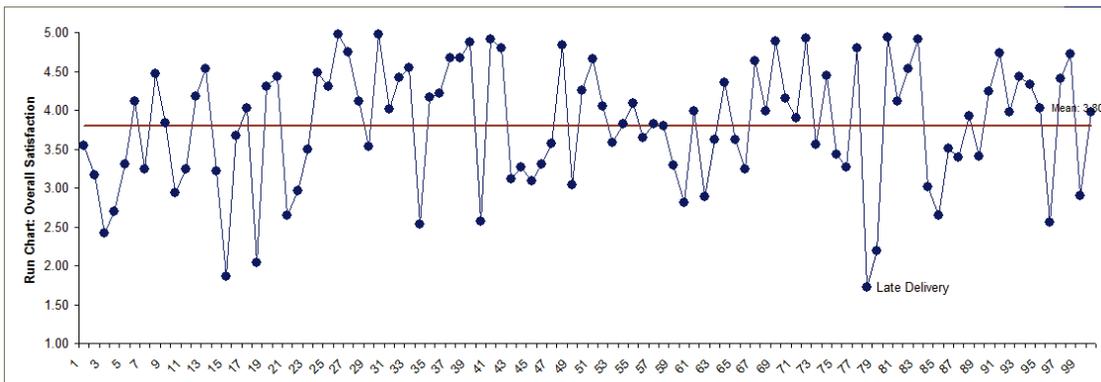
- 7. A label can be added to a data point by one single-click on the data point, followed by a Right mouse click, and select **Add Data Label**. See also **SigmaXL Chart Tools > Add Data Label** in Control Phase Tools, Part B - X-Bar & Range Charts.



- 8. Click **OK**. Resulting Run Chart with label attached to data point:

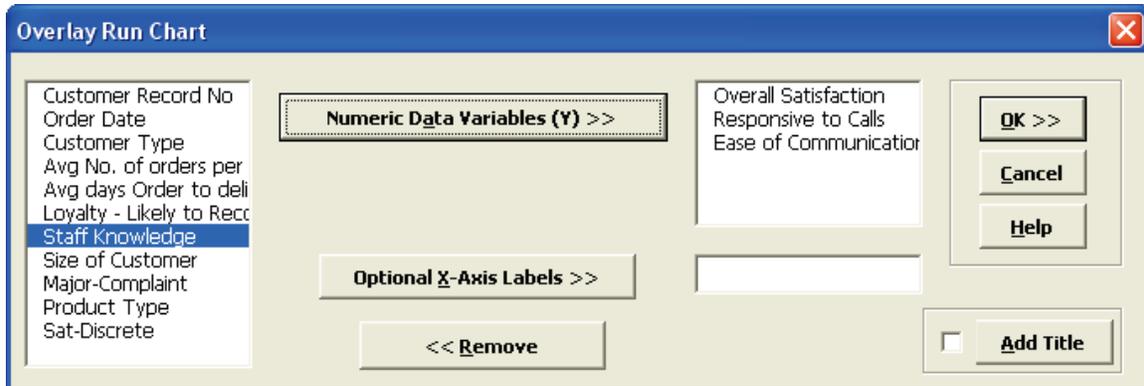


- 9. This label can be changed to a text comment. Single-click three times on the label and type in a comment as shown:

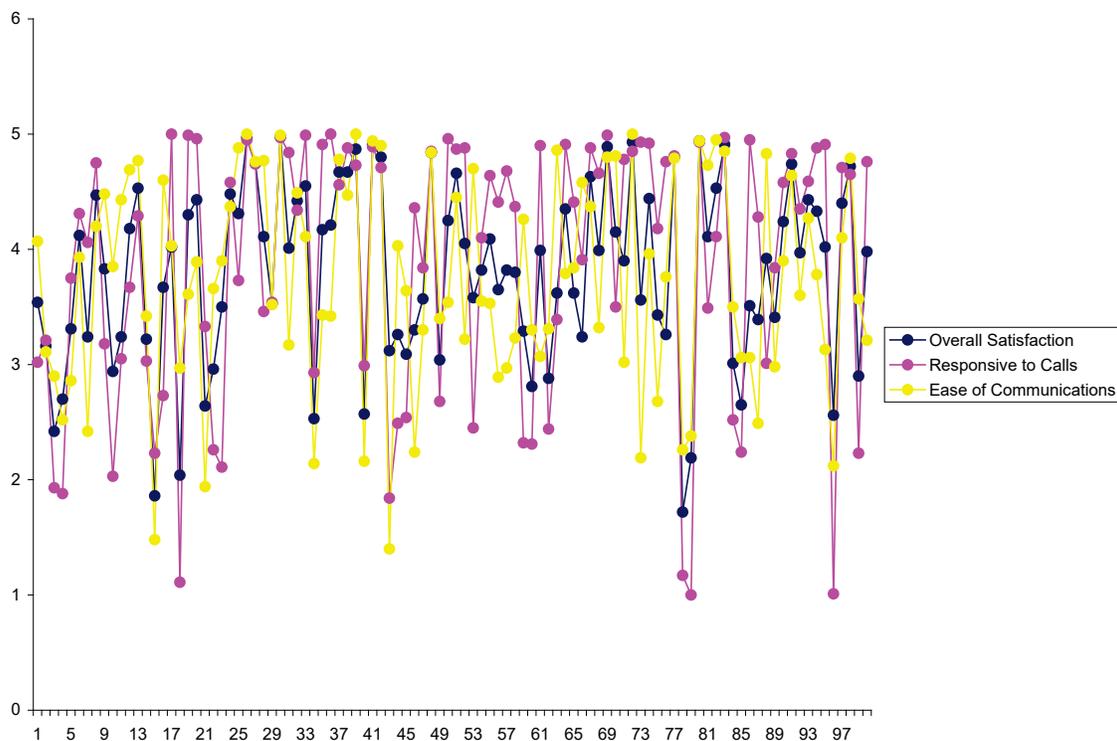


Overlay Run Charts

1. Click **Sheet 1** Tab of **Customer Data.xlsx** (or press F4 to activate last worksheet). Click **SigmaXL > Graphical Tools > Overlay Run Chart**.
2. Ensure that entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
3. Select *Overall Satisfaction, Responsive to Calls and Ease of Communications*. Click **Numeric Data Variable (Y) >>**.



4. Click **OK**. An Overlay Run Chart of Overall Satisfaction, Responsive to Calls and Ease of Communications is produced.



Part I – Measurement Systems Analysis

Type 1 Gage Study Template

Click **SigmaXL > Templates and Calculators > Basic MSA Templates > Type 1 Gage Study** or **SigmaXL > Measurement Systems Analysis > Basic MSA Templates > Type 1 Gage Study**.

See [Basic MSA Templates – Type 1 Gage Study](#) for an example of the Type 1 Gage Study template.

Gage Bias and Linearity Study Template

Click **SigmaXL > Templates and Calculators > Basic MSA Templates > Gage Bias and Linearity Study** or **SigmaXL > Measurement Systems Analysis > Basic MSA Templates > Gage Bias and Linearity Study**.

See [Basic MSA Templates – Gage Bias and Linearity Study](#) for an example of the Gage Bias and Linearity Study template.

Gage R&R Study (MSA) Template

Click **SigmaXL > Templates and Calculators > Basic MSA Templates > Gage R&R Study (MSA)** or **SigmaXL > Measurement Systems Analysis > Basic MSA Templates > Gage R&R Study (MSA)**.

See [Basic MSA Templates – Gage R&R Study \(MSA\)](#) for an example of the Gage R&R Study (MSA) template. If you wish to analyze the template data with the Gage R&R Analysis Tool, click Create Stacked Column Format for “Analyze Gage R&R” >>.

After entering the data into the Gage R&R Study template, you can create Gage Multi-Vari & X-bar R charts. Click **SigmaXL > Templates and Calculators > Basic MSA Templates > Gage R&R: Multi-Vari & X-bar R Charts** or **SigmaXL > Measurement Systems Analysis > Basic MSA Templates > Gage R&R: Multi-Vari & X-bar R Charts**. See [Basic MSA Templates - Gage R&R: Multi-Vari & X-bar R Charts](#) for an example.

Attribute Gage R&R (MSA) Template

Click **SigmaXL > Templates and Calculators > Basic MSA Templates > Attribute Gage R&R (MSA)** or **SigmaXL > Measurement Systems Analysis > Basic MSA Templates > Attribute Gage R&R (MSA)**.

See [Basic MSA Templates – Attribute Gage R&R \(MSA\)](#) for an example of the Attribute Gage R&R (MSA) template.

Create Gage R&R (Crossed) Worksheet

1. To create a blank Gage R&R worksheet, click **SigmaXL > Measurement Systems Analysis > Create Gage R&R (Crossed) Worksheet**. A crossed Gage R&R study means that every operator/appraiser will measure every part/sample.

The dialog box titled "Create Gage R&R (Crossed) Worksheet" contains the following elements:

- Number of Parts/Samples:** 10
- Number of Operators/Appraisers:** 3
- Number of Replicates/Trials:** 3
- Randomize Parts/Samples
- Randomize Operators/Appraisers
- Part/Sample Names:**
 - 1: Part 01
 - 2: Part 02
 - 3: Part 03
- Operator/Appraiser Names:**
 - 1: Operator A
 - 2: Operator B
 - 3: Operator C
- Buttons:** OK >>, Cancel, Help, Reset

2. The default settings are 10 Parts, 3 Operators, 3 Replicates with Parts and Operators randomized. Clicking the **Reset** button will restore the dialog to these default settings.
3. Part and Operator names can be edited.
4. Click **OK**. The following worksheet is produced:

Gage R&R Study (Crossed) Worksheet

Gage Name:	
Date of Study:	
Performed By:	
Notes:	

Run Order	Std. Order	Part	Operator	Measurement
1	86	Part 09	Operator C	
2	61	Part 01	Operator C	
3	65	Part 02	Operator C	
4	74	Part 05	Operator C	
5	64	Part 02	Operator C	
6	83	Part 08	Operator C	
7	71	Part 04	Operator C	
8	66	Part 02	Operator C	
9	62	Part 01	Operator C	
10	90	Part 10	Operator C	
11	76	Part 06	Operator C	
12	80	Part 07	Operator C	
13	73	Part 05	Operator C	
14	79	Part 07	Operator C	
15	63	Part 01	Operator C	

5. Your worksheet order will be different due to the randomization. The results of the Gage R&R study are entered in the yellow highlighted Measurement cells.

Analyze Gage R&R (Crossed)

1. Open the file **Gage RR – AIAG.xlsx**. This is an example from the Automotive Industry Action Group (AIAG) MSA Reference Manual, 3rd Edition, page 101. Note that parts were measured in random order, but the worksheet is given in standard order.
2. Click **SigmaXL > Measurement Systems Analysis > Analyze Gage R&R (Crossed)**. The data worksheet is recognized by SigmaXL and highlighted automatically. Click **Next**.
3. Select *Part*, *Operator* and *Measurement* as shown. Check **Display Multi-Vari & X-Bar R Charts**. Check **Tolerance/Historical StDev (Optional)** and enter 8 for the Process Tolerance **Upper-Lower Spec**.

Analyze Gage R&R (Crossed)

Run Order
Std. Order

Part >> Part

Operator >> Operator

Measurement >> Measurement

<< Remove

OK >>

Cancel

Help

Standard Deviation Multiplier: 6

Alpha to Remove Interaction: 0.1

Confidence Level: 90.0 %

Display Multi-Vari & X-Bar R Charts

Report Information (Optional)

Tolerance/Historical StDev (Optional)

Gage Name: AIAG Example, MSA Reference

Performed By:

Date:

Notes: Parts were measured in random

Upper-Lower Spec: 8

Upper Spec:

Lower Spec:

Historical Process Standard Deviation:

SigmaXL: Measure Phase Tools

4. Click **OK**. The following Gage R&R Study Report is produced:

Gage R&R Study (Crossed) Report

Gage Name:	AIAG Example, MSA Reference Manual, 3rd Edition, Page 101
Date of Study:	
Performed By:	
Notes:	Parts were measured in random order, but worksheet is given in standard order.

Process Tolerance (USL - LSL):	8
Historical Process Standard Deviation:	
Standard Deviation Multiplier:	6
Alpha to Remove Interaction:	0.1
Confidence Level:	90.0
Number of Parts:	10
Number of Operators:	3
Number of Replicates:	3
Design Type:	Balanced

Analysis of Variance with Part * Operator Interaction:

Source	DF	SS	MS	F	P
Part:	9	88.362	9.8180	492.29	0.0000
Operator:	2	3.1673	1.5836	79.406	0.0000
Part * Operator:	18	0.358982	0.019943457	0.433721	0.9741
Repeatability:	60	2.7589	0.045982222		
Total:	89	94.647	1.06		

Analysis of Variance without Part * Operator Interaction (P for Interaction >= 0.1):

Source	DF	SS	MS	F	P
Part:	9	88.362	9.8180	245.61	0.0000
Operator:	2	3.1673	1.5836	39.617	0.0000
Repeatability:	78	3.1179	0.039973276		
Total:	89	94.64711222	1.063450699		

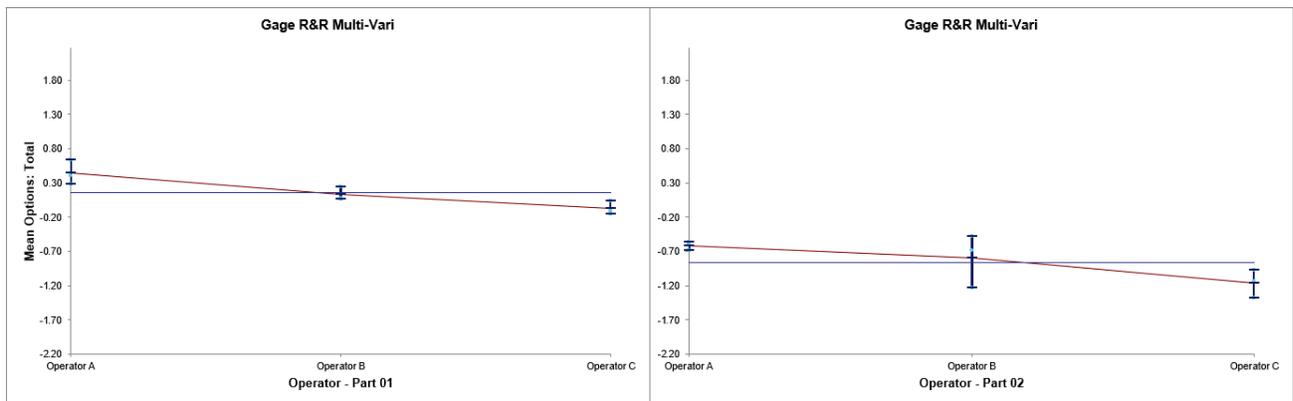
Gage R&R Metrics	StDev	StDev Lower 90% CI	StDev Upper 90% CI	6 * StDev	% Total Variation (TV)	% TV Lower 90% CI	% TV Upper 90% CI
Gage R&R:	0.302372	0.235108	1.03	1.8142	27.86	15.46	70.93
Operator (AV Appraiser Variation):	0.226838	0.127545	1.01	1.3610	20.90		
Part * Operator (INT Interaction):	0	0	0	0	0.00		
Reproducibility (SQRT(AV^2 + INT^2)):	0.226838	0.127545	1.01	1.3610	20.90		
Repeatability (EV Equipment Variation):	0.199933	0.176915	0.230560	1.1996	18.42		
Part Variation (PV):	1.04	0.758821	1.7170	6.2540	96.04		
Total Variation (TV):	1.09	0.816099	1.8111	6.5118	100.00		

Gage R&R Metrics	% Tolerance	% Tolerance Lower 90% CI	% Tolerance Upper 90% CI
Gage R&R:	22.68	17.63	77.50
Operator (AV Appraiser Variation):	17.01	9.57	76.03
Part * Operator (INT Interaction):	0.00	0.00	0.00
Reproducibility (SQRT(AV^2 + INT^2)):	17.01	9.57	76.03
Repeatability (EV Equipment Variation):	14.99	13.27	17.29
Part Variation (PV):	78.17	56.91	128.78
Total Variation (TV):	81.40	61.21	135.83

Gage R&R Metrics	Variance Component	% Contribution of Variance Component
Gage R&R:	0.091428538	7.76
Operator:	0.051455261	4.37
Part * Operator:	0	0.00
Reproducibility:	0.051455261	4.37
Repeatability:	0.039973276	3.39
Part Variation:	1.09	92.24
Total Variation:	1.1779	100.00

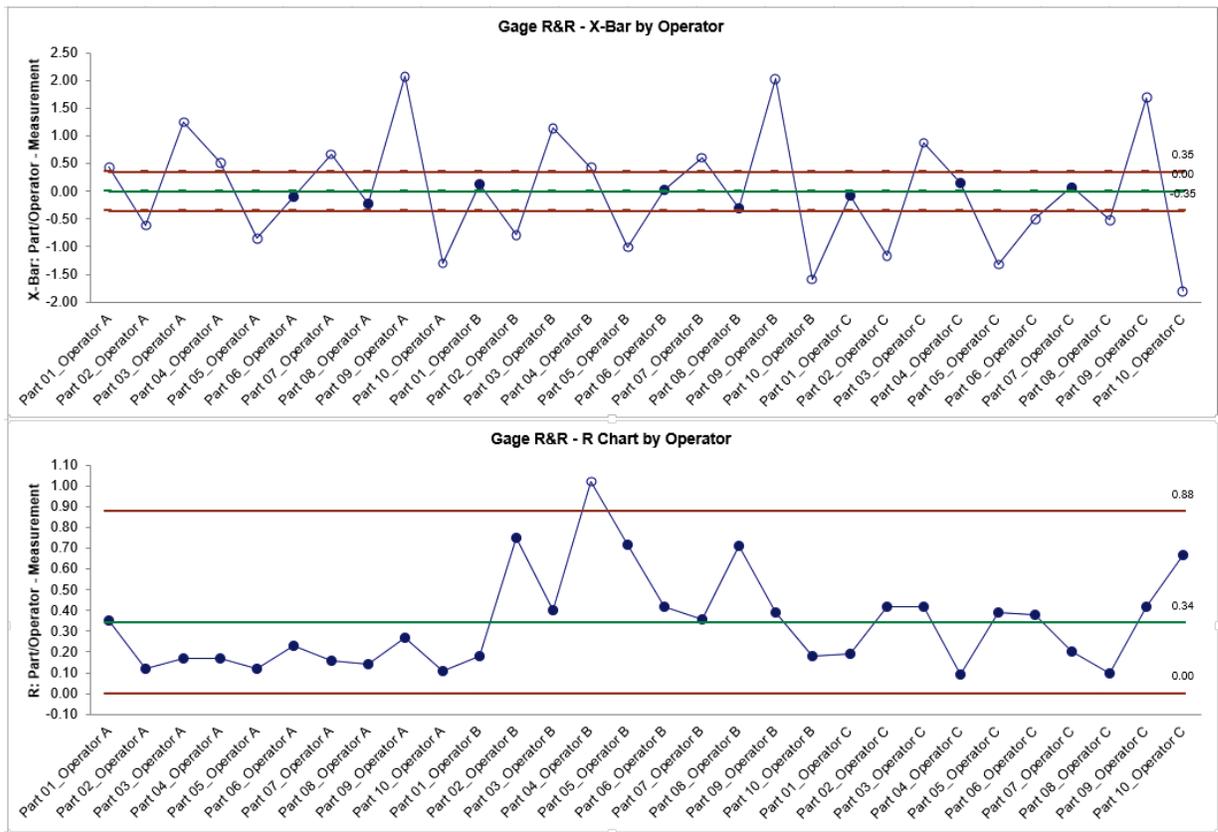
Gage R&R Metrics	NDC	NDC Lower 90% CI	NDC Upper 90% CI
Number of Distinct Categories (Signal-to-Noise Ratio: 1.41 * PV/R&R):	4.9	1.4	9.0

5. The first ANOVA table shows that the Operator * Part Interaction is not significant with the P-Value = 0.9741 (greater than the alpha-to-remove value of 0.1). The second ANOVA table shows the results excluding the Part * Operator Interaction.
6. The cells shaded in light blue highlight the critical metrics: Gage R&R % Total Variation (also known as %R&R) and %Tolerance: < 10% indicates a good measurement system; > 30% indicates an unacceptable measurement system. The Number of Distinct Categories should be at least 5. Initial analysis shows that this is a marginal measurement system. (Traditionally NDC is truncated to an integer value, but SigmaXL reports a more informative one decimal place).
7. SigmaXL also reports the confidence intervals for the Gage R&R metrics. Note that when the confidence intervals are taken into account, it is possible that this measurement system is totally inadequate. This strongly suggests that the measurement system needs to be improved but it also points out the weakness of the traditional Gage R&R study with 10 Parts, 3 Operators, and 3 Replicates. These sample sizes will typically yield wide confidence intervals. For statistical details of the Gage R&R confidence intervals see Burdick, R. K., Borror, C. M., and Montgomery, D. C., “Design and Analysis of Gauge R&R Studies: Making Decisions with Confidence Intervals in Random and Mixed ANOVA Models”, ASA-SIAM Series on Statistics and Applied Probability, 2005.
8. Click on the **Gage R&R – Multi-Vari Sheet** tab to view the Multi-Vari Chart:



The Multi-Vari chart shows each Part as a separate graph. Each Operator’s response readings are denoted as a vertical line with the top tick corresponding to the Maximum value, bottom tick is the Minimum, and the middle tick is the Mean. The horizontal line across each graph is the overall average for each part.

9. Click on the **Gage R&R – X-Bar R Sheet** tab to view the X-Bar & R Chart:



When interpreting the X-bar and R chart for a Gage R&R study, it is desirable that the X-bar chart be out-of-control, and the Range chart be in-control. The control limits are derived from within Operator repeatability. The Range chart indicates a problem with Operator B being inconsistent.

Attribute MSA (Binary)

Attribute MSA is also known as Attribute Agreement Analysis. Use the Binary option if the assessed result is text or numeric binary (e.g., 0/1, Pass/Fail, Good/Bad, G/NG, Yes/No).

1. Open the file **Attribute MSA – AIAG.xlsx**. This is an example from the Automotive Industry Action Group (AIAG) MSA Reference Manual, 3rd edition, page 127 (4th Edition, page 134). There are 50 samples, 3 appraisers and 3 trials with a 0/1 response. A “good” sample is denoted as a 1. A “bad” sample is denoted as a 0. Note that the worksheet data must be in stacked column format and the known reference values must be consistent for each sample.
2. Click **SigmaXL > Measurement Systems Analysis > Attribute MSA (Binary)**. Ensure that the entire data table is selected. Click **Next**.
3. Select *Part*, *Appraiser*, *Assessed Result* and *Reference* as shown. Check **Report Information** and enter *AIAG Attribute MSA Binary* for **Product/Unit Name**. Select **Percent Confidence Interval Type – Exact**. The default **Good Level** of “1” will be used as specified in the AIAG manual:

The screenshot shows the 'Attribute MSA (Binary)' dialog box. It has a title bar with a close button. The main area contains several input fields and buttons. On the left is a large empty box. To its right are four rows of buttons and text boxes: 'Part/Sample >>' with 'Part', 'Appraiser >>' with 'Appraiser', 'Assessed Result >>' with 'Assessed Result', and 'True Standard (Optional) >>' with 'Reference'. Below these is a '<< Remove' button. To the right of these fields are three buttons: 'OK >>', 'Cancel', and 'Help'. Below the main area is a 'Good Level >>' section with a list box containing '0' and '1', and a text box containing '1'. At the bottom, there is a section for 'Report Information (Optional)' which is checked. To its right is a 'Confidence Level' field with '95.0'. Below 'Report Information' are fields for 'Product/Unit Name' (containing 'AIAG Attribute MSA Binary'), 'Performed By', 'Date', and 'Notes'. To the right of these is a 'Percent Confidence Interval Type' section with two radio buttons: 'Wilson Score' (unselected) and 'Exact' (selected).

Tip: The **Good Level** definition is used to determine Type I and Type II error rates. It is applicable only when a **True Standard** is selected.

Tip: **Percent Confidence Interval Type** applies to the Percent Agreement and Percent Effectiveness Confidence Intervals. These are binomial proportions that have an "oscillation" phenomenon where the coverage probability varies with the sample size and proportion value. **Exact** is strictly conservative and will guarantee the specified confidence level as a minimum coverage probability, but results in wider intervals. **Wilson Score** has a mean coverage

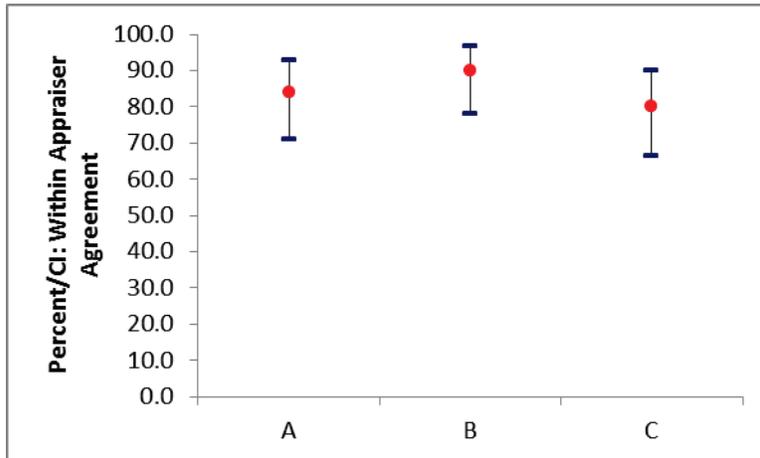
probability that matches the specified confidence interval. Since the intervals are narrower and thereby more powerful, Wilson Score is recommended for use in attribute MSA studies due to the small sample sizes typically used. **Exact** is selected in this example for continuity with the results from SigmaXL Version 6.

4. Click **OK**. The Attribute MSA Binary Analysis Report is produced. The tables and associated graphs are described separately by section for clarity.

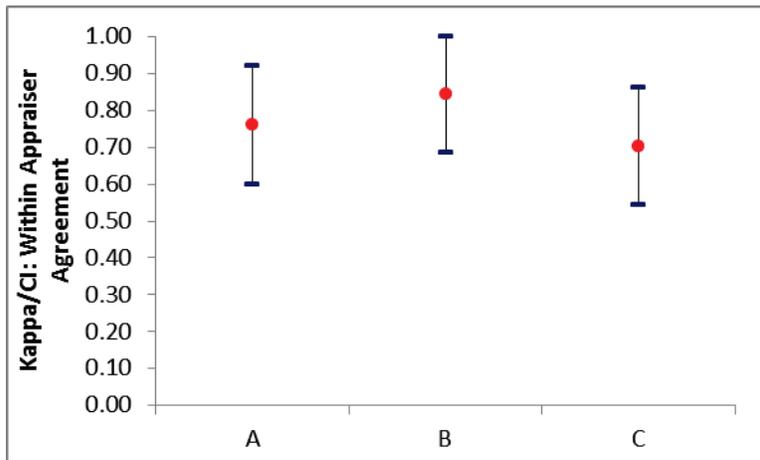
Tip: While this report is quite extensive, a quick assessment of the attribute measurement system can be made by viewing the Kappa color highlights: **Green** - very good agreement (Kappa ≥ 0.9); **Yellow** - marginally acceptable, improvement should be considered (Kappa 0.7 to < 0.9); **Red** - unacceptable (Kappa < 0.7). Further details on Kappa are given below.

Within Appraiser Agreement is an assessment of each appraiser’s consistency of ratings across trials and requires at least two trials. This is analogous to Gage R&R Repeatability. Note that the reference standard is not considered, so an appraiser may be perfectly consistent but consistently wrong!

Percent/CI: Within Appraiser Agreement Graph:



Kappa/CI: Within Appraiser Agreement Graph:



Within Appraiser Agreement Table:

Attribute Agreement Report:

Within Appraiser Agreement	# Inspected	# Matched	Percent	95.0% LC (Exact)	95.0% UC (Exact)	Fleiss' Kappa	Fleiss' Kappa P-Value	Fleiss' Kappa 95.0% LC	Fleiss' Kappa 95.0% UC
A	50	42	84.00	70.89	92.83	0.7600	0.0000	0.6000	0.9200
B	50	45	90.00	78.19	96.67	0.8451	0.0000	0.6850	1.0000
C	50	40	80.00	66.28	89.97	0.7029	0.0000	0.5429	0.8629

Within Appraiser Percent Agreement will decrease as the number of trials increase because a match occurs only if an appraiser is consistent across all trials. Use the **Kappa/CI: Within Appraiser Agreement Graph** to determine adequacy of the Within Appraiser agreement. See below for additional interpretation guidelines.

Tip: Hover the mouse pointer over the heading cells to view the following report comments.

Inspected: Number of parts or samples.

Matched: A match occurs only if an appraiser is consistent across all trials.

Percent Agreement = (# Matched / # Inspected) * 100

LC = Percent Lower Confidence Limit. **UC = Percent** Upper Confidence Limit. Confidence intervals (CI) for binomial proportions have an "oscillation" phenomenon where the coverage probability varies with n and p. **Exact** is strictly conservative and will guarantee the specified confidence level as a minimum coverage probability, but results in wide intervals. Wilson **Score** has a mean coverage probability that matches the specified confidence interval. Since the intervals are narrower and thereby more powerful, they are recommended for use in attribute MSA studies due to the small sample sizes typically used. See Appendix [Percent Confidence Intervals \(Exact Versus Wilson Score\)](#) for references.

Fleiss' Kappa statistic is a measure of agreement that is analogous to a "correlation coefficient" for discrete data. Kappa ranges from -1 to +1: A Kappa value of +1 indicates perfect agreement. If Kappa = 0, then agreement is the same as would be expected by chance. If Kappa = -1, then there is perfect disagreement. **"Rule-of-thumb" interpretation guidelines: >= 0.9 very good agreement (green); 0.7 to < 0.9 marginally acceptable, improvement should be considered (yellow); < 0.7 unacceptable (red).** See Appendix [Kappa](#) for further details on the Kappa calculations and "rule-of-thumb" interpretation guidelines.

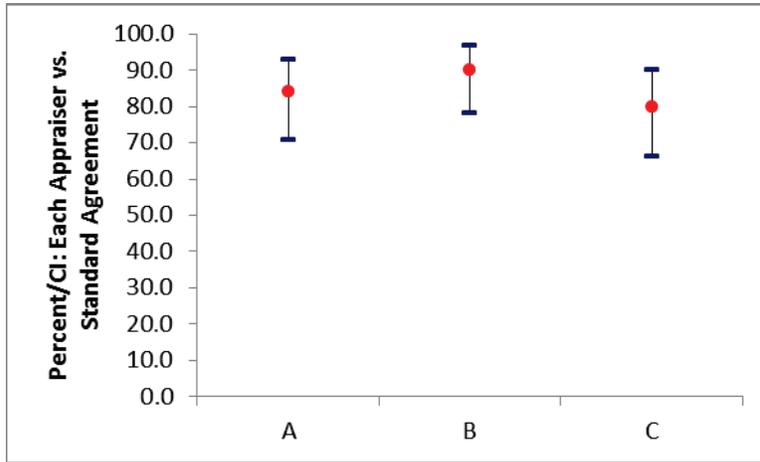
Fleiss' Kappa P-Value: H0: Kappa = 0. If P-Value < alpha (.05 for specified 95% confidence level), reject H0 and conclude that agreement is not the same as would be expected by chance. Significant P-Values are highlighted in red.

Fleiss' Kappa LC (Lower Confidence) and **Fleiss' Kappa UC** (Upper Confidence) limits use a kappa normal approximation. Interpretation Guidelines: Kappa lower confidence limit >= 0.9: very good agreement. Kappa upper confidence limit < 0.7: the attribute agreement is unacceptable. Wide confidence intervals indicate that the sample size is inadequate.

In this example, we have marginal **Within Appraiser Agreement** for each of the appraisers.

Each Appraiser vs. Standard Agreement is an assessment of each appraiser’s ratings across trials compared to a known reference standard. This is analogous to Gage R&R Accuracy.

Percent/CI: Each Appraiser vs. Standard Agreement Graph:



Each Appraiser vs. Standard Agreement Table:

Each Appraiser vs. Standard Agreement	# Inspected	# Matched	Percent	95.0% LC (Exact)	95.0% UC (Exact)	Fleiss' Kappa	Fleiss' Kappa P-Value	Fleiss' Kappa 95.0% LC	Fleiss' Kappa 95.0% UC
A	50	42	84.00	70.89	92.83	0.8802	0.0000	0.7202	1.0000
B	50	45	90.00	78.19	96.67	0.9226	0.0000	0.7626	1.0000
C	50	40	80.00	66.28	89.97	0.7747	0.0000	0.6147	0.9347

Tip: The **Percent/CI Each Appraiser vs. Standard Agreement Graph** can be used to compare agreement to standard across the appraisers, but should not be used as an absolute measure of agreement. **Each Appraiser vs. Standard Agreement** will decrease as the number of trials increase because a match occurs only if an appraiser agrees with the standard consistently across all trials. Use **Fleiss' Kappa** in the **Each Appraiser vs. Standard Agreement Table** to determine the adequacy of Each Appraiser versus Standard agreement.

Inspected: Number of parts or samples.

Matched: A match occurs only if an appraiser agrees with the standard consistently across all trials.

Percent Agreement = (# Matched / # Inspected) * 100

Kappa is interpreted as above: **>= 0.9 very good agreement (green); 0.7 to < 0.9 marginally acceptable, improvement should be considered (yellow); < 0.7 unacceptable (red).**

Appraisers A and C have marginal agreement versus the standard values. Appraiser B has very good agreement to the standard.

Each Appraiser vs. Standard Disagreement is a breakdown of each appraiser’s rating misclassifications (compared to a known reference standard). This table is applicable only to binary two-level responses (e.g., 0/1, G/NG, Pass/Fail, True/False, Yes/No).

Each Appraiser vs. Standard Disagreement	Type I Error	Type I Error %	Type II Error	Type II Error %	Mixed Error	Mixed Error %
A	0	0.00	0	0.00	8	16.00
B	0	0.00	0	0.00	5	10.00
C	0	0.00	0	0.00	10	20.00

Disagreement Legend		
Error Type	Assessment	True Standard
Type I:	0	1
Type II:	1	0
Mixed:	Assessments across trials are not identical	

A **Type I Error** occurs when the appraiser consistently assesses a good part/sample as bad. "Good" is defined by the user in the Attribute MSA analysis dialog.

Type I Error % = (Type I Error / # Good Parts or Samples) * 100

A **Type II Error** occurs when a bad part/sample is consistently assessed as good.

Type II Error % = (Type II Error / # Bad Parts or Samples) * 100

A **Mixed Error** occurs when the assessments across trials are not identical.

Mixed Error % = (Mixed Error / # Parts or Samples) * 100

Between Appraiser Agreement is an assessment of the appraisers’ consistency of ratings across trials and between each other. At least two appraisers are required. This is analogous to Gage R&R Reproducibility. Note that the reference standard is not considered, so the appraisers may be perfectly consistent, but consistently wrong!

Between Appraiser Agreement	# Inspected	# Matched	Percent	95.0% LC (Exact)	95.0% UC (Exact)	Fleiss' Kappa	Fleiss' Kappa P-Value	Fleiss' Kappa 95.0% LC	Fleiss' Kappa 95.0% UC
	50	39	78.00	64.04	88.47	0.7936	0.0000	0.7474	0.8398

All Appraisers vs. Standard Agreement is an assessment of all appraisers’ ratings across trials compared to a known reference standard. This is analogous to Gage R&R Accuracy.

All Appraisers vs. Standard Agreement	# Inspected	# Matched	Percent	95.0% LC (Exact)	95.0% UC (Exact)	Fleiss' Kappa	Fleiss' Kappa P-Value	Fleiss' Kappa 95.0% LC	Fleiss' Kappa 95.0% UC
	50	39	78.00	64.04	88.47	0.8592	0.0000	0.7668	0.9516

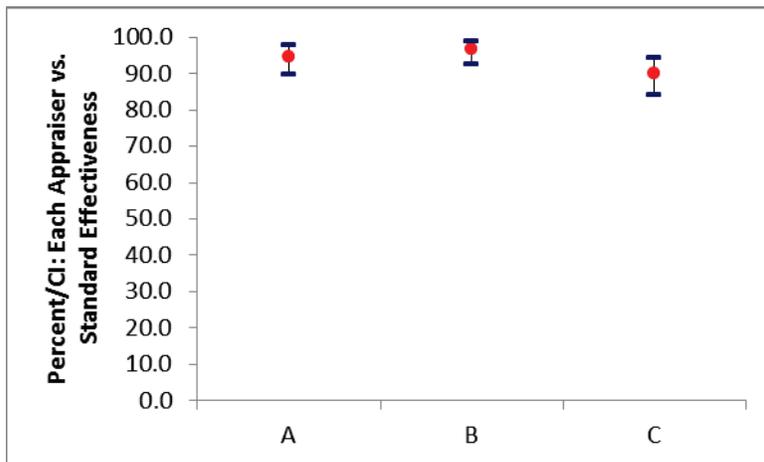
Kappa is interpreted as above: **>= 0.9 very good agreement (green); 0.7 to < 0.9 marginally acceptable (yellow); < 0.7 unacceptable (red).**

Since the **Between Appraiser Agreement** and **All Appraisers vs. Standard Agreement** are marginally acceptable, improvements to the attribute measurement should be considered. Look for unclear or confusing operational definitions, inadequate training, operator distractions or poor lighting. Consider the use of pictures to clearly define a defect.

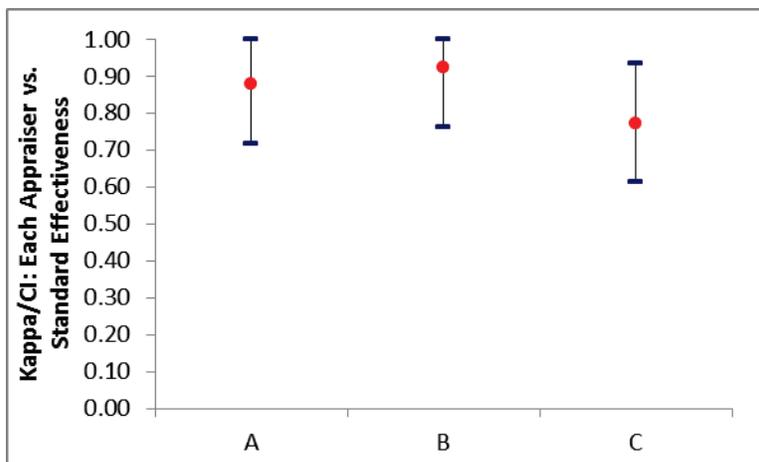
The **Attribute Effectiveness Report** is similar to the **Attribute Agreement Report**, but treats each trial as an opportunity. Consistency across trials or appraisers is not considered. This has the benefit of providing a Percent measure that is unaffected by the number of trials or appraisers. Also, the increased sample size for # Inspected results in a reduction of the width of the Percent confidence interval. The Misclassification report shows all errors classified as Type I or Type II. Mixed errors are not relevant here. This report requires a known reference standard and includes: **Each Appraiser vs. Standard Effectiveness, All Appraisers vs. Standard Effectiveness, and Effectiveness and Misclassification Summary.**

Each Appraiser vs. Standard Effectiveness is an assessment of each appraiser’s ratings compared to a known reference standard. This is analogous to Gage R&R Accuracy. Unlike the **Each Appraiser vs. Standard Agreement** table above, consistency across trials is not considered here - each trial is considered as an opportunity. This has the benefit of providing a Percent measure that is unaffected by the number of trials. Also, the increased sample size for # Inspected results in a reduction of the width of the Percent confidence interval.

Percent/CI: Each Appraiser vs. Standard Effectiveness Graph:



Kappa/CI: Each Appraiser vs. Standard Effectiveness Graph:



Each Appraiser vs. Standard Effectiveness Table:

Attribute Effectiveness Report:

Each Appraiser vs. Standard Effectiveness	# Inspected	# Matched	Percent	95.0% LC (Exact)	95.0% UC (Exact)	Fleiss' Kappa	Fleiss' Kappa P-Value	Fleiss' Kappa 95.0% LC	Fleiss' Kappa 95.0% UC
A	150	142	94.67	89.76	97.67	0.8788	0.0000	0.7187	1.0000
B	150	145	96.67	92.39	98.91	0.9230	0.0000	0.7629	1.0000
C	150	135	90.00	84.04	94.29	0.7739	0.0000	0.6138	0.9339

Inspected = # Parts or Samples * # Trials

Matched: A match occurs if an appraiser agrees with the standard (consistency across trials is not considered here).

Percent Effectiveness = (# Matched / # Inspected) * 100. Interpretation Guidelines for Percent Effectiveness: => 95% very good; 85% to <95% marginal, may be acceptable but improvement should be considered; < 85% unacceptable. These guidelines assume an equal number of known good and known bad parts/samples.

Kappa is interpreted as above: **>= 0.9 very good agreement (green); 0.7 to < 0.9 marginally acceptable (yellow); < 0.7 unacceptable (red).**

Tip: The Kappa values in the Effectiveness tables are very similar to those in the Agreement tables (the slight difference is due to average Kappa for unstacked versus Kappa for stacked data). This is why the **Kappa/CI Each Appraiser vs. Standard Agreement** graph is not shown. It would essentially be a duplicate of the **Kappa/CI Each Appraiser vs. Standard Effectiveness** graph.

Appraisers A and C have marginal agreement versus the standard values, with less than 95% Effectiveness and Kappa < 0.9. Appraiser B has very good agreement to the standard.

All Appraisers vs. Standard Effectiveness is an assessment of all appraisers' ratings compared to a known reference standard. This is analogous to Gage R&R Accuracy. Unlike the **All Appraiser vs. Standard Agreement** table above, consistency across trials and appraisers is not considered here - each trial is considered as an opportunity. This has the benefit of providing a Percent measure that is unaffected by the number of trials or appraisers. Also, the increased sample size for # Inspected results in a reduction of the width of the Percent confidence interval.

All Appraisers vs. Standard Effectiveness	# Inspected	# Matched	Percent	95.0% LC (Exact)	95.0% UC (Exact)	Fleiss' Kappa	Fleiss' Kappa P-Value	Fleiss' Kappa 95.0% LC	Fleiss' Kappa 95.0% UC
	450	422	93.78	91.13	95.83	0.8581	0.0000	0.7657	0.9505

Inspected = # Parts or Samples * # Trials * # Appraisers

Matched: A match occurs if an appraiser agrees with the standard (consistency across trials and appraisers is not considered here).

Percent Effectiveness = (# Matched / # Inspected) * 100.

The interpretation guidelines for Kappa and Percent Effectiveness are the same as noted above. This measurement system is marginally acceptable.

Each Appraiser vs. Standard Misclassification is a breakdown of each appraiser’s rating misclassifications (compared to a known reference standard). This table is applicable only to binary two-level responses (e.g., 0/1, G/NG, Pass/Fail, True/False, Yes/No). Unlike the **Each Appraiser vs. Standard Disagreement** table above, consistency across trials is not considered here. All errors are classified as Type I or Type II. Mixed errors are not relevant.

Each Appraiser vs. Standard Misclassification	Type I Error	# Inspected (1)	Type I Error %	Type II Error	# Inspected (0)	Type II Error %
A	5	102	4.90	3	48	6.25
B	2	102	1.96	3	48	6.25
C	9	102	8.82	6	48	12.50

Misclassification Legend		
Error Type	Assessment	True Standard
Type I:	0	1
Type II:	1	0

A **Type I** error occurs when the appraiser assesses a good part/sample as bad (consistency across trials is not considered here). "Good" is defined by the user in the Attribute MSA analysis dialog. See **Misclassification Legend** for specific definition of **Type I** and **Type II** Errors.

Inspected = # Good Parts or Samples * # Trials

Type I Error % = (Type I Error / # Inspected Good) * 100

A **Type II** error occurs when a bad part/sample is assessed as good. See **Misclassification Legend** for specific definition of **Type I** and **Type II** Errors.

Inspected = # Bad Parts or Samples * # Trials

Type II Error % = (Type II Error / # Inspected Bad) * 100

All Appraisers vs. Standard Misclassification is a breakdown of all appraisers’ rating misclassifications (compared to a known reference standard). This table is applicable only to binary two-level responses (e.g., 0/1, G/NG, Pass/Fail, True/False, Yes/No). Unlike the **All Appraisers vs. Standard Disagreement table** above, consistency across trials and appraisers is not considered here. All errors are classified as Type I or Type II. Mixed errors are not relevant.

All Appraisers vs. Standard Misclassification	Type I Error	# Inspected (1)	Type I Error %	Type II Error	# Inspected (0)	Type II Error %
	16	306	5.23	12	144	8.33

A **Type I** error occurs when the appraiser assesses a good part/sample as bad (consistency across trials is not considered here). "Good" is defined by the user in the Attribute MSA analysis dialog. See **Misclassification Legend** for specific definition of **Type I** and **Type II** Errors.

Inspected = # Good Parts or Samples * # Trials * # Appraisers

Type I Error % = (Type I Error / # Inspected Good) * 100

A **Type II** error occurs when a bad part/sample is assessed as good. See **Misclassification Legend** for specific definition of **Type I** and **Type II** Errors.

Inspected = # Bad Parts or Samples * # Trials * # Appraisers

Type II Error % = (Type II Error / # Inspected Bad) * 100

Effectiveness and Misclassification Summary is a summary table of all appraisers' correct rating counts and misclassification counts compared to the known reference standard values.

Effectiveness and Misclassification Summary		Standard (0)	Standard (1)
Appraiser (0)		132	16
Appraiser (1)		12	290

Attribute MSA Data is a summary showing the original data in unstacked format. This makes it easy to visually compare appraiser results by part. If a reference standard is provided, the cells are color highlighted as follows: agrees with reference (green); Type I error (yellow); Type II error (red):

Attribute MSA Data:

Part	Reference	A 1	A 2	A 3	B 1	B 2	B 3	C 1	C 2	C 3
1	1	1	1	1	1	1	1	1	1	1
2	1	1	1	1	1	1	1	1	1	1
3	0	0	0	0	0	0	0	0	0	0
4	0	0	0	0	0	0	0	0	0	0
5	0	0	0	0	0	0	0	0	0	0
6	1	1	1	0	1	1	0	1	0	0
7	1	1	1	1	1	1	1	1	1	1
8	1	1	1	1	1	1	1	1	1	1
9	0	0	0	0	0	0	0	0	0	0
10	1	1	1	1	1	1	1	1	1	1
11	1	1	1	1	1	1	1	1	1	1
12	0	0	0	0	0	0	0	0	0	1
13	1	1	1	1	1	1	1	1	1	1
14	1	1	1	0	1	1	1	1	0	0
15	1	1	1	1	1	1	1	1	1	1
16	1	1	1	1	1	1	1	1	1	1
17	1	1	1	1	1	1	1	1	1	1
18	1	1	1	1	1	1	1	1	1	1
19	1	1	1	1	1	1	1	1	1	1
20	1	1	1	1	1	1	1	1	1	1
21	1	1	1	0	1	0	1	0	1	0
22	0	0	0	1	0	1	0	1	1	0
23	1	1	1	1	1	1	1	1	1	1
24	1	1	1	1	1	1	1	1	1	1
25	0	0	0	0	0	0	0	0	0	0
26	0	0	1	0	0	0	0	0	0	1
27	1	1	1	1	1	1	1	1	1	1
28	1	1	1	1	1	1	1	1	1	1
29	1	1	1	1	1	1	1	1	1	1
30	0	0	0	0	0	0	1	0	0	0
31	1	1	1	1	1	1	1	1	1	1
32	1	1	1	1	1	1	1	1	1	1
33	1	1	1	1	1	1	1	1	1	1
34	0	0	0	1	0	0	1	0	1	0
35	1	1	1	1	1	1	1	1	1	1
36	1	1	1	0	1	1	1	1	0	1
37	0	0	0	0	0	0	1	0	0	0
38	1	1	1	1	1	1	1	1	1	1
39	0	0	0	0	0	0	0	0	0	0
40	1	1	1	1	1	1	1	1	1	1
41	1	1	1	1	1	1	1	1	1	1
42	0	0	0	0	0	0	0	0	0	0
43	1	1	0	1	1	1	1	1	1	0
44	1	1	1	1	1	1	1	1	1	1
45	0	0	0	0	0	0	0	0	0	0
46	1	1	1	1	1	1	1	1	1	1
47	1	1	1	1	1	1	1	1	1	1
48	0	0	0	0	0	0	0	0	0	0
49	1	1	1	1	1	1	1	1	1	1
50	0	0	0	0	0	0	0	0	0	0

In conclusion, with the Kappa scores in the “yellow zone” (< 0.9) and Percent Effectiveness less than 95% this measurement system is marginal and should be improved. Appraiser B is the exception and does well against the standard. Look for unclear or confusing operational definitions, inadequate training, operator distractions or poor lighting. Consider the use of pictures to clearly define a defect. Use Attribute MSA as a way to “put your stake in the ground” and track the effectiveness of improvements to the measurement system.

Attribute MSA (Ordinal)

Attribute MSA is also known as Attribute Agreement Analysis. Use the Ordinal option if the assessed result is numeric ordinal (e.g., 1, 2, 3, 4, 5). There must be at least 3 response levels in the assessed result, otherwise it is binary. Examples of ordinal responses used elsewhere in this workbook include:

- *Customer Loyalty –Likely to Recommend* score which contains ordinal integer values from 1 to 5, where a 1 indicates that the customer is very unlikely to recommend and a 5 indicates that the customer is very likely to recommend. This example is used in [Ordinal Logistic Regression](#).
- *Taste Score* on a scale of 1-7 where 1 is "awful" and 7 is "delicious." This is used in the cake bake taste test Design of Experiments: [Part B – Three Factor Full Factorial Example Using DOE Template](#).

An Ordinal Attribute MSA study should be done prior to formal ordinal data collection for use in hypothesis testing, regression or design of experiments.

1. Open the file **Attribute MSA – Ordinal.xlsx**. This is an Ordinal MSA example with 50 samples, 3 appraisers and 3 trials. The response is 1 to 5, grading product quality. One denotes "Very Poor Quality," 2 is "Poor," 3 is "Fair," 4 is "Good" and a 5 is "Very Good Quality." The *Expert Reference* column is the reference standard from an expert appraisal. Note that the worksheet data must be in stacked column format and the reference values must be consistent for each sample.
2. Click **SigmaXL > Measurement Systems Analysis > Attribute MSA (Ordinal)**. Ensure that the entire data table is selected. Click **Next**.

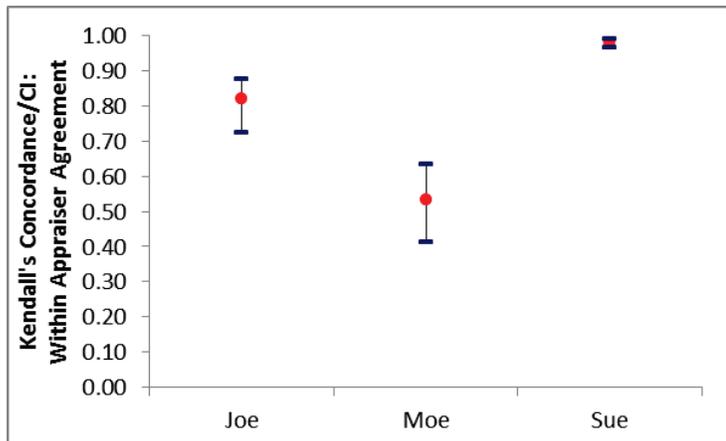
3. Select *Sample No.*, *Appraiser*, *Assessed Result* and *Expert Reference* as shown. Check **Report Information** and enter *Attribute MSA Ordinal* for **Product/Unit Name**. Select **Percent Confidence Interval Type – Wilson Score**:

4. Click **OK**. The Attribute MSA Ordinal Analysis Report is produced.

Tip: While this report is quite extensive, a quick assessment of the attribute measurement system can be made by viewing the **Kendall Concordance** and **Kendall Correlation** color highlights: **Green** - very good agreement; **Yellow** - marginally acceptable, improvement should be considered; **Red** - unacceptable. Further details on the Kendall Coefficients are given below.

Tip: **Fleiss' Kappa** and **Percent Agreement** are included in the report for completeness but not recommended for use with Ordinal response data because they treat each response level as nominal. Kendall's Concordance and Correlation take the order of the data into account, so a deviation of 1 is not as bad as a deviation of 2 or more. See **Attribute MSA – Nominal** for a discussion of the **Fleiss' Kappa** report.

Kendall's Concordance/CI Within Appraiser Agreement Graph:



Within Appraiser Agreement Table:

Attribute Agreement Report:									
Within Appraiser Agreement	# Inspected	# Matched	Percent	95.0% LC (Score)	95.0% UC (Score)	Kendall's Coeff. Concordance	Kendall's Concordance P-Value	Kendall's Concordance 95.0% LC	Kendall's Concordance 95.0% UC
Joe	50	14	28.00	17.47	41.67	0.8204	0.0000	0.7234	0.8772
Moe	50	2	4.00	1.10	13.46	0.5332	0.0048	0.4112	0.6351
Sue	50	40	80.00	66.96	88.76	0.9754	0.0000	0.9670	0.9908

Between Appraiser Agreement Table:

Between Appraiser Agreement	# Inspected	# Matched	Percent	95.0% LC (Score)	95.0% UC (Score)	Kendall's Coeff. Concordance	Kendall's Concordance P-Value	Kendall's Concordance 95.0% LC	Kendall's Concordance 95.0% UC
	50	1	2.00	0.35	10.50	0.0495	0.0000	0.5649	0.7214

Kendall's Coefficient of Concordance (Kendall's W) is a measure of association for discrete ordinal data, used for assessments that do not include a known reference standard. Kendall's coefficient of concordance ranges from 0 to 1: A coefficient value of 1 indicates perfect agreement. If the coefficient = 0, then the agreement is random, i.e., the same as would be expected by chance. **“Rule-of-thumb” interpretation guidelines: >= 0.9 very good agreement (green); 0.7 to < 0.9 marginally acceptable, improvement should be considered (yellow); < 0.7 unacceptable (red).** See Appendix [Kendall's Coefficient of Concordance](#) for further details on the Kendall Concordance calculations and “rule-of-thumb” interpretation guidelines.

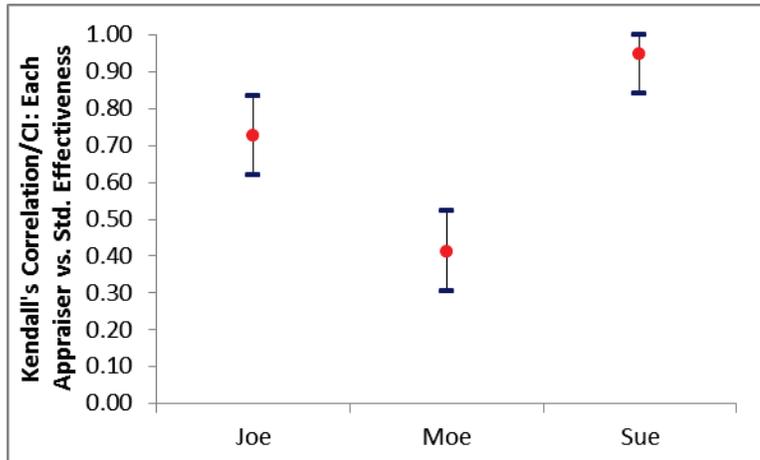
Kendall's Concordance P-Value: H0: Kendall's Coefficient of Concordance = 0. If P-Value < alpha (.05 for specified 95% confidence level), reject H0 and conclude that agreement is not the same as would be expected by chance. Significant P-Values are highlighted in red.

Kendall's Concordance LC (Lower Confidence) limit and Kendall's Concordance UC (Upper Confidence) limit cannot be solved analytically, so are estimated using bootstrapping. See Appendix [Kendall's Coefficient of Concordance](#) for further details on the bootstrap confidence intervals. Interpretation Guidelines: Concordance lower confidence limit >= 0.9: very good agreement. Concordance upper confidence limit < 0.7: the attribute agreement is unacceptable. Wide confidence intervals indicate that the sample size is inadequate.

The **Within Appraiser Agreement** for Joe is marginal, Moe is unacceptable and Sue is very good.

The **Between Appraiser Agreement** is unacceptable.

Kendall's Correlation/CI Each Appraiser vs. Standard Effectiveness Graph:



Each Appraiser vs. Standard Agreement Table:

Each Appraiser vs. Standard Agreement	# Inspected	# Matched	Percent	95.0% LC (Score)	95.0% UC (Score)	Kendall's Correlation Coeff.	Kendall's Correlation P-Value	Kendall's Correlation 95.0% LC	Kendall's Correlation 95.0% UC
Joe	50	11	22.00	12.75	35.24	0.7323	0.0000	0.6218	0.8427
Moe	50	2	4.00	1.10	13.46	0.4161	0.0000	0.3057	0.5265
Sue	50	40	80.00	66.96	88.76	0.9495	0.0000	0.8391	1.0000

All Appraisers vs. Standard Agreement Table:

All Appraisers vs. Standard Agreement	# Inspected	# Matched	Percent	95.0% LC (Score)	95.0% UC (Score)	Kendall's Correlation Coeff.	Kendall's Correlation P-Value	Kendall's Correlation 95.0% LC	Kendall's Correlation 95.0% UC
	50	1	2.00	0.35	10.50	0.6993	0.0000	0.6355	0.7631

Each Appraiser vs. Standard Effectiveness Table:

Each Appraiser vs. Standard Effectiveness	# Inspected	# Matched	Percent	95.0% LC (Score)	95.0% UC (Score)	Kendall's Correlation Coeff.	Kendall's Correlation P-Value	Kendall's Correlation 95.0% LC	Kendall's Correlation 95.0% UC
Joe	150	81	54.00	46.02	61.78	0.7280	0.0000	0.6201	0.8359
Moe	150	47	31.33	24.45	39.14	0.4137	0.0000	0.3058	0.5217
Sue	150	134	89.33	83.38	93.33	0.9478	0.0000	0.8399	1.0000

All Appraisers vs. Standard Effectiveness Table:

All Appraisers vs. Standard Effectiveness	# Inspected	# Matched	Percent	95.0% LC (Score)	95.0% UC (Score)	Kendall's Correlation Coeff.	Kendall's Correlation P-Value	Kendall's Correlation 95.0% LC	Kendall's Correlation 95.0% UC
	450	262	58.22	53.61	62.69	0.6830	0.0000	0.6211	0.7448

Kendall's Correlation Coefficient (Kendall's tau-b) is a measure of association for discrete ordinal data, used for assessments that include a known reference standard. Kendall's correlation coefficient ranges from -1 to 1: A coefficient value of 1 indicates perfect agreement. If the coefficient = 0, then the agreement is random, i.e., the same as would be expected by chance. A coefficient value of -1 indicates perfect disagreement. **“Rule-of-thumb” interpretation guidelines:** **>= 0.8 very good agreement (green); 0.6 to < 0.8 marginally acceptable, improvement should be considered (yellow); < 0.6 unacceptable (red).** See Appendix [Kendall's Correlation Coefficient](#) for further details on the Kendall Correlation calculations and “rule-of-thumb” interpretation guidelines.

Kendall's Correlation P-Value: H0: Kendall's Correlation Coefficient = 0. If P-Value < alpha (.05 for specified 95% confidence level), reject H0 and conclude that agreement is not the same as would be expected by chance. Significant P-Values are highlighted in red.

Kendall's Correlation LC (Lower Confidence) and **Kendall's Correlation UC** (Upper Confidence) limit use a normal approximation. Interpretation Guidelines: Correlation lower confidence limit ≥ 0.8 : very good agreement. Correlation upper confidence limit < 0.6 : the attribute agreement is unacceptable. Wide confidence intervals indicate that the sample size is inadequate.

Tip: Kendall's Correlation values in the Effectiveness tables are very similar to those in the Agreement tables (the slight difference is due to average Kendall for unstacked versus Kendall for stacked data). This is why the **Kendall's Correlation/CI Each Appraiser vs. Standard Agreement** graph is not shown. It would essentially be a duplicate of the **Kendall's Correlation/CI Each Appraiser vs. Standard Effectiveness** graph.

Appraiser Joe has marginal agreement versus the standard values. Appraiser Moe has unacceptable agreement to the standard. Sue has very good agreement to the standard.

Overall, the appraisers have marginal agreement to the standard.

Note that the Percent Agreement results in **All Appraisers vs. Standard Agreement Table** show only 2% agreement! This is due to the requirement that all appraisers agree with the standard across all trials for a 5-level response, which is very unlikely to occur. This highlights the problem with using Percent Agreement in an Ordinal MSA. Kendall's coefficients are the key metric to assess an Ordinal MSA.

Effectiveness and Misclassification Summary is a summary table of all appraisers' correct rating counts and misclassification counts compared to the known reference standard values.

Effectiveness and Misclassification Summary	Standard (1)	Standard (2)	Standard (3)	Standard (4)	Standard (5)
Appraiser Response (1)	102	20	3	2	0
Appraiser Response (2)	38	37	3	6	5
Appraiser Response (3)	8	14	19	20	6
Appraiser Response (4)	12	7	6	26	19
Appraiser Response (5)	2	3	5	9	78

Attribute MSA Data is a summary showing the original data in unstacked format. This makes it easy to visually compare appraiser results by part. If a reference standard is provided, the cells are color highlighted as follows: absolute deviation = 0 (green); absolute deviation = 1 (yellow); absolute deviation ≥ 2 (red):

Attribute MSA Data:

Sample No.	Expert Reference	Joe 1	Joe 2	Joe 3	Moe 1	Moe 2	Moe 3	Sue 1	Sue 2	Sue 3
1	4	5	4	5	4	4	3	0	4	4
2	2	3	3	2	3	3	2	2	2	2
3	5	5	5	5	5	4	4	5	5	5
4	4	3	3	3	3	3	4	3	4	4
5	3	4	3	3	2	1	3	3	3	3
6	1	2	1	2	2	4	2	1	1	1
7	1	1	1	1	2	1	1	1	1	1
8	1	1	2	1	2	1	1	1	1	1
9	5	4	5	4	5	5	5	5	5	5
10	1	1	2	2	1	4	1	1	1	1

In conclusion, this measurement system is marginal and should be improved. Appraiser Moe needs training and Appraiser Joe needs a refresher. Sue has very good agreement based on Kendall's Concordance and Correlation, but would have been considered marginal based on Kappa ($< .9$) and Percent Effectiveness ($< 95\%$). As discussed above, Kappa and Percent Effectiveness do not take the order of the response data into account, so are not as useful as Kendall's coefficients in an Ordinal MSA study.

Attribute MSA (Nominal)

Attribute MSA is also known as Attribute Agreement Analysis. Use the Nominal option if the assessed result is numeric or text nominal (e.g., Defect Type 1, Defect Type 2, Defect Type 3). There must be at least 3 response levels in the assessed result, otherwise it is binary.

1. Open the file **Attribute MSA – Nominal.xlsx**. This is a Nominal MSA example with 30 samples, 3 appraisers and 2 trials. The response is text “Type_1”, “Type_2” and “Type_3.” The *Expert Reference* column is the reference standard from an expert appraisal. Note that the worksheet data must be in stacked column format and the reference values must be consistent for each sample.
2. Click **SigmaXL > Measurement Systems Analysis > Attribute MSA (Nominal)**. Ensure that the entire data table is selected. Click **Next**.
3. Select *Sample No.*, *Appraiser*, *Assessed Result* and *Expert Reference* as shown. Check **Report Information** and enter *Attribute MSA – Nominal* for **Product/Unit Name**. Select **Percent Confidence Interval Type – Wilson Score**:

4. Click **OK**. The Attribute MSA Nominal Analysis Report is produced.

Tip: While this report is quite extensive, a quick assessment of the attribute measurement system can be made by viewing the Kappa color highlights: **Green** - very good agreement (Kappa ≥ 0.9); **Yellow** - marginally acceptable, improvement should be considered (Kappa 0.7 to < 0.9); **Red** - unacceptable (Kappa < 0.7). See [Attribute MSA \(Binary\)](#) for a detailed discussion of the report graphs and tables. Here we will just look at Kappa.

Attribute Agreement Report:

Attribute Agreement Report:										
Within Appraiser Agreement	# Inspected	# Matched	Percent	95.0% LC (Score)	95.0% UC (Score)	Fleiss' Kappa Overall Response	Fleiss' Kappa Overall Response P-Value	Fleiss' Kappa Overall Response 95.0% LC	Fleiss' Kappa Overall Response 95.0% UC	Fleiss' Kappa Response (Type_1)
A	30	30	100.00	88.65	100.00	1.0000	0.0000	0.7333	1.0000	1.0000
B	30	29	96.67	83.33	99.41	0.9467	0.0000	0.6799	1.0000	1.0000
C	30	26	86.67	70.32	94.69	0.7849	0.0000	0.5184	1.0000	1.0000
Each Appraiser vs. Standard Agreement	# Inspected	# Matched	Percent	95.0% LC (Score)	95.0% UC (Score)	Fleiss' Kappa Overall Response	Fleiss' Kappa Overall Response P-Value	Fleiss' Kappa Overall Response 95.0% LC	Fleiss' Kappa Overall Response 95.0% UC	Fleiss' Kappa Response (Type_1)
A	30	29	96.67	83.33	99.41	0.9467	0.0000	0.7581	1.0000	1.0000
B	30	29	96.67	83.33	99.41	0.9733	0.0000	0.7847	1.0000	1.0000
C	30	25	83.33	66.44	92.66	0.8391	0.0000	0.6506	1.0000	1.0000
Between Appraiser Agreement	# Inspected	# Matched	Percent	95.0% LC (Score)	95.0% UC (Score)	Fleiss' Kappa Overall Response	Fleiss' Kappa Overall Response P-Value	Fleiss' Kappa Overall Response 95.0% LC	Fleiss' Kappa Overall Response 95.0% UC	Fleiss' Kappa Response (Type_1)
	30	25	83.33	66.44	92.66	0.9110	0.0000	0.8421	0.9798	1.0000
All Appraisers vs. Standard Agreement	# Inspected	# Matched	Percent	95.0% LC (Score)	95.0% UC (Score)	Fleiss' Kappa Overall Response	Fleiss' Kappa Overall Response P-Value	Fleiss' Kappa Overall Response 95.0% LC	Fleiss' Kappa Overall Response 95.0% UC	Fleiss' Kappa Response (Type_1)
	30	25	83.33	66.44	92.66	0.9197	0.0000	0.8108	1.0000	1.0000

Attribute Effectiveness Report:

Attribute Effectiveness Report:										
Each Appraiser vs. Standard Effectiveness	# Inspected	# Matched	Percent	95.0% LC (Score)	95.0% UC (Score)	Fleiss' Kappa Overall Response	Fleiss' Kappa Overall Response P-Value	Fleiss' Kappa Overall Response 95.0% LC	Fleiss' Kappa Overall Response 95.0% UC	Fleiss' Kappa Response (Type_1)
A	60	58	96.67	88.64	99.08	0.9467	0.0000	0.7581	1.0000	1.0000
B	60	59	98.33	91.14	99.71	0.9733	0.0000	0.7847	1.0000	1.0000
C	60	54	90.00	79.85	95.34	0.8391	0.0000	0.6513	1.0000	1.0000
All Appraisers vs. Standard Effectiveness	# Inspected	# Matched	Percent	95.0% LC (Score)	95.0% UC (Score)	Fleiss' Kappa Overall Response	Fleiss' Kappa Overall Response P-Value	Fleiss' Kappa Overall Response 95.0% LC	Fleiss' Kappa Overall Response 95.0% UC	Fleiss' Kappa Response (Type_1)
	180	171	95.00	90.77	97.35	0.9200	0.0000	0.8111	1.0000	1.0000
Effectiveness and Misclassification Summary	Standard (Type_1)	Standard (Type_2)	Standard (Type_3)							
Appraiser Response (Type_1)	30	0	0							
Appraiser Response (Type_2)	0	70	1							
Appraiser Response (Type_3)	0	8	71							

Fleiss' Kappa statistic is a measure of agreement that is analogous to a “correlation coefficient” for discrete data. Kappa ranges from -1 to +1: A Kappa value of +1 indicates perfect agreement. If Kappa = 0, then agreement is the same as would be expected by chance. If Kappa = -1, then there is perfect disagreement. **“Rule-of-thumb” interpretation guidelines: ≥ 0.9 very good agreement (green); 0.7 to < 0.9 marginally acceptable, improvement should be considered (yellow); < 0.7 unacceptable (red).** See Appendix [Kappa](#) for further details on the Kappa calculations and “rule-of-thumb” interpretation guidelines.

Fleiss' Kappa P-Value: H_0 : Kappa = 0. If P-Value < alpha (.05 for specified 95% confidence level), reject H_0 and conclude that agreement is not the same as would be expected by chance. Significant P-Values are highlighted in red.

Fleiss' Kappa LC (Lower Confidence) and **Fleiss' Kappa UC** (Upper Confidence) limits use a kappa normal approximation. Interpretation Guidelines: Kappa lower confidence limit ≥ 0.9 : very good agreement. Kappa upper confidence limit < 0.7 : the attribute agreement is unacceptable. Wide confidence intervals indicate that the sample size is inadequate.

Fleiss' Kappa Overall is an overall Kappa for all of the response levels.

Fleiss' Kappa Individual gives Kappa for each response level (*Type_1*, *Type_2* and *Type_3*). This is useful to identify if an appraiser has difficulty assessing a particular defect type.

This is a very good attribute measurement system, but Appraiser C is marginal so a refresher would be helpful.

Part J – Process Capability

Process Capability Templates

Click **SigmaXL > Templates and Calculators > Basic Process Capability Templates > Process Sigma Level – Discrete Data** or
Process Sigma Level – Continuous Data or
Process Capability Indices or
Process Capability & Confidence Intervals or
Tolerance Interval Calculator (Normal Exact).

These templates are also located at **SigmaXL > Process Capability > Basic Process Capability Templates**.

See **Measure Phase Part B – Templates and Calculators** for Process Capability Template examples:

[Basic Process Capability Templates – Process Sigma Level – Discrete Data](#)

[Basic Process Capability Templates – Process Sigma Level – Continuous Data](#)

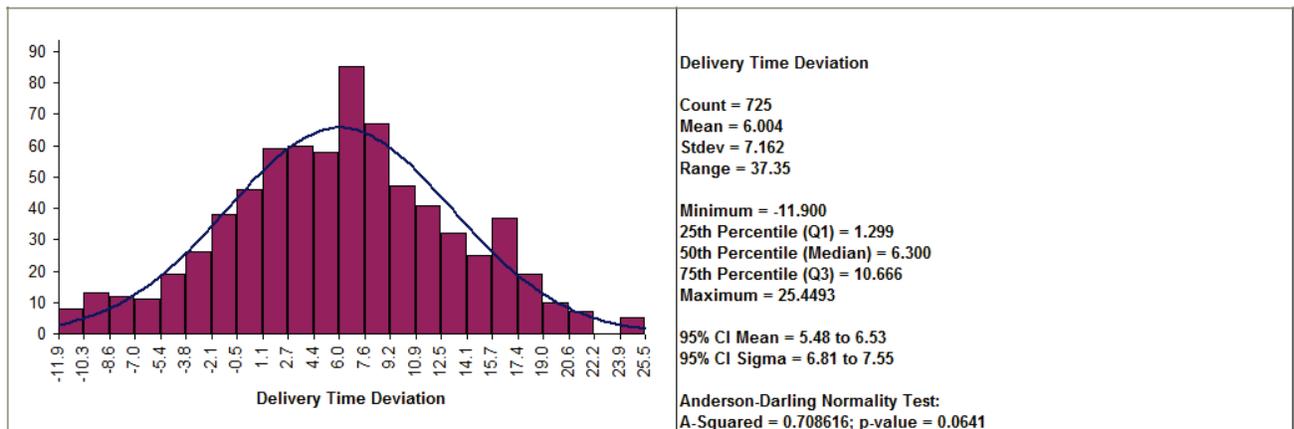
[Basic Process Capability Templates – Process Capability Indices](#)

[Basic Process Capability Templates – Process Capability & Confidence Intervals](#)

[Basic Process Capability Templates – Tolerance Interval Calculator \(Normal Exact\)](#)

Histograms and Process Capability

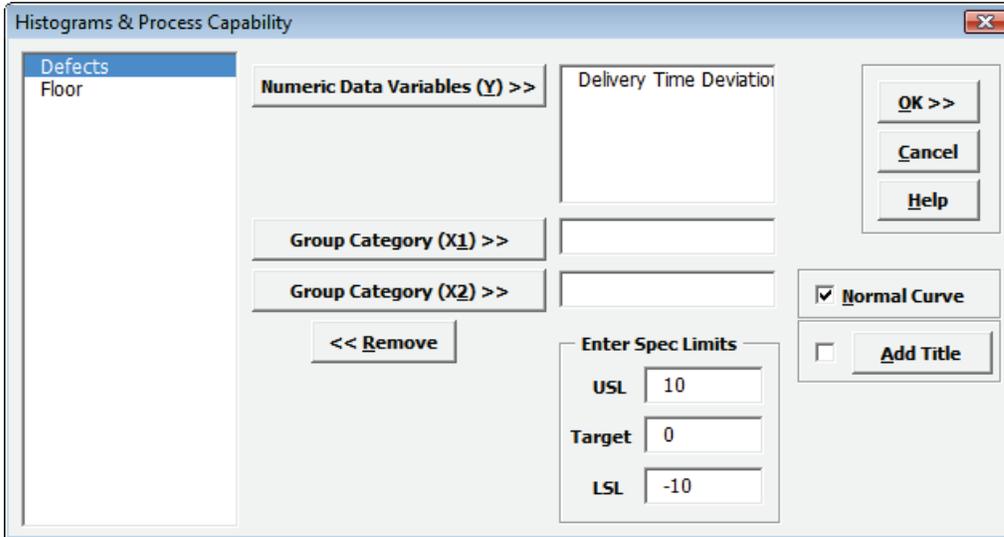
1. Open the file **Delivery Times.xlsx**. This contains continuous data of hotel breakfast delivery times. Deviation Time is the deviation around targeted delivery time in minutes. The Critical Customer Requirements (CCR's) are as follows: USL = 10 minutes late, LSL = -10 minutes (early).
2. Let's begin with a view of the data using Histograms and Descriptive Statistics. Click **SigmaXL > Graphical Tools > Histograms & Descriptive Statistics**.
3. Ensure that entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
4. Select *Delivery Time Deviation*, click **Numeric Data Variable (Y) >>**. Click **OK**.



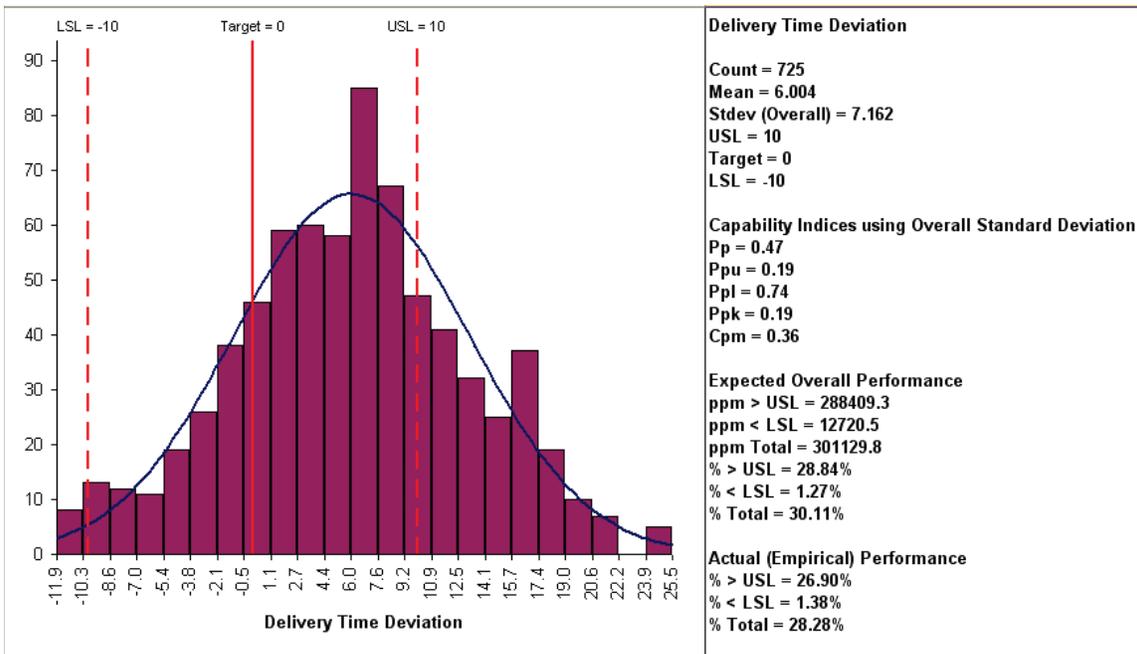
Note that the data has an average of 6, so on average the breakfast delivery is 6 minutes late. Also note the wide variation in delivery times (StDev = 7.2 minutes). With the AD P-Value > .05, this data can be assumed to have a normal distribution.

SigmaXL: Measure Phase Tools

- Click **Sheet 1** Tab. Click **SigmaXL > Process Capability > Histograms & Process Capability (or SigmaXL > Graphical Tools > Histograms & Process Capability)**.
- Select *Delivery Time Deviation*, click **Numeric Data Variable (Y) >>**. Enter **USL = 10, Target = 0, LSL = -10**, check **Normal Curve** as shown below:



- Click **OK**. Resulting Histogram & Process Capability report is shown below:



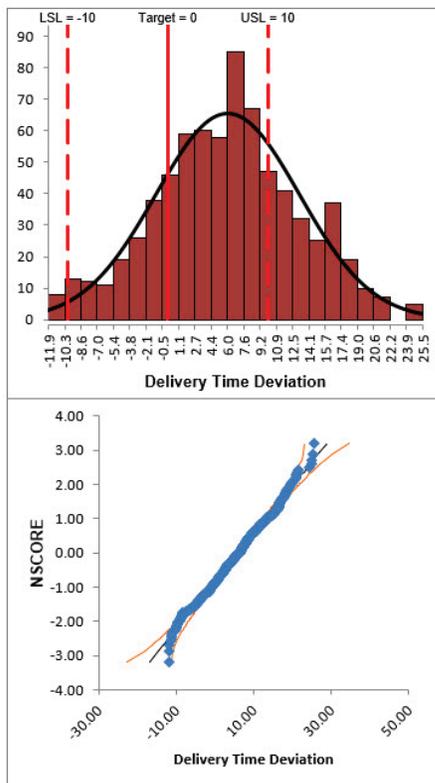
With the Process Performance indices Pp and Ppk < 1, this process is clearly in need of improvement. Note that the difference between Pp and Ppk is due to the off-center process mean. (Cp and Cpk indices are provided in the Capability Combination Report or optionally when creating Control Charts).

Capability Combination Report (Individuals)

1. Open the file **Delivery Times.xlsx**.
2. Click **SigmaXL >Process Capability > Capability Combination Report (Individuals)**.
3. Select *Delivery Time Deviation*, click **Numeric Data Variable (Y) >>**. Enter **USL = 10, Target = 0, LSL = -10**, as shown below:

The screenshot shows the 'Capability Combination Report (Individuals)' dialog box. In the top left, a list contains 'Delivery Time Deviation', 'Defects', and 'Floor', with 'Delivery Time Deviation' selected. The 'Numeric Data Variable (Y) >>' field contains 'Delivery Time Deviation'. Below it, the 'Optional X-Axis Labels >>' field is empty. The '<< Remove' button is visible. On the right, there are 'OK >>', 'Cancel', and 'Help' buttons. The main area is divided into two sections: 'Individuals/Moving Range Control Chart Options' and 'Spec Limits for Capability'. In the first section, 'Calculate Control Limits' is selected, and 'Tests for Special Causes' and 'Add Title' are unchecked. The 'Individuals' section has empty fields for UCL, CL, and LCL. The 'Moving Range' section has empty fields for UCL and LCL. The 'Estimate' section has 'Average MR' selected, and the 'Individual CL' section has 'Mean' selected. The 'Spec Limits for Capability' section shows USL = 10, Target = 0, and LSL = -10.

4. Click **OK**. The resulting Process Capability Combination report is shown below:



Process Capability Report: Delivery Time Deviation	
Count	725
Mean	6.004
StDev (Overall, Long Term)	7.162
StDev (Within, Short Term)	7.202
USL	10
Target	0
LSL	-10

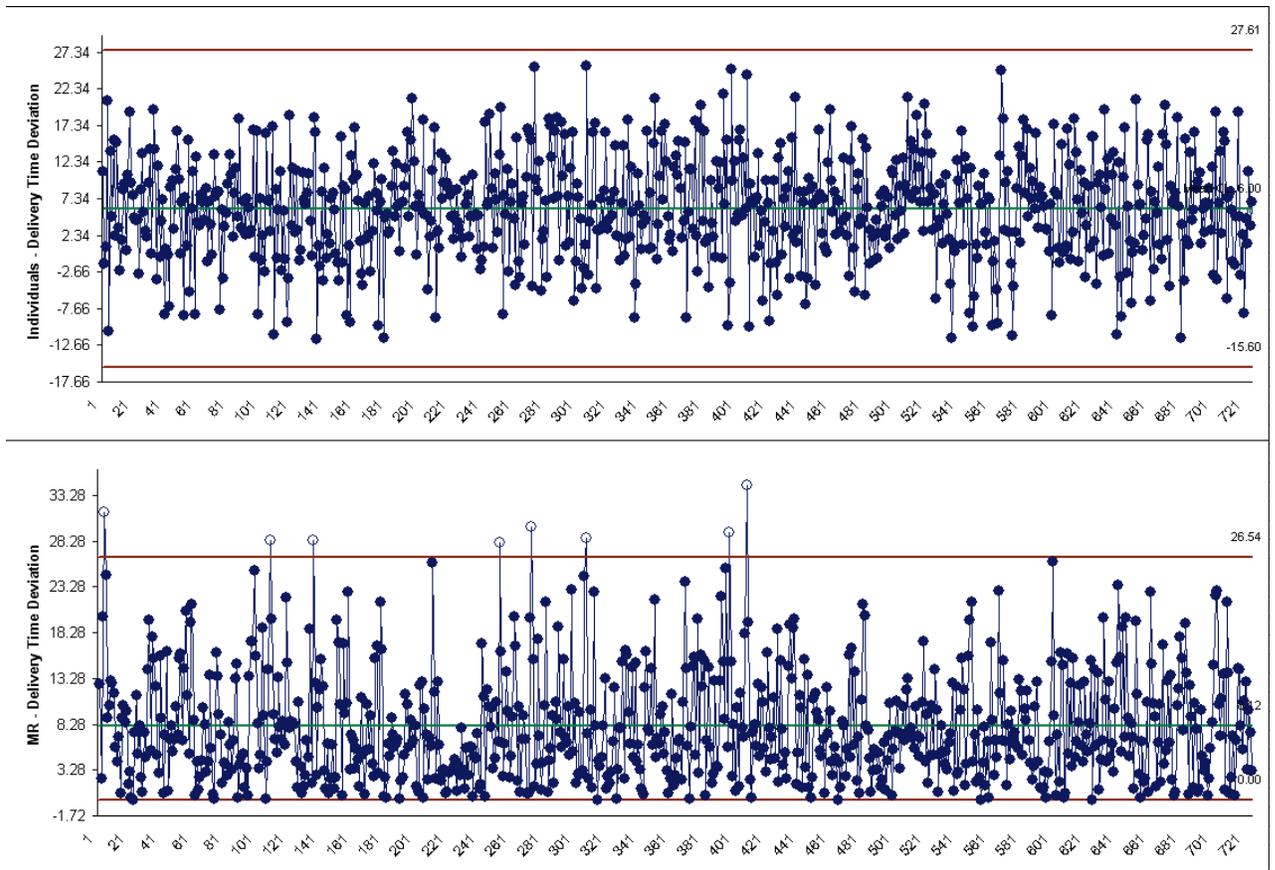
Capability Indices using Overall StDev	
Pp	0.47
Ppu	0.19
Ppl	0.74
Ppk	0.19
Cpm	0.36

Potential Capability Indices using Within StDev	
Cp	0.46
Cpu	0.18
Cpl	0.74
Cpk	0.18

Expected Overall Performance	
ppm > USL	288409.3
ppm < LSL	12720.5
ppm Total	301130
% > USL	28.84%
% < LSL	1.27%
% Total	30.11%

Actual (Empirical) Performance	
% > USL	26.90%
% < LSL	1.38%
% Total	28.28%

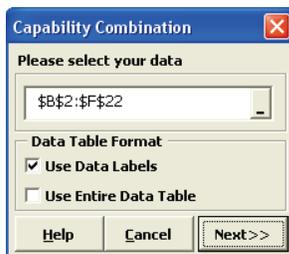
Anderson-Darling Normality Test	
A-Squared	0.708616
P-Value	0.0641



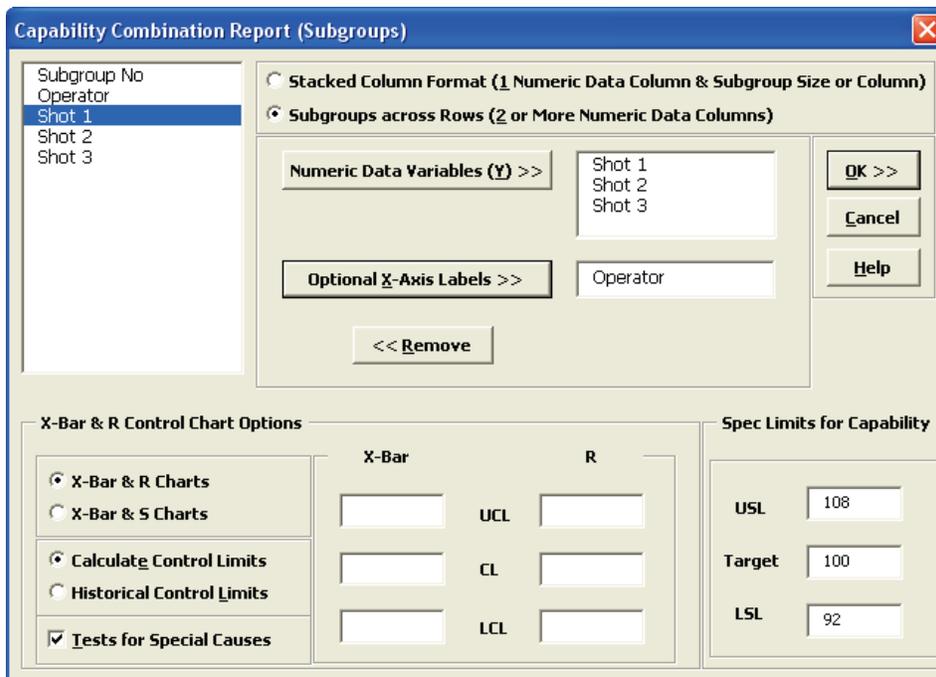
5. Cp is the “best case” potential capability index. Ppk is the “worst case” actual performance. If Cp is excellent (say $C_p = 2$) and Ppk is poor (say $P_{pk} < 1$), this indicates that the process can be dramatically improved through centering and stabilizing. If Cp is poor, you will have to implement fundamental improvements using DMAIC to address common cause variation. If $P_{pk} = C_p$, this indicates that the process is centered and stable, and the short-term (within subgroup) variation is the same as the long-term (overall) variation.
6. Cpm is a similar statistic but it incorporates a penalty for deviation from target.
7. On a technical note regarding compatibility with other software tools, please be advised that SigmaXL does not use unbiasing constants when calculating Overall StDev, Pp, Ppk and Cpm. This is done in order to ensure that the calculated Overall StDev matches the results given in Descriptive Statistics. Note that Within StDev, Cp and Cpk are based on control chart techniques (MR-bar/d2 for Individuals Chart, R-bar/d2 for X-Bar & R, S-bar/c4 for X-Bar & S).

Capability Combination Report (Subgroups)

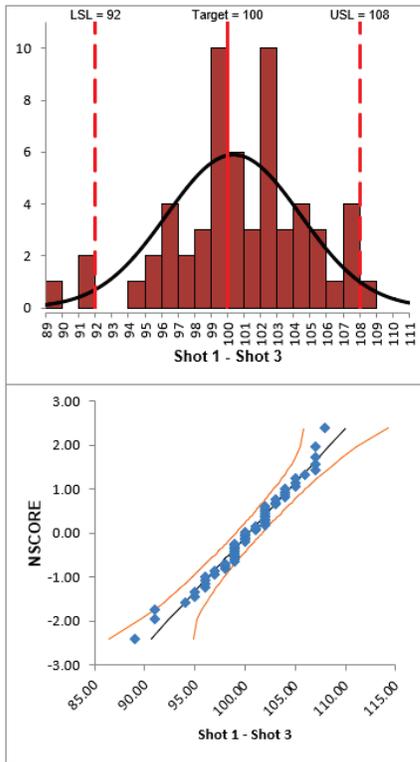
1. Open the file **Catapult Data – Xbar Control Charts.xlsx**. Each operator fires the catapult ball 3 times. The target distance is 100 inches. The Upper Specification Limit (USL) is 108 inches. The Lower Specification Limit (LSL) is 92 inches.
2. Select **B2:F22**; here, we will only use the first 20 subgroups to determine the process capability (subgroups 21 to 25 are studied later in **Control Phase Tools: Statistical Process Control (SPC) Charts: Part B - X-Bar & Range Charts**).
3. Select **SigmaXL > Process Capability > Capability Combination Report (Subgroups)**.
4. Do not check **Use Entire Data Table**!



5. Click **Next**. Select **Subgroups across Rows**, select *Shot 1*, *Shot 2*, *Shot 3*, click **Numeric Data Variables (Y) >>**. Select *Operator*, click **Optional X-Axis Labels >>**. Select **X-Bar & R Charts**. Check **Tests for Special Causes**. Enter **USL = 108**, **Target = 100**, **LSL = 92**, as shown:



6. Click **OK**. The resulting Process Capability Combination report is shown below:



Process Capability Report: Shot 1 - Shot 3	
Count	60
Mean	100.37
StDev (Overall, Long Term)	4.055
StDev (Within, Short Term)	3.721
USL	108
Target	100
LSL	92

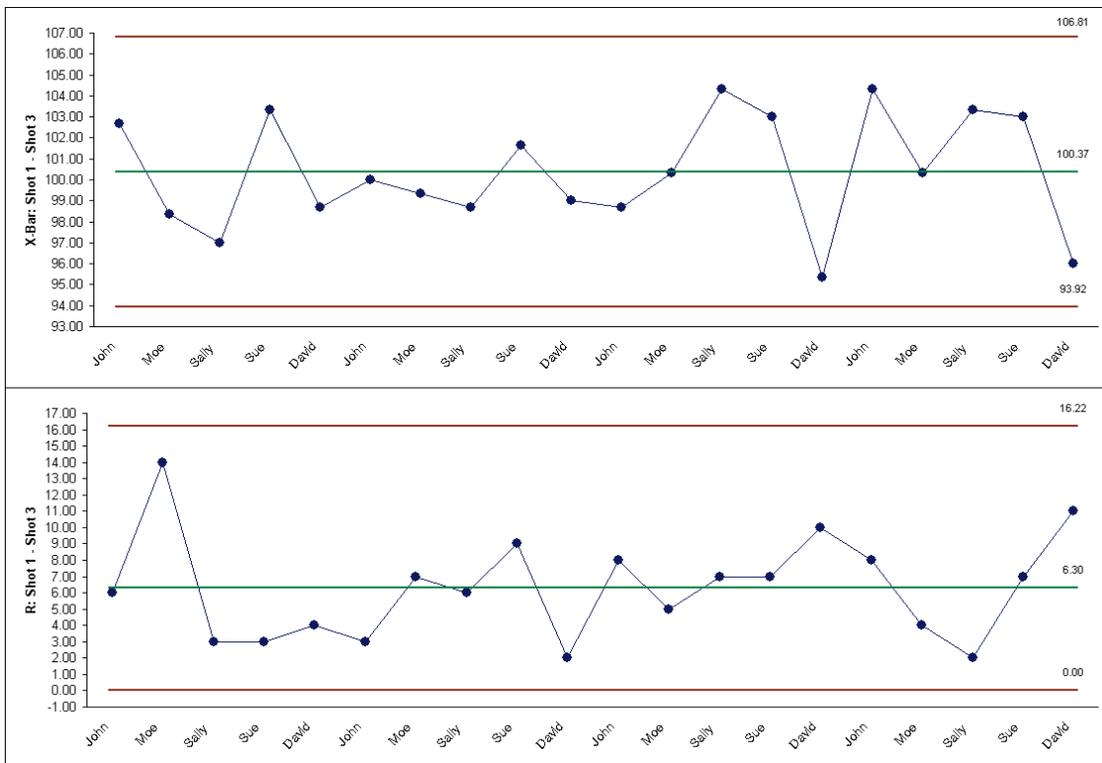
Capability Indices using Overall StDev	
Pp	0.66
Ppu	0.63
Ppl	0.69
Ppk	0.63
Cpm	0.66

Potential Capability Indices using Within StDev	
Cp	0.72
Cpu	0.68
Cpl	0.75
Cpk	0.68

Expected Overall Performance	
ppm > USL	29874.0
ppm < LSL	19531.9
ppm Total	49406
% > USL	2.99%
% < LSL	1.95%
% Total	4.94%

Actual (Empirical) Performance	
% > USL	
% < LSL	5.00%
% Total	5.00%

Anderson-Darling Normality Test	
A-Squared	0.530155
P-Value	0.1689



Tests for Special Causes - X-Bar: Shot 1 - Shot 3
 Number of Data Points Failing Tests = 0

Capability Combination Report (Individuals Nonnormal)

An important assumption for process capability analysis is that the data be normally distributed. The Capability Combination Report (Individuals Nonnormal) allows you to transform the data to normality or utilize nonnormal distributions, including:

Box-Cox Transformation (includes an automatic threshold option so that data with negative values can be transformed)

Johnson Transformation

Distributions supported:

- Half-Normal
- Lognormal (2 & 3 parameter)
- Exponential (1 & 2 parameter)
- Weibull (2 & 3 parameter)
- Beta (2 & 4 parameter)
- Gamma (2 & 3 parameter)
- Logistic
- Loglogistic (2 & 3 parameter)
- Largest Extreme Value
- Smallest Extreme Value

Automatic Best Fit based on AD P-Value

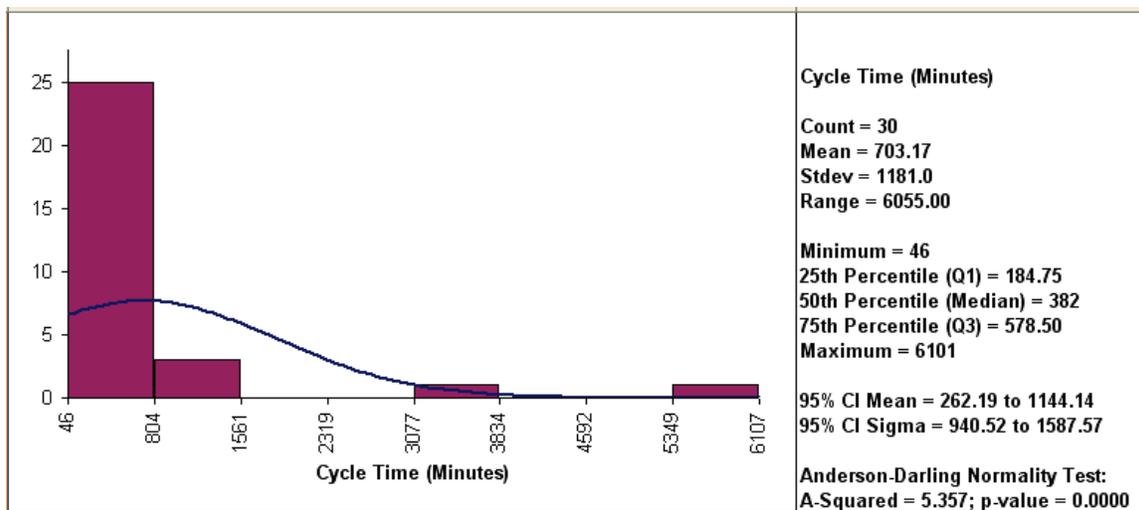
For technical details, see Appendix: [Statistical Details for Nonnormal Distributions and Transformations](#). See also Andrew Sleeper, *Six Sigma Distribution Modeling*, for further information on these methods.

Note that these transformations and distributions are particularly effective for inherently skewed data but should not be used with bimodal data or where the nonnormality is due to outliers (typically identified with a Normal Probability Plot). In these cases, you should identify the reason for the bimodality or outliers and take corrective action. Another common reason for nonnormal data is poor measurement discrimination leading to “chunky” data. In this case, attempts should be made to improve the measurement system.

Box-Cox Transformation

SigmaXL’s default setting is to use the Box-Cox transformation which is the most common approach to dealing with nonnormal data. Box-Cox is used to convert nonnormal data to normal by applying a power transformation, Y^λ , where λ varies from -5 to +5. You may select rounded or optimal λ . Rounded is typically preferred since it will result in a more “intuitive” transformation such as $\ln(Y)$ ($\lambda=0$) or \sqrt{Y} ($\lambda=0.5$). If the data includes zero or negative values, select λ & Threshold. SigmaXL will solve for an optimal threshold which is a shift factor on the data so that all of the values are positive.

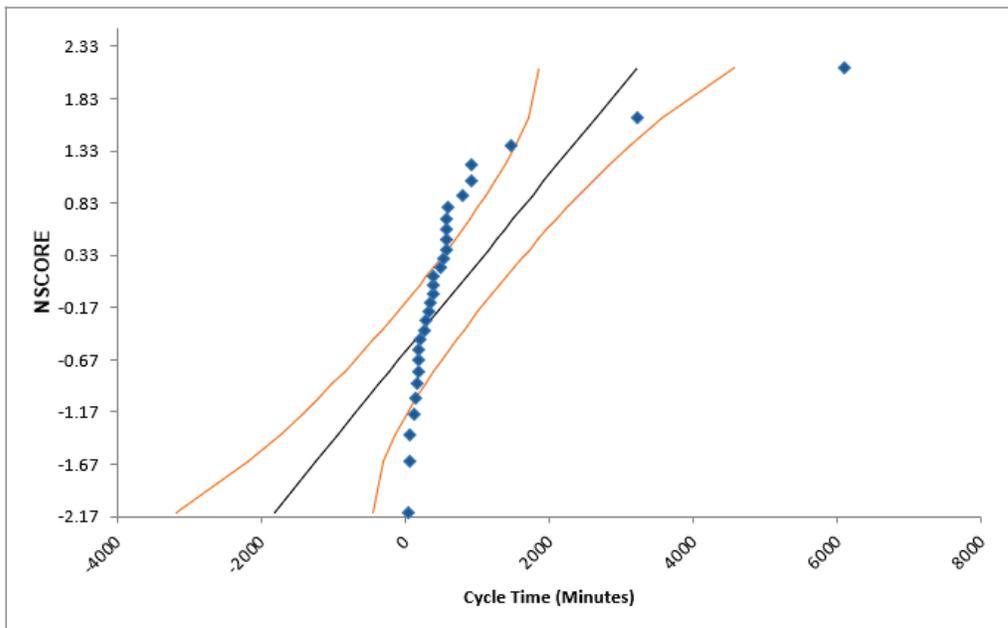
1. Open the file **Nonnormal Cycle Time2.xlsx**. This contains continuous data of process cycle times. The Critical Customer Requirement is: USL = 1000 minutes.
2. Let’s begin with a view of the data using Histograms and Descriptive Statistics. Click **SigmaXL > Graphical Tools > Histograms & Descriptive Statistics**.
3. Ensure that entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
4. Select *Cycle Time (Minutes)*, click **Numeric Data Variable (Y) >>**. Click **OK**.



Clearly this is a process in need of improvement. To start, we would like to get a baseline process capability. The problem with using regular Capability analysis is that the results will be incorrect due to the nonnormality in the data. The Histogram and AD P-Value < .05 clearly show that this data is not normal.

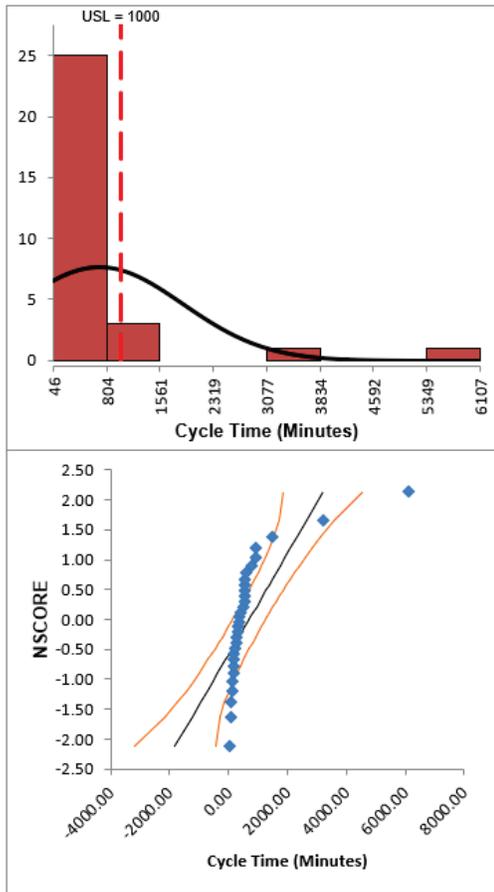
5. We will confirm the nonnormality by using a Normal Probability Plot. Click **Sheet 1 Tab (or F4)**. Click **SigmaXL > Graphical Tools > Normal Probability Plots**.
6. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.

7. Select *Cycle Time (Minutes)*, click **Numeric Data Variable (Y) >>**. Click **OK**. A Normal Probability Plot of Cycle Time data is produced:



8. The curvature in this normal probability plot confirms that this data is not normal.
9. For now, let us ignore the nonnormal issue and perform a Process Capability study assuming a normal distribution. Click **Sheet 1** Tab. Click **SigmaXL > Process Capability > Capability Combination Report (Individuals)**.
10. Select *Cycle Time (Minutes)*, click **Numeric Data Variable (Y) >>**. Enter **USL = 1000**; delete previous **Target** and **LSL** settings.

11. Click **OK**. The resulting Process Capability Report is shown below:



Process Capability Report: Cycle Time (Minutes)	
Count	30
Mean	703.17
StDev (Overall, Long Term)	1181.0
StDev (Within, Short Term)	709.10
USL	1000
Target	
LSL	

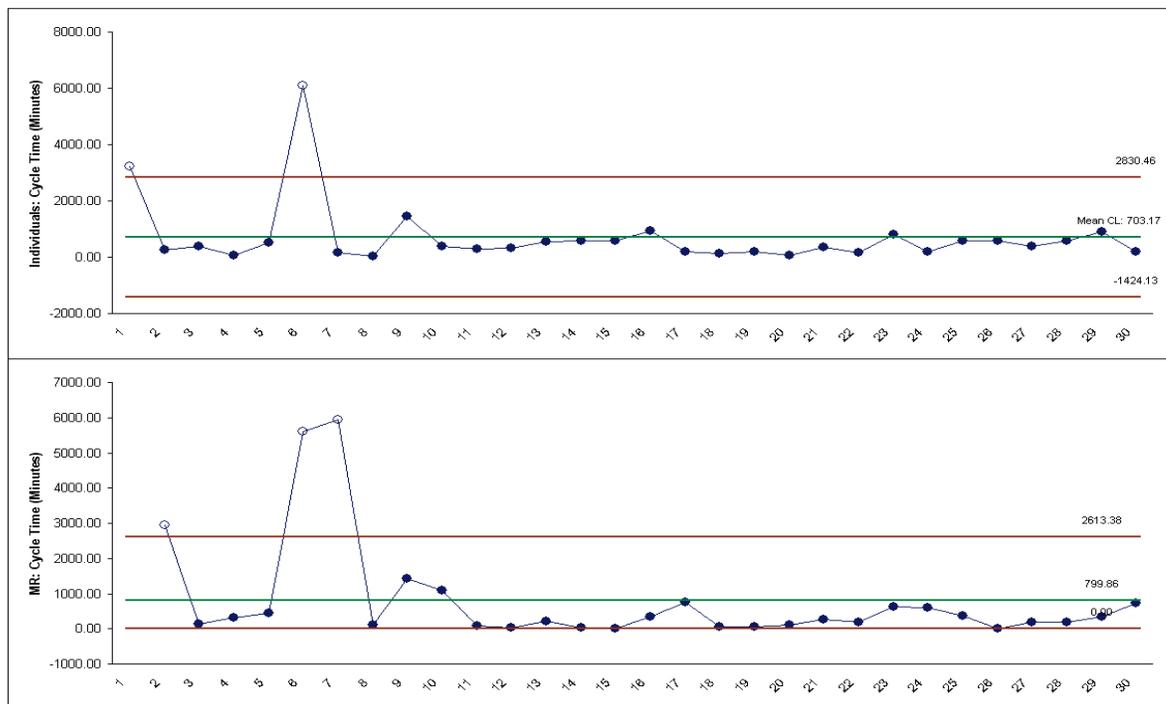
Capability Indices using Overall StDev	
Pp	
Ppu	0.08
Ppl	
Ppk	0.08
Cpm	

Potential Capability Indices using Within StDev	
Cp	
Cpu	0.14
Cpl	
Cpk	0.14

Expected Overall Performance	
ppm > USL	400771.3
ppm < LSL	
ppm Total	400771
% > USL	40.08%
% < LSL	
% Total	40.08%

Actual (Empirical) Performance	
% > USL	10.00%
% < LSL	
% Total	10.00%

Anderson-Darling Normality Test	
A-Squared	5.357
P-Value	0.0000



Notice the discrepancy between the Expected Overall (Theoretical) Performance and Actual (Empirical) Performance. This is largely due to the nonnormality in the data, since the expected performance assumes that the data is normal. So why not just use the actual performance and disregard the expected? This would not be reliable because the sample size, $n = 30$, is too small to estimate performance using pass/fail (discrete) criteria.

Also note that the process appears to be out-of-control on both the individuals and moving range charts.

12. We will now perform a process capability analysis using the Capability Combination Report for Nonnormal Individuals. Click **Sheet 1** Tab (or **F4**). Click **SigmaXL > Process Capability > Nonnormal > Capability Combination Report (Individuals Nonnormal)**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
13. Select *Cycle Time (Minutes)*, click **Numeric Data Variable (Y) >>**. Enter **USL = 1000**. We will use the default selection for **Transformation/Distribution Options: Box-Cox Transformation with Rounded Lambda**. Check **Tests for Special Causes** as shown:

Capability Combination Report (Individuals - Nonnormal)

Cycle Time (Minutes)

Numeric Data Variable (Y) >> Cycle Time (Minutes) OK >>

Optional X-Axis Labels >> Cancel

<< Remove Help

Control Chart Options

Individuals - Original Data

Individuals/Moving Range - Normalized Data

Tests for Special Causes

Add Title

Spec Limits for Capability

USL 1000

Target

LSL

Capability Index Options

Z-Score Method

Percentile (ISO) Method

Transformation/Distribution Options

Box-Cox Transformation

Johnson Transformation

Automatic Best Fit

Specify Distribution

Box-Cox Transformation Options

Rounded Lambda

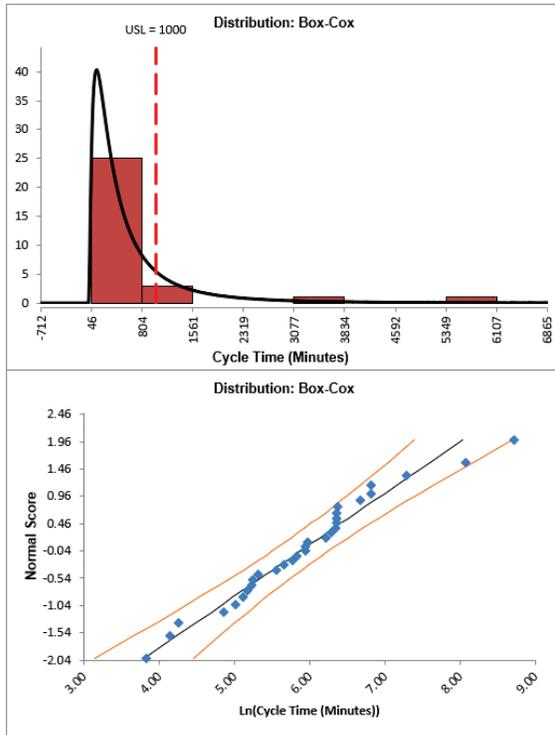
Optimal Lambda

Lambda & Threshold (Shift)

Optional Threshold Value

Optional Lambda Value

14. Click **OK**. The resulting Process Capability Combination report is shown below:



Process Capability Report (Nonnormal): Cycle Time (Minutes)	
Box-Cox Transformation: $Z = \ln(Y)$	
Sample Count	30
Sample Mean	703.17
USL	1000
Target	
LSL	
Lambda (Rounded)	0
Mean (Transformed)	5.920
StDev (Transformed Overall, Long Term)	1.061030703
StDev (Transformed Within, Short Term)	1.004837519
USL (Transformed)	6.908
Target (Transformed)	
LSL (Transformed)	

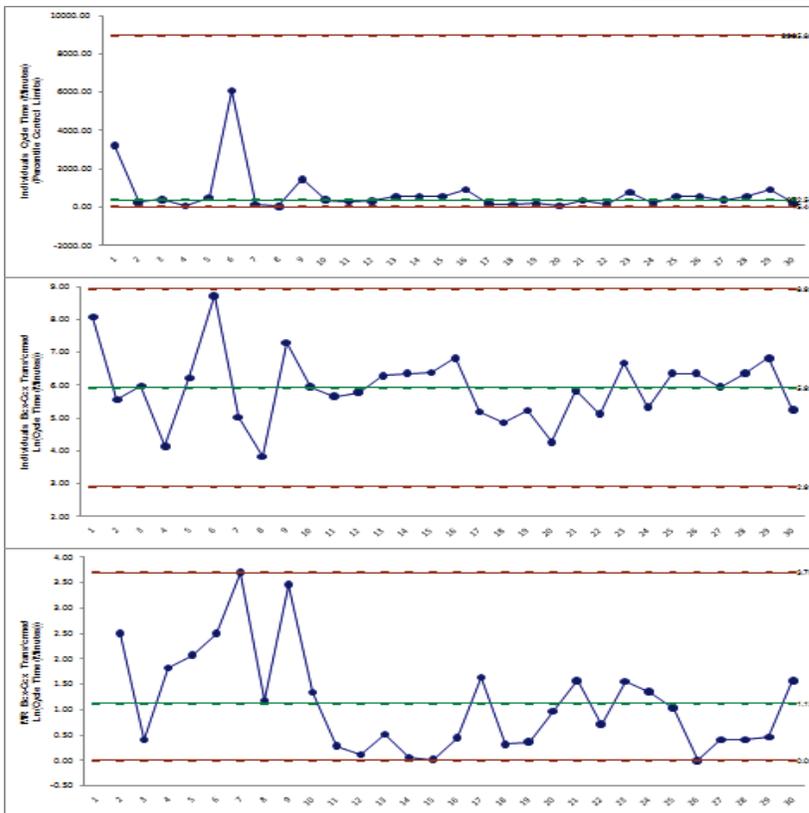
Capability Indices using Transformed Overall StDev	
Pp	
Ppu	0.31
Ppl	
Ppk	0.31

Potential Capability Indices using Transformed Within StDev	
Cp	
Cpu	0.33
Cpl	
Cpk	0.33

Expected Overall Performance	
ppm > USL	176026.1
ppm < LSL	
ppm Total	176026.1
% > USL	17.60%
% < LSL	
% Total	17.60%

Actual (Empirical) Performance	
% > USL	10.00%
% < LSL	
% Total	10.00%

Anderson-Darling Normality Tests	
AD Normality Transformed Data	0.369554
AD Normality P-Value Transformed Data	0.4041
AD Normality Original Data	5.357
AD Normality P-Value Original Data	0.0000



Tests for Special Causes - Indv: Cycle Time (Minutes)
Number of Data Points Failing Tests = 0

The **AD Normality P-Value Transformed Data** value of 0.404 confirms that the Box-Cox transformation to normality was successful. The process capability indices and expected performance can now be used to establish a baseline performance. Note that there are no out-of-control signals on the control charts, so the signals observed earlier when normality was assumed were false alarms.

The **Individuals – Original Data** chart displays the untransformed data with control limits calculated as:

UCL = 99.865 percentile

CL = 50th percentile

LCL = 0.135 percentile

The benefit of displaying this chart is that one can observe the original untransformed data. Since the control limits are based on percentiles, this represents the overall, long term variation rather than the typical short-term variation. The limits will likely be nonsymmetrical.

The **Individuals/Moving Range – Normalized Data** chart displays the transformed z-values with control limits calculated using the standard Shewhart formulas for Individuals and Moving Range charts. The benefit of using this chart is that tests for special causes can be applied and the control limits are based on short term variation. The disadvantage is that one is observing transformed data on the chart rather than the original data.

Automatic Best Fit

Now we will redo the capability analysis using the **Automatic Best Fit option**.

15. Click **Recall SigmaXL Dialog** menu or press **F3** to recall last dialog. Select **Automatic Best Fit** as shown:

SigmaXL: Measure Phase Tools

Capability Combination Report (Individuals - Nonnormal)

Cycle Time (Minutes)

Numeric Data Variable (Y) >> Cycle Time (Minutes) **OK >>**

Optional X-Axis Labels >> **Cancel**

<< Remove **Help**

Control Chart Options

Individuals - Original Data

Individuals/Moving Range - Normalized Data

Tests for Special Causes

Add Title

Spec Limits for Capability

USL 1000

Target

LSL

Capability Index Options

Z-Score Method

Percentile (ISO) Method

Transformation/Distribution Options

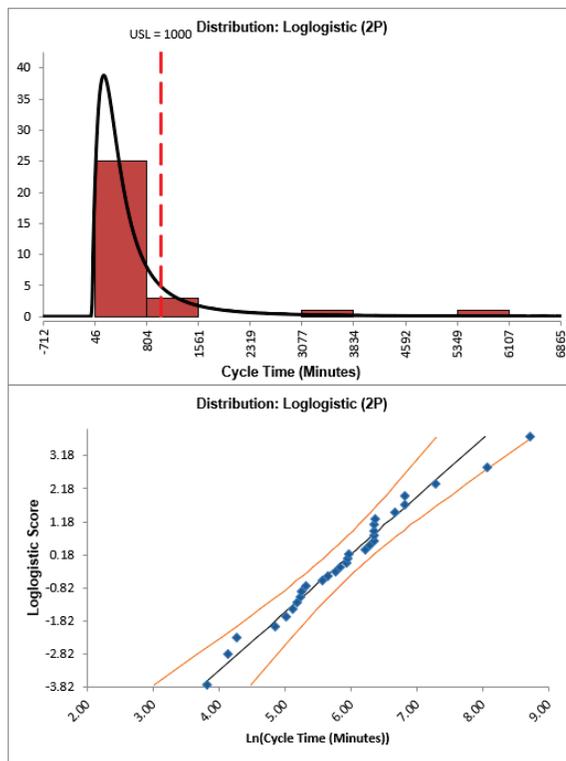
Box-Cox Transformation

Johnson Transformation

Automatic Best Fit

Specify Distribution

16. Click **OK**. The resulting Process Capability Combination report is shown below. Please note that due to the extensive computations required, this could take up to 1 minute (or longer for large datasets):



Process Capability Report (Nonnormal): Cycle Time (Minutes)	
Distribution: Logistic (2P)	
Sample Count	30
Sample Mean	703.17
Location	5.895
Scale	0.571471221
StDev (Transformed Overall, Long Term)	1
StDev (Normalized Within, Short Term)	0.964388
USL	1000
Target	
LSL	

Capability Indices (Z-Score Method)	
Pp	
Ppu	0.35
Ppl	
Ppk	0.35

Potential Capability Indices using Normalized Within StDev	
Cp	
Cpu	0.37
Cpl	
Cpk	0.37

Expected Overall Performance	
ppm > USL	145274.6
ppm < LSL	
ppm Total	145274.6
% > USL	14.53%
% < LSL	
% Total	14.53%

Actual (Empirical) Performance	
% > USL	10.00%
% < LSL	
% Total	10.00%

Anderson-Darling Goodness-of-Fit Tests	
AD Logistic (2P)	0.245264
AD Logistic (2P) P-Value	> 0.2500
AD Normality Original Data	5.357
AD Normality P-Value Original Data	0.0000

The 2 Parameter Loglogistic distribution was selected as the best fit distribution. For details on how this selection was made, see Appendix: [Statistical Details for Nonnormal Distributions and Transformations](#).

The Anderson Darling statistic for the Loglogistic distribution is 0.245 which is less than the 0.37 value for the AD Normality test of the Box-Cox transformation indicating a better fit. (Note that published AD P-Values for this distribution are limited to a maximum value of 0.25. The best fit selection uses a P-Value estimate that is obtained by transforming the data to normality and then using a modified Anderson Darling Normality test on the transformed data).

Distribution Fitting

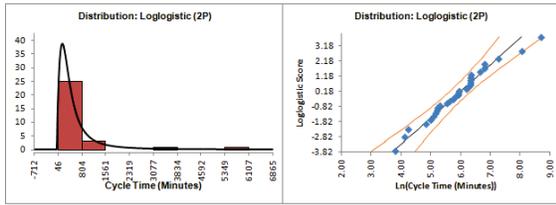
Another helpful tool to evaluate transformations and distributions is **Distribution Fitting**.

1. Click **Sheet 1** Tab (or **F4**). Click **SigmaXL > Process Capability > Nonnormal > Distribution Fitting**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
2. Select *Cycle Time (Minutes)*, click **Numeric Data Variable (Y) >>**. We will use the default selection for **Transformation/Distribution Options: All Transformations & Distributions** as shown:

The screenshot shows the 'Distribution Fitting' dialog box with the following settings:

- Numeric Data Variable (Y) >>**: Cycle Time (Minutes)
- Display Options**:
 - Histograms
 - Probability Plots
 - Add Title
- Confidence Level**: 95.0 %
- Estimate Percentiles (optional - enter percents separated by a space)**: (Empty text box)
- Transformation/Distribution Options**:
 - All Transformations & Distributions
 - Box-Cox Transformation
 - Johnson Transformation
 - Automatic Best Fit
 - Specify Distribution

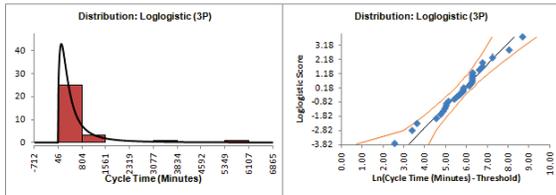
3. Click **OK**. The resulting Distribution Fitting report is shown below. Please note that due to the extensive computations required, this could take up to 1 minute (or longer for large datasets):



Distribution Fitting Report: Cycle Time (Minutes)				
Distribution: Logistic (2P)				
	Sample Count	30		
	Sample Mean	703.17		
Parameter Estimates				
Parameter	Estimate	SE Estimate	Lower 95.0 CI	Upper 95.0 CI
Location	5.895	0.179714	5.543	6.247
Scale	0.571471	0.087949896	0.422862	0.772671

Model Summary and Goodness-of-Fit		
Log-Likelihood	-221.011	
AD Logistic (2P)	0.245264	
AD Logistic (2P) p-value	> 0.2500	

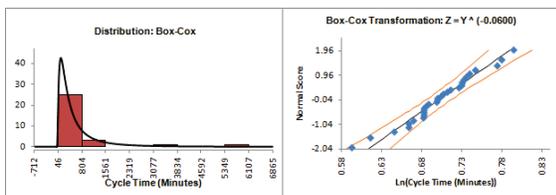
Percentile Report				
Percentage	Cycle Time (Minutes)	SE Percentile	Lower 95.0 CI	Upper 95.0 CI
0.135	9.329	5.067		
50	363.22	65.277	255.39	516.59
99.865	15841	9631.1	4811.0	52156



Distribution Fitting Report: Cycle Time (Minutes)				
Distribution: Logistic (3P)				
	Sample Count	30		
	Sample Mean	703.17		
Parameter Estimates				
Parameter	Estimate	SE Estimate	Lower 95.0 CI	Upper 95.0 CI
Location	5.771	0.232566	5.315	6.227
Scale	0.671824	0.151455	0.431879	1.04507697
Threshold (ML Estimate)	33.025	23.549	-13.129	79.150

Model Summary and Goodness-of-Fit		
Log-Likelihood	-220.631	
AD Logistic (3P)	0.299483	
AD Logistic (3P) p-value	0.562 (Z-Score Est.)	

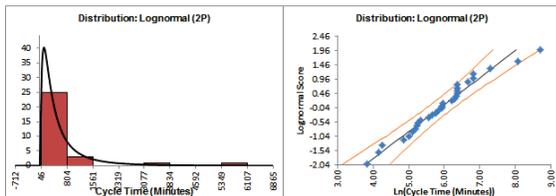
Percentile Report				
Percentage	Cycle Time (Minutes)	SE Percentile	Lower 95.0 CI	Upper 95.0 CI
0.135	36.816	20.605	12.293	110.27
50	353.84	67.889	242.94	515.38
99.865	27183	25811	4227.5	174793



Distribution Fitting Report: Cycle Time (Minutes)				
Box-Cox Transformation: Z = Y^(-0.0600)				
	Sample Count	30		
	Sample Mean	703.17		
Parameter Estimates				
Parameter	Estimate	SE Estimate	Lower 95.0 CI	Upper 95.0 CI
Lambda (Optimal)	-0.06		-0.573963	0.423205
Mean (Transformed Data)	0.702382	0.007937578	0.688925	0.717839
StDev (Transformed Data)	0.044219135	0.00570856	0.034333661	0.056950541

Model Summary and Goodness-of-Fit		
Log-Likelihood	51.498	
AD Box-Cox	0.335933	
AD Box-Cox p-value	0.485000	

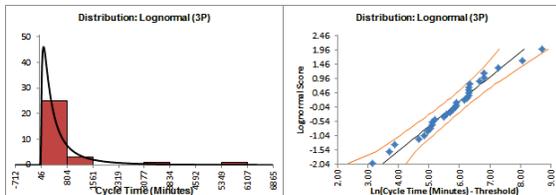
Percentile Report				
Percentage	Cycle Time (Minutes)	SE Percentile	Lower 95.0 CI	Upper 95.0 CI
0.135	11810		36159	4136.7
50	360.67		523.88	250.34
99.865	20.179		42.944	9.797



Distribution Fitting Report: Cycle Time (Minutes)				
Distribution: Lognormal (2P)				
	Sample Count	30		
	Sample Mean	703.17		
Parameter Estimates				
Parameter	Estimate	SE Estimate	Lower 95.0 CI	Upper 95.0 CI
Location	5.920	0.193717	5.541	6.300
Scale	1.061030703	0.140537	0.818434	1.376

Model Summary and Goodness-of-Fit		
Log-Likelihood	-221.456	
AD Lognormal (2P)	0.369554	
AD Lognormal (2P) p-value	0.425000	

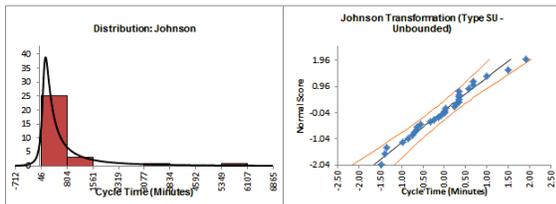
Percentile Report				
Percentage	Cycle Time (Minutes)	SE Percentile	Lower 95.0 CI	Upper 95.0 CI
0.135	15.445	7.168	6.221	28.348
50	372.54	72.167	254.85	544.59
99.865	8985.9	4169.3	3619.2	22310



Distribution Fitting Report: Cycle Time (Minutes)				
Distribution: Lognormal (3P)				
	Sample Count	30		
	Sample Mean	703.17		
Parameter Estimates				
Parameter	Estimate	SE Estimate	Lower 95.0 CI	Upper 95.0 CI
Location	5.805	0.254377	5.307	6.304
Scale	1.160	0.221685	0.797715	1.687
Threshold (ML Estimate)	22.909	23.784	-23.708	69.525

Model Summary and Goodness-of-Fit		
Log-Likelihood	-221.179	
AD Lognormal (3P)	0.385496	
AD Lognormal (3P) p-value	0.370 (Z-Score Est.)	

Percentile Report				
Percentage	Cycle Time (Minutes)	SE Percentile	Lower 95.0 CI	Upper 95.0 CI
0.135	33.134	18.057	11.386	95.418
50	354.92	73.985	235.88	534.03
99.865	10804	6560.8	3285.8	35521



Distribution Fitting Report: Cycle Time (Minutes)				
Johnson Transformation (Type SU - Unbounded)				
Z = Shape1 + Shape2 * Asinh((Y - Location) / Scale)				
	Sample Count	30		
	Sample Mean	703.17		
Parameter Estimates				
Parameter	Estimate	SE Estimate	Lower 95.0 CI	Upper 95.0 CI
Shape1	-1.0451			
Shape2	0.613400			
Location	119.33			
Scale	96.387			
Mean (Transformed Data)	-0.047151462	0.144403	0.330234	0.235931
StDev (Transformed Data)	0.864576	0.103876532	0.624699	1.036246798

Model Summary and Goodness-of-Fit		
Log-Likelihood	-35.536	
AD Johnson	0.286526	
AD Johnson p-value	0.589000	

Percentile Report				
Percentage	Cycle Time (Minutes)	SE Percentile	Lower 95.0 CI	Upper 95.0 CI
0.135	380.45		-1331.1	-28.013
50	355.08		258.88	502.38
99.865	12664		4305.3	37712

The distributions and transformations are sorted in descending order using the AD Normality P-Value on the transformed z-score values. Note that the first distribution shown may not be the selected “best fit”, because the best fit procedure also looks for models that are close but with fewer parameters.

The reported AD P-Values are those derived from the particular distribution. The AD P-Value is not available for distributions with a threshold (except Weibull), so the AD Normality P-Value on the transformed z-score values is used (labeled as Z-Score Est.).

Since the sort order is based on the AD P-Values from Z-Score estimates, it is possible that the reported distribution-based AD P-Values may not be in perfect descending order. However, any discrepancies based on sort order will likely not be statistically or practically significant.

Some data will have distributions and transformations where the parameters cannot be solved (e.g., 2-parameter Weibull with negative values). These are excluded from the Distribution Fitting report.

The parameter estimates and percentile report includes a confidence interval as specified in the **Distribution Fitting** dialog, with 95% being the default. Note that the wide intervals here are due to the small sample size, $n = 30$.

The control limits for the percentile-based Individuals chart will be the 0.135% (lower control limit), 50% (center line, median) and 99.865% (upper control limit). Additional percentiles may be entered in the **Distribution Fitting** dialog.

After reviewing this report, if you wish to perform a process capability analysis with a particular transformation or distribution, simply select **Specify Distribution** from the **Transformation/Distribution Options** in the **Capability Combination Report (Individuals - Nonnormal)** dialog as shown below (using 2 Parameter Loglogistic):

SigmaXL: Measure Phase Tools

Capability Combination Report (Individuals - Nonnormal)

Cycle Time (Minutes)

Numeric Data Variable (Y) >> Cycle Time (Minutes)

Optional X-Axis Labels >>

<< Remove

OK >>

Cancel

Help

Control Chart Options

Individuals - Original Data

Individuals/Moving Range - Normalized Data

Tests for Special Causes

Add Title

Spec Limits for Capability

USL 1000

Target

LSL

Capability Index Options

Z-Score Method

Percentile (ISO) Method

Transformation/Distribution Options

Box-Cox Transformation

Johnson Transformation

Automatic Best Fit

Specify Distribution

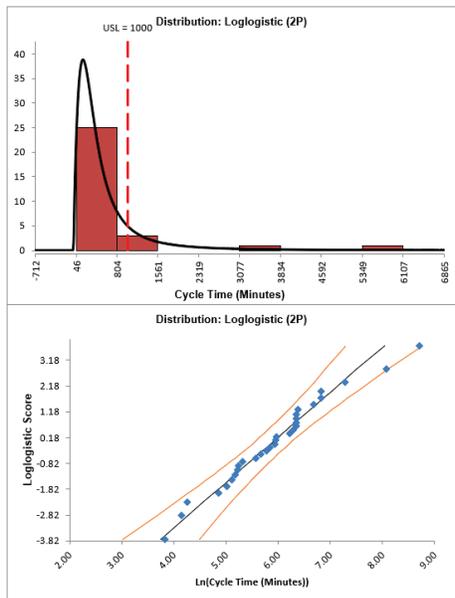
Gamma with Threshold (3 Parameter)

Logistic (2 Parameter)

Loglogistic (2 Parameter)

Selected Distribution

Loglogistic (2 Parameter)

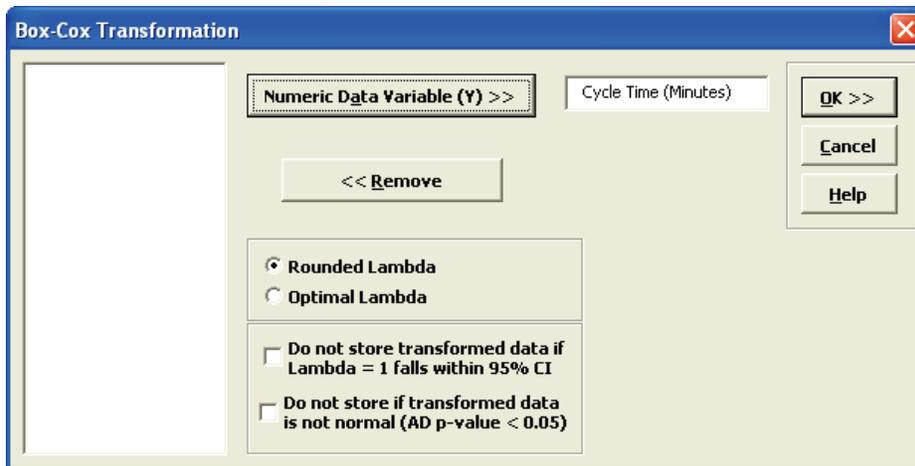


Process Capability Report (Nonnormal): Cycle Time (Minutes)	
Distribution: Loglogistic (2P)	
Sample Count	30
Sample Mean	703.17
Location	5.895
Scale	0.571471221
StDev (Transformed Overall, Long Term)	1
StDev (Normalized Within, Short Term)	0.964388
USL	1000
Target	
LSL	
Capability Indices (Z-Score Method)	
Pp	
Ppu	0.35
Ppl	
Ppk	0.35
Potential Capability Indices using Normalized Within StDev	
Cp	
Cpu	0.37
Cpl	
Cpk	0.37
Expected Overall Performance	
ppm > USL	145274.6
ppm < LSL	
ppm Total	145274.6
% > USL	14.53%
% < LSL	
% Total	14.53%
Actual (Empirical) Performance	
% > USL	10.00%
% < LSL	
% Total	10.00%
Anderson-Darling Goodness-of-Fit Tests	
AD Loglogistic (2P)	0.245264
AD Loglogistic (2P) P-Value	> 0.2500
AD Normality Original Data	5.357
AD Normality P-Value Original Data	0.0000

Box-Cox Transformation

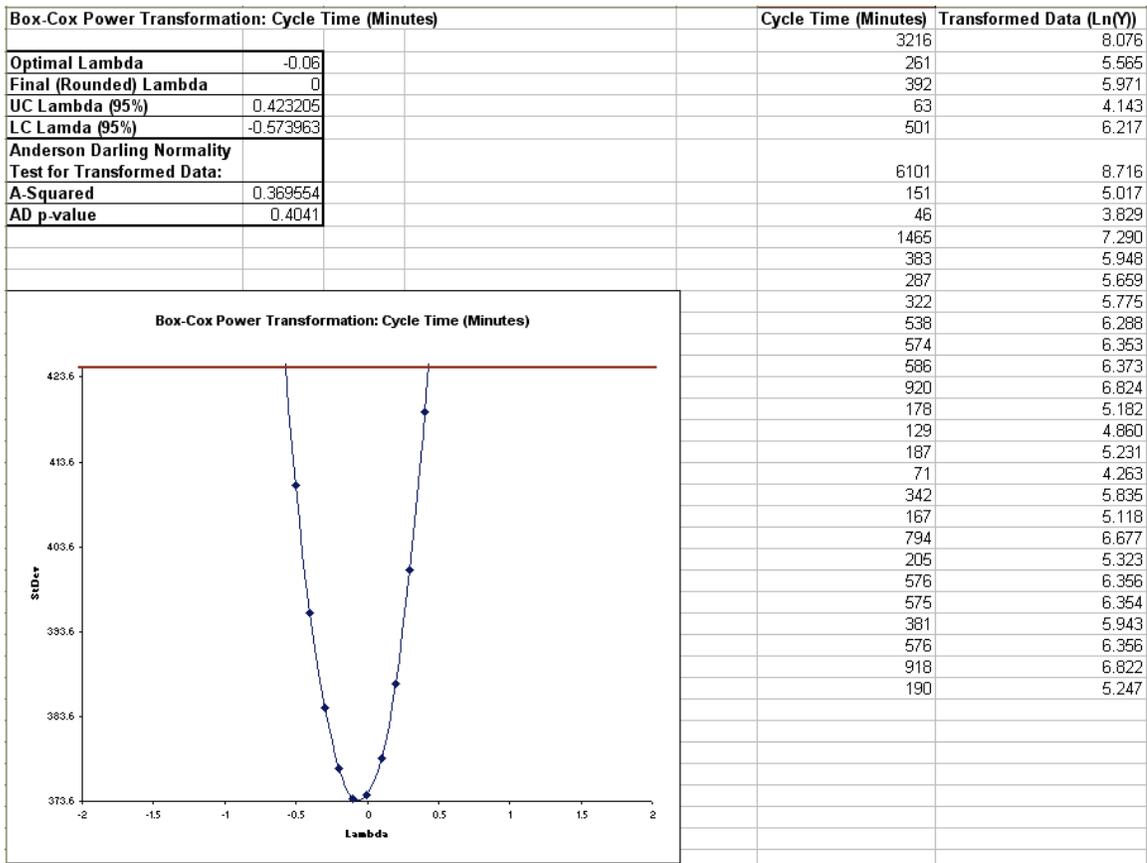
This is a standalone tool that allows you to visually see how the Box-Cox transformation selects a rounded or optimal lambda value.

1. Open the file **Nonnormal Cycle Time2.xlsx**. Select **Sheet 1** Tab.
2. Click **SigmaXL > Process Capability > Nonnormal > Box-Cox Transformation** (or **SigmaXL > Data Manipulation > Box-Cox Transformation** or **SigmaXL > Control Charts > Nonnormal > Box-Cox Transformation**). Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
3. Select *Cycle Time (Minutes)*, click **Numeric Data Variable (Y) >>**. The selected variable must contain all positive values.



Tip: Note that while this tool is often successful to transform the data to normality, there may not be a suitable transformation to make the data normal. The output report indicates the Anderson-Darling P-Value for the transformed data. You may wish to check **Do not store if transformed data is not normal**. Another option is **Do not store transformed data if Lambda = 1 falls within 95%CI**. This latter option prevents you from using transformations that do not result in a statistically significant improvement towards normality.

4. Click **OK**. The resulting report is shown:



5. The fact that Lambda=1 falls outside of the 95% confidence interval tells us that the transformation is statistically significant. The Anderson-Darling P-Value of 0.4041 indicates that we cannot reject the null hypothesis that the transformed data is normal, so the Ln transformation has successfully converted the data to normality.

Part K – Reliability/Weibull Analysis

Reliability/Weibull Analysis

The Weibull distribution is a continuous distribution that was publicized by Waloddi Weibull in 1951. It has become widely used, especially in the reliability field. The Weibull distribution's popularity resulted from its ability to be used with small sample sizes and its flexibility. See Dodson, B. “The Weibull Analysis Handbook” Second Edition, ASQ, for further information. For SigmaXL technical details, see Appendix: [Statistical Details for Nonnormal Distributions and Transformations](#).

The Weibull distribution can also be used as an alternative to the Box-Cox Transformation to determine Process Capability Indices or Control Limits for nonnormal data.

Weibull is particularly effective when the data are censored, with exact failure times unknown or suspension of test. Suspension of test is also referred to as right censored. SigmaXL can analyze complete or right censored data.

1. Open the file **Weibull Dodson.xlsx**. This contains right censored time-to-failure data from Dodson, page 28.
2. Click **SigmaXL > Reliability/Weibull Analysis**. Ensure that the entire data table is selected. Click **Next**.
3. Select *Time-to-Fail*, click **Numeric Response (Y) >>**. Note that the selected variable must contain all positive values. Select *Censor Code*, click **Right Censor Column >>**. The default censor value of 1 will be used. We will use the default method of estimation, **Maximum Likelihood** (some reliability practitioners prefer **Least Squares Regression (X on Y)** for small sample sizes). Other defaults include **Threshold = 0**, and **Confidence Level = 90%**. Enter time values to estimate survival probabilities as shown:

Reliability/Weibull Analysis

Censor Description

Numeric Response (Y) >> Time-to-Fail
(Time to Event)

Right Censor Column >> Censor Code
(Optional: Text or Numeric)

<< Remove

0
1

Censor Value >> 1

Method of Estimation:

Maximum Likelihood
 Least Squares Regression (X on Y)

Threshold: 0.0
Confidence Level: 90.0 %

Estimate Additional Time Percentiles (enter percents separated by a space):

Estimate Survival Probabilities (enter time values separated by a space):

10 20 30 40 50 60 70 80 90 100 150 200 250 300

OK >>
Cancel
Help

4. The resulting Weibull analysis report is shown:

Weibull Analysis: Time-to-Fail

User Settings:

Estimation Method	Maximum Likelihood
Confidence Level	90.0
Threshold	0

Censoring Information:

Number of Uncensored Observations	11
Number of Right Censored Observations	9
Total	20

Model Summary and Goodness-of-Fit:

Log-Likelihood	-66.176
Hollander-Proschan Test	0.029
Hollander-Proschan p-value	0.9771

Parameter Estimates:

Parameter	Estimate	SE Estimate	Lower 90% CI	Upper 90% CI
Shape	2.9744	0.716995	2.0008	4.4218
Scale	203.29	20.851	171.74	240.65

Distribution Characteristics:

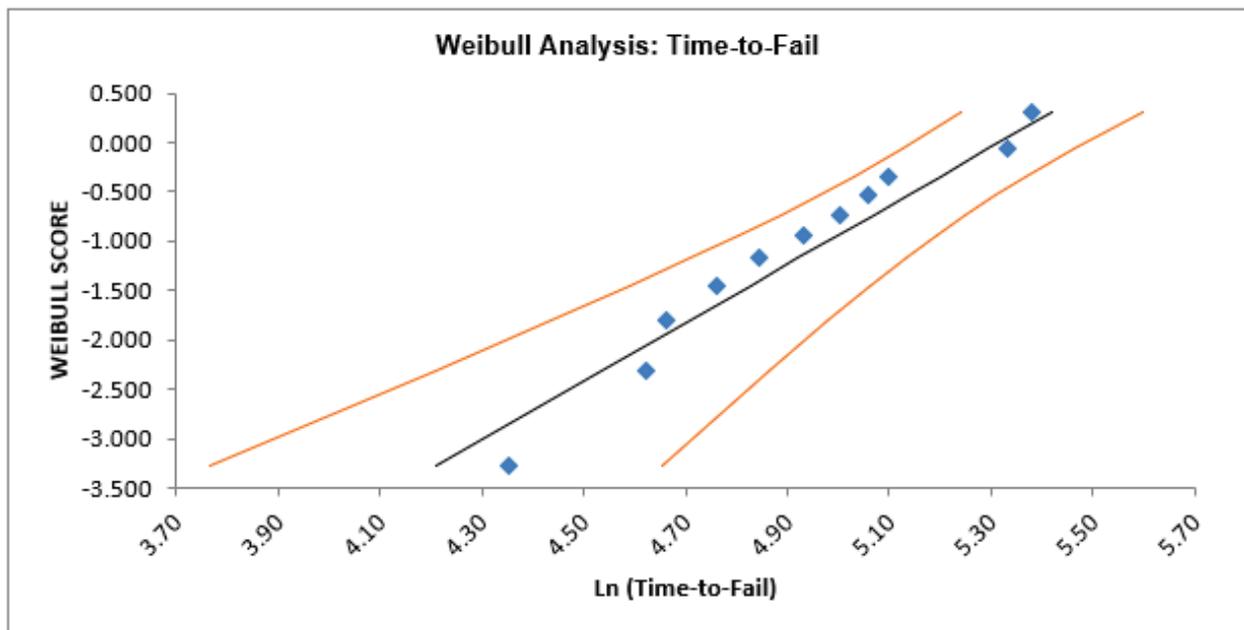
	Estimate	SE Estimate	Lower 90% CI	Upper 90% CI
Mean (MTTF)	181.47	18.420	153.57	214.44
Standard Deviation	66.463	16.172	44.541	99.174

Percentile Report:

Percentage	Percentile (Time)	SE Percentile	Lower 90% CI	Upper 90% CI
0.1	19.935	11.035	8.0201	49.549
0.135	22.052	11.680	9.2279	52.699
0.5	34.268	14.592	17.010	69.037
1	43.298	16.077	23.508	79.747
5	74.895	18.491	49.898	112.41
10	95.401	18.620	69.204	131.52
25	133.73	17.724	107.53	166.30
50	179.73	18.395	151.88	212.68
75	226.89	24.902	189.42	271.78
90	269.09	35.286	216.88	333.87
95	293.99	42.803	231.38	373.54
99	339.71	58.540	255.86	451.03
99.5	356.11	64.682	264.14	480.10
99.865	383.55	75.460	277.51	530.11
99.9	389.32	77.800	280.26	540.83

Survival Probability Report:

Time	Survival Probability	Lower 90.0% CI	Upper 90.0% CI
10	0.999871	0.995704	0.999996
20	0.998990	0.985014	0.999932
30	0.996631	0.968919	0.999639
40	0.992091	0.947956	0.998821
50	0.984698	0.922545	0.997055
60	0.973826	0.893053	0.993800
70	0.958916	0.859814	0.988415
80	0.939499	0.823131	0.980189
90	0.915219	0.783272	0.968381
100	0.885858	0.740474	0.952268
150	0.667092	0.488546	0.795498
200	0.385752	0.208191	0.560906
250	0.157255	0.035498215	0.358758
300	0.04151304	0.00136017	0.215707



5. The Hollander-Proschan test is used when there are right censored observations (see Dodson). If the P-Value is less than 0.05, this indicates significant lack of fit, so Weibull would not be an appropriate distribution. Given the P-Value of 0.9771, there is no lack-of-fit with this time-to-fail data. This is confirmed in the Weibull probability plot with the data roughly following a straight line.
6. If the data is uncensored then the Anderson-Darling unadjusted test is used.
7. The shape parameter, also referred to as Beta, is the slope of the line on the Weibull probability plot. The Weibull distribution can be used in a wide variety of situations and dependent on the value of Beta, is equal to or can approximate several other distributions. For example, if Beta = 1, the Weibull distribution is identical to the exponential distribution; if Beta = 2.5, the Weibull distribution approximates the lognormal distribution; and if Beta = 3.6, the Weibull distribution approximates the normal distribution. Because of this flexibility, there are few observed failure rates that cannot be accurately modeled by the Weibull distribution. Furthermore, if Beta is < 1, the system is suffering from infant mortality failures; if Beta is = 1, failures are random – they are occurring at a constant failure rate and are independent of time; if Beta is > 1, components are wearing out with age.
8. The scale parameter is the characteristic life. This is the point at which we could expect 63.2% of the population under study to have failed. The survival probability report shows that the survival probability at time 200 = 0.386. The cumulative failure probability at 200 is $1 - 0.386 = 0.614$.
9. The percentile report tells us that we can expect a 50% failure rate (one-half of the population will fail) at time = 179.73, i.e. the Median = 179.73.
10. The percentile report can also be used to determine control limits for nonnormal data. The lower control limit is the 0.135 % value = 22.052; the center line is the 50% value = 179.73; the upper control limit is the 99.865 % value = 383.55.

SigmaXL: Analyze Phase Tools

Part A – Stratification with Pareto

SigmaXL's Pareto tool allows you to create Basic (Single) or Advanced (Multiple) Pareto Charts. Advanced Pareto charts are particularly useful in the Analyze Phase because of the ease with which you can slice and dice (or stratify) your data. Of course, Pareto charts are not limited to the Analyze Phase – they can also be used to aid project selection and to prioritize in the Measure Phase.

Consider the following guidelines to help ensure that your Pareto analysis is successful:

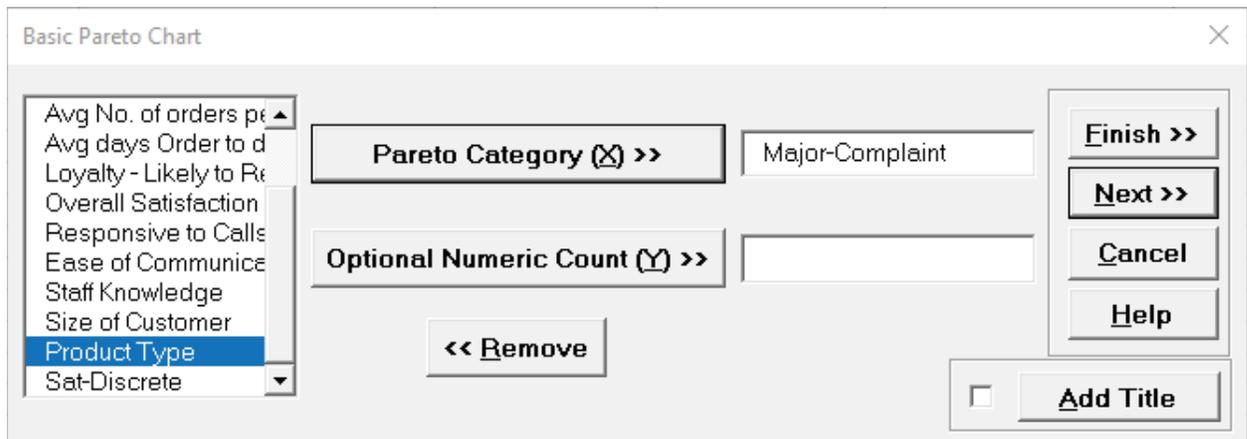
- Your Pareto analysis will only be as good as the quality of the data collected. Ensure that you have the right data and that the data is correct. Use other graphs such as run charts to apply a sanity check to your data.
- Check process stability using appropriate control charts. If the process is not in control, your prioritization of defects and root causes could be invalid.
- Avoid collecting data over too short a time period. Your data may not be representative of the process as a whole. Also keep in mind that since the data is discrete, a minimum sample size of 500 is recommended with 1000 preferred.
- Conversely, data gathered over too long a time period may include process changes that could lead to incorrect conclusions. SigmaXL provides a date subsetting feature that allows you to easily explore different time periods.
- If your initial Pareto analysis does not yield useful results, explore other categories that may be important. SigmaXL's Advanced Charts makes it easy for you to 'slice and dice' your data with different X categories.
- Consider Pareto charting measures such as cost and severity, in addition to defect counts. SigmaXL enables you to chart multiple Y responses.

Basic Pareto Chart Template

Click **SigmaXL > Templates and Calculators > Basic Graphical Templates > Pareto Chart** or **SigmaXL > Graphical Tools > Basic Graphical Templates > Pareto Chart**. See **Measure Phase Part B – Templates and Calculators** for a Pareto Template example.

Basic (Single) Pareto Charts

1. Open the file **Customer Data.xlsx**. Click **SigmaXL>Graphical Tools >Basic Pareto Chart**.
2. Ensure that entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
3. Select *Major-Complaint*, click **Pareto Category (X) >>**.



Tip: SigmaXL will automatically count the number of unique items in the Pareto Category. If we had a separate column with a count (or cost), this count column would be selected as the Optional Numeric Count (Y).

4. Click **Next**. Set **Basic Chart Options** as follows:

Tab “Other” Bar/Cum Sum:
Cum Sum Line – On Top of First Bar



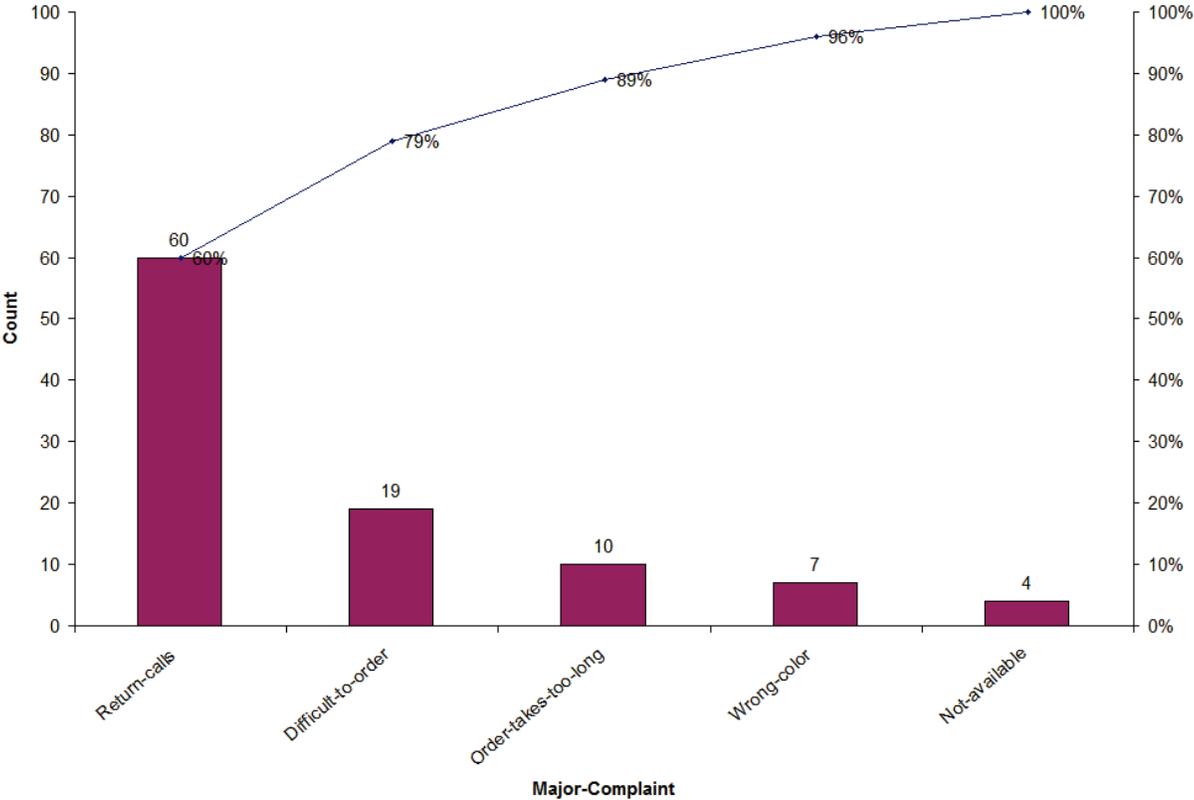
Tab **Chart Options**:
Category (X) Font – Slanted
Data Labels – Check Bars, Cum Sum
Check **Save Defaults**.



Tip: After you have saved your defaults, you can bypass the above options, by clicking **Finish** instead of **Next** at the original Basic Pareto Chart dialog box. The saved defaults will automatically be applied.

Note: For SigmaXL Mac version, click **SigmaXL > Graphical Tools > Basic Pareto Chart Options** to set options. The Basic Pareto Chart is displayed after you click **Finish**.

5. Click **Finish**. The Pareto Chart is produced:

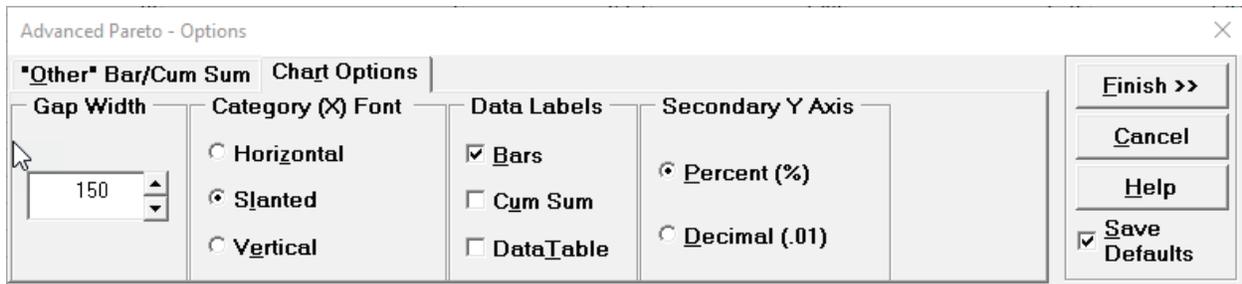


Advanced (Multiple) Pareto Charts

1. Click **Sheet1** Tab of **Customer Data.xlsx** (or press **F4** to activate last worksheet).
2. Click **SigmaXL > Graphical Tools > Advanced Pareto Options**.
Note that the Sample Charts have nothing to do with the data set being evaluated. They are used to dynamically illustrate how your options affect the charts to be produced. (The Sample Charts are not displayed in Excel for Mac).
3. Set **Order of Bars** to **Same Order** on the **"Other Bar"/Cum Sum** options tab.
This is typically used for comparative purposes. The **Descending Order** option makes each Chart a true Pareto Chart, but is less useful for comparison.

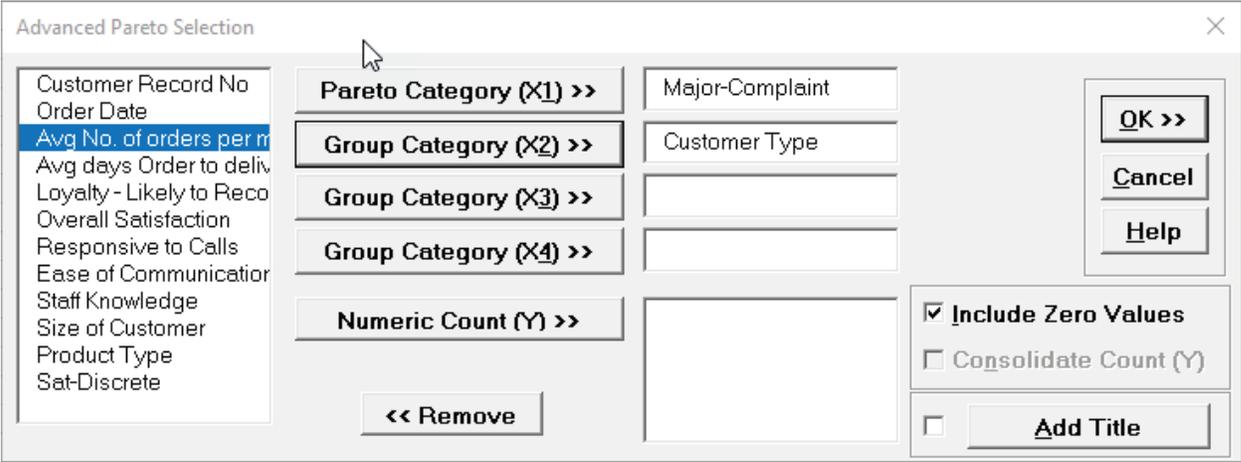


4. Click **Chart Options** tab. Set according to choice – in this case we have selected **Data Labels** for the **Bars** but not for the **Cum Sum** line.

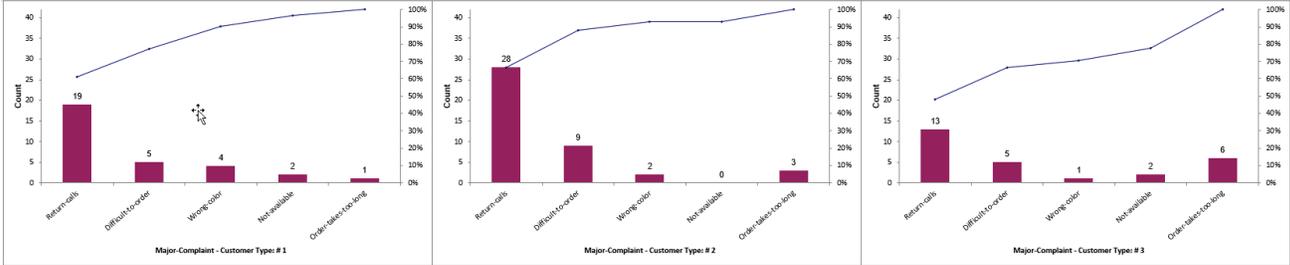


5. Ensure that **Save Defaults** is checked. Note that these options will be saved and applied to all Advanced Pareto Charts. Click **Finish**
6. SigmaXL automatically takes you to the next step of Chart Generation (This is equivalent to clicking **SigmaXL > Graphical Tools > Advanced Pareto Charts**). If necessary, check **Use Entire Data Table**.
7. Click **Next**.

- 8. Select *Major Complaint*, click **Pareto Category (X1) >>**; select *Customer Type*, click **Group Category (X2) >>**.



- 9. Click **OK**. A Pareto Chart of Major Customer Complaints is produced for each Customer Type.



- 10. Click **Sheet 1** Tab, Click **SigmaXL > Graphical Tools > Advanced Pareto Charts**.

- 11. Ensure that entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**. (Steps 10 and 11 can be bypassed with the **Recall SigmaXL Dialog** menu or by pressing **F3** to recall last dialog).

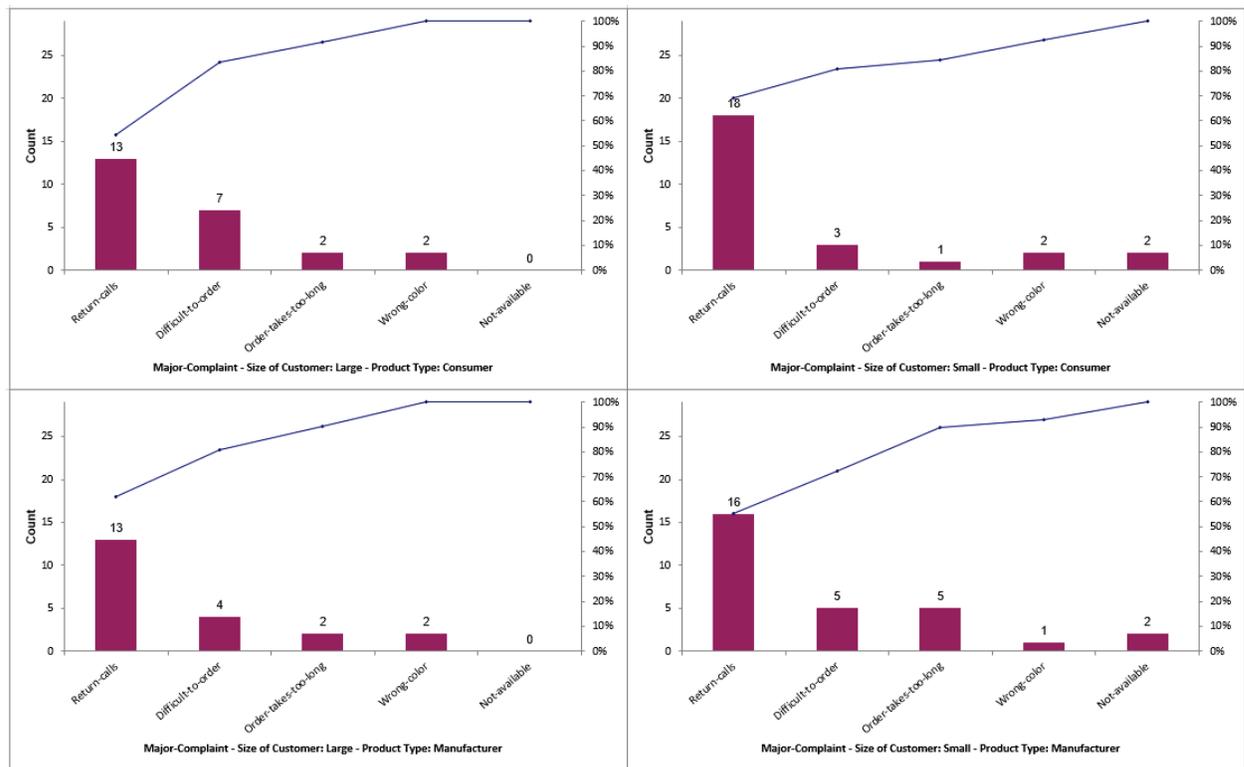
12. Select *Major Complaint*, click **Pareto Category (X1) >>**; select *Size of Customer*, click **Group Category (X2) >>**; select *Product Type*, click **Group Category (X3) >>**.

If a **Numeric Count (Y)** variable is not specified, SigmaXL automatically determines the counts from the **Pareto Category (X1)**.

Normally we would use a text column of discrete Xs, but be aware that numeric columns are also allowed. **Be careful here – this could easily generate a very large number of charts.**

The total number of charts generated = (# of levels in X2) * (# of levels in X3) * (# of levels in X4) * (# of Y variables). Please note, Group Category (X3) and Group Category (X4)

13. Click **OK**. Multiple Pareto's are generated:

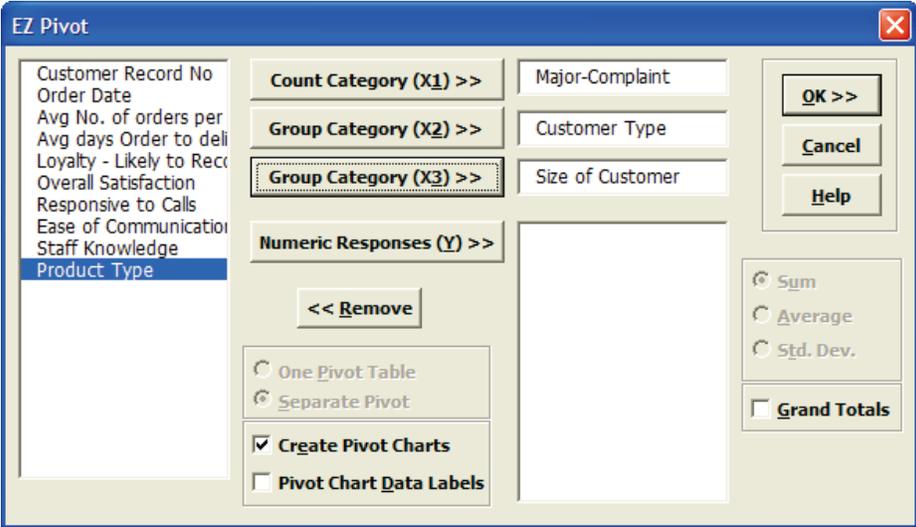


Part B - EZ-Pivot

One of the most powerful features in Excel is the Pivot table. SigmaXL’s EZ-Pivot tool simplifies the creation of Pivot tables using the familiar X and Y dialog box found in the previous Pareto tools.

Example of Three X’s, No Response Y’s

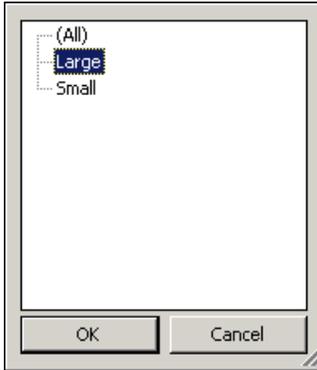
1. Open **Customer Data.xlsx**, click **Sheet 1** (or press **F4** to activate last worksheet). Select **SigmaXL > Graphical Tools > EZ-Pivot/Pivot Charts**.
2. Ensure that entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
3. Select *Major Complaint*, click **Count Category (X1) >>**. Note that if Y is not specified, the Pivot Table Data is based on a count of X1, hence the name **Count Category**.
4. Select *Customer Type*, click **Group Category (X2) >>**; select *Size of Customer*, click **Group Category (X3) >>** as shown.



5. Click **OK**. Resulting Pivot Table of Major Complaint by Customer Type is shown:

Size of Customer	(All)		
Count of Major-Complaint	Customer Type		
Major-Complaint		1	2
Difficult-to-order		5	9
Not-available		2	2
Order-takes-too-long		1	3
Return-calls		19	28
Wrong-color		4	2

- This Pivot table shows the counts for each Major Complaint (X1), broken out by Customer Type (X2), for all Sizes of Customers (X3). (Grand Totals can be added to the Pivot Table by using **Pivot Table Toolbar > Table Options**. Check **Grand Totals for Columns, Grand Totals for Rows**).
- To display counts for a specific Customer Size, click the arrow adjacent to **Size of Customer (All)**. Select **Large**.



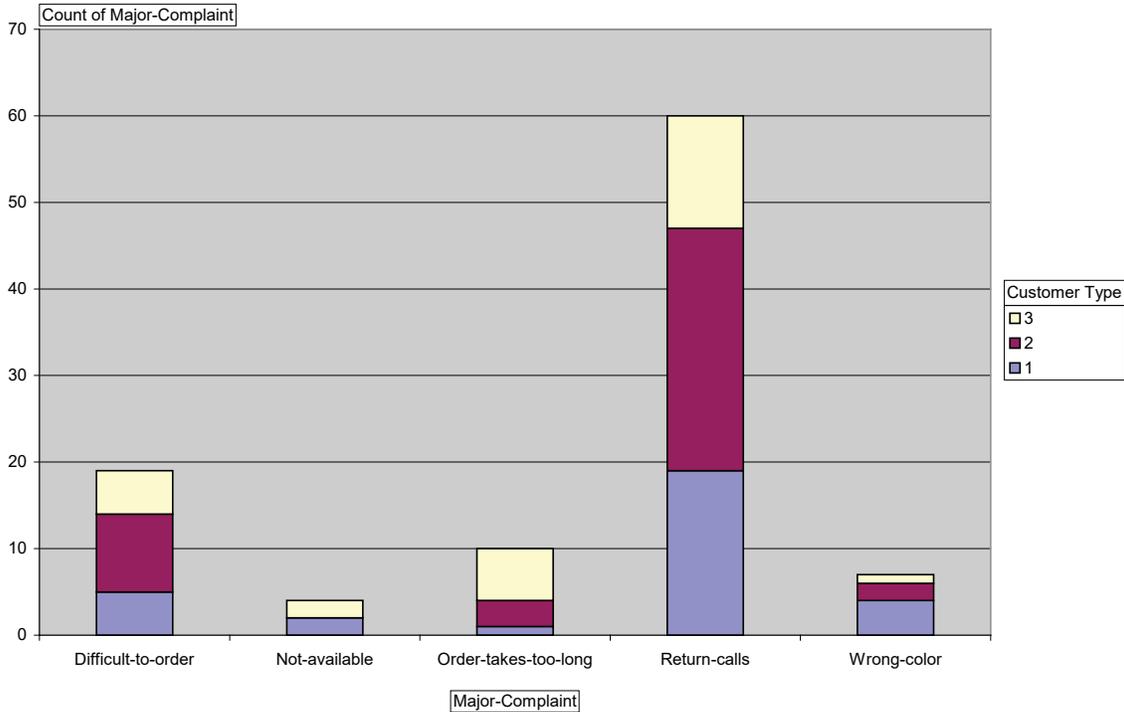
- Click **OK**. Resulting Pivot Table is:

Size of Customer	Large		
Count of Major-Complaint	Customer Type		
Major-Complaint		1	2 3
Difficult-to-order		3	5 3
Order-takes-too-long			2 2
Return-calls		9	14 3
Wrong-color		2	1 1

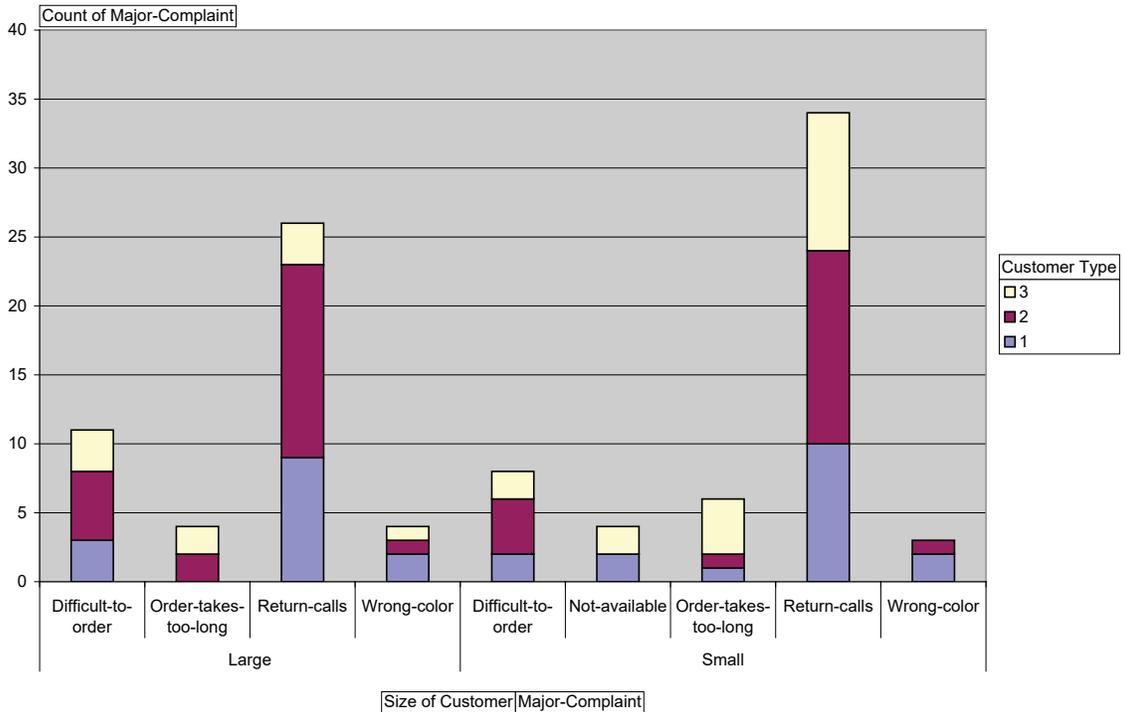
Note that the Major Complaint “Not-Available” is not shown. Pivot tables only show rows where there is at least a count of one.

- The Pivot Chart can be seen by clicking the **EZ Pivot Chart (1)** tab; reset **Size of Customer to All** as shown below:

Size of Customer (All)

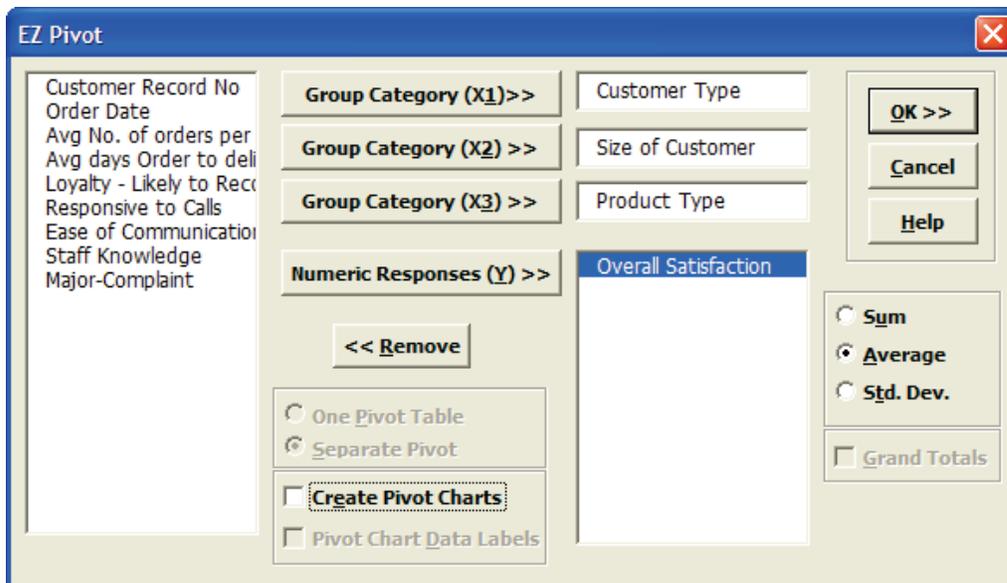


10. Drag **Size of Customer** from **Report Filter** to **Axis Fields** (shown under **PivotChart Fields** to the right of the chart) and Excel will automatically split the Pivot Chart showing both Large and Small Customers.



Example of Three X's and One Y

1. Select **Sheet 1 of Customer Data.xlsx**; click **SigmaXL > Graphical Tools > EZ-Pivot/Pivot Charts**; click **Next** (alternatively, click **Recall SigmaXL Dialog** menu or press **F3** to recall last dialog).
2. Select *Customer Type*, click **Count Category (X1) >>**; select *Size of Customer*, click **Group Category (X2) >>**; select *Product Type*, click **Group Category (X3)**; select *Overall Satisfaction*, click **Numeric Responses (Y) >>**. Note that the Label for **X1** changed from **Count Category** to **Group Category**. The Pivot Table data will now be based on Y data.
3. The Response default uses a Sum of Y. This however can be changed to Average or Standard Deviation. Select **Average**. Uncheck **Create Pivot Charts** (Since we are looking at averages, the stacked bar Pivot Charts would not be very useful, unless they are changed to **clustered column** format using **Chart > Chart Type**).



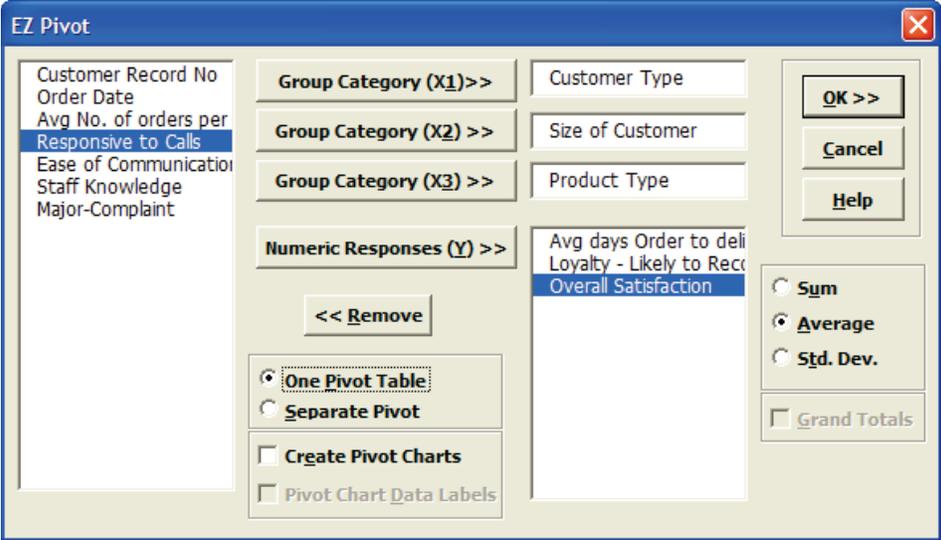
4. Click **OK**. The resulting Pivot Table is:

Product Type	(All)	
Average of Overall Satisfaction	Size of Customer	
Customer Type	Large	Small
1	3.237857143	3.521764706
2	4.309090909	4.091
3	3.56	3.681666667

5. Note that the table now contains Averages of the Customer Satisfaction scores (Y). Again, Product Type (X3) can be varied to show Consumer, Manufacturer, or All. Double clicking on **Average of Overall Satisfaction** allows you to switch to Standard Deviation (StDev).

Example of 3 X's and 3 Y's

1. Click **Recall SigmaXL Dialog** menu or press **F3** to recall last dialog.
2. Select *Customer Type*, click **Group Category (X1) >>**; select *Size of Customer*, click **Group Category (X2) >>**; select *Product Type*, click **Group Category (X3) >>**. Select *Avg Days Order to Delivery, Loyalty – Likely to Respond, Overall Satisfaction*, click **Numeric Responses (Y) >>**. Select **Average** and **One Pivot Table** (default is separate Pivot Tables for each Y). Uncheck **Create Pivot Charts**.



3. Click **OK**. Resulting Pivot Table:

Product Type		(All)		
Customer Type		Data	Size of Customer	
			Large	Small
1	Average of Avg days Order to delivery time		50.85714286	49.76470588
	Average of Loyalty - Likely to Recommend		2.857142857	3.470588235
	Average of Overall Satisfaction		3.237857143	3.521764706
2	Average of Avg days Order to delivery time		48.63636364	49.65
	Average of Loyalty - Likely to Recommend		3.818181818	3.85
	Average of Overall Satisfaction		4.309090909	4.091
3	Average of Avg days Order to delivery time		47	47.66666667
	Average of Loyalty - Likely to Recommend		3.333333333	3.5
	Average of Overall Satisfaction		3.56	3.681666667

Again, Product Type (X3) can be varied.

Part C – Confidence Intervals

One Sample Confidence Interval Templates

Click **SigmaXL > Templates and Calculators > Basic Statistical Templates >**

1 Sample Z-Test and Confidence Interval for Mean or

1 Sample t-Test and Confidence Interval for Mean or

1 Sample Chi-Square Test and CI for Standard Deviation or

1 Proportion Test and Confidence Interval or

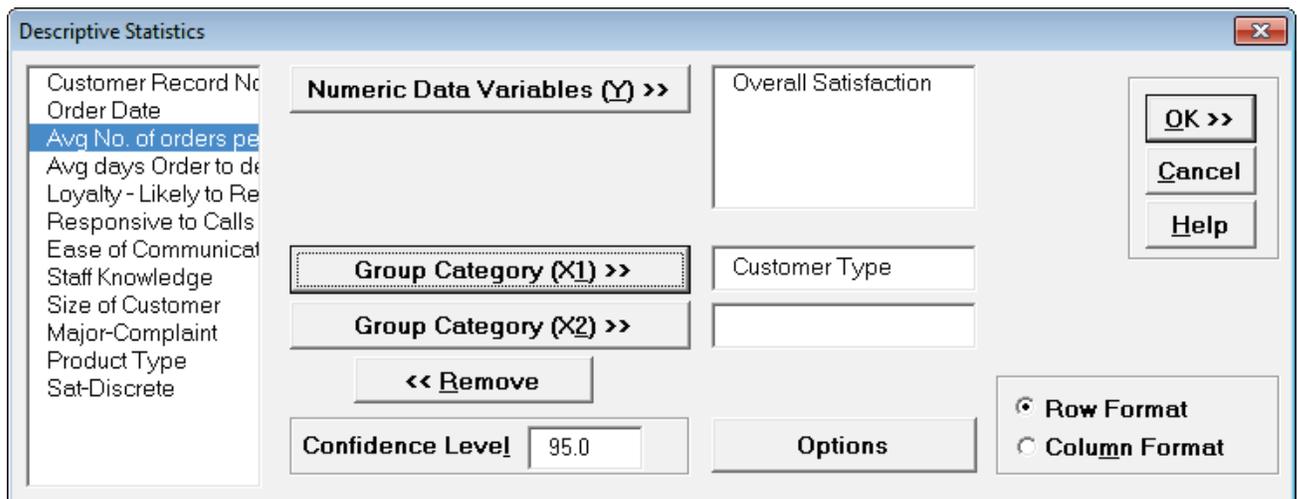
1 Poisson Rate Test and Confidence Interval.

These confidence interval templates are also located at **SigmaXL > Statistical Tools > Basic Statistical Templates.**

See **Measure Phase Part B – Templates and Calculators** for Confidence Interval template examples.

Confidence Intervals – Descriptive Statistics

1. Open **Customer Data.xlsx**. Click **Sheet 1** Tab (or press **F4** to activate last worksheet).
2. Click **SigmaXL > Statistical Tools > Descriptive Statistics.**
3. Check **Use Entire Data Table**, click **Next**.
4. Select **Overall Satisfaction**, click **Numeric Data Variables (Y) >>**. Select **Customer Type**, click **Group Category (X1) >>**; **Confidence Level** default is 95%:



5. Click **OK**. Descriptive Statistics are given for Customer Satisfaction grouped by Customer Type:

Overall Satisfaction by Customer Type			
Descriptive Statistics	Customer Type = 1	Customer Type = 2	Customer Type = 3
Count	31	42	27
Mean	3.394	4.205	3.641
StDev	0.824680	0.621200	0.670478
Range	3.080	2.560	2.740
Minimum	1.720	2.420	2.190
25th Percentile (Q1)	2.810	3.828	3.240
50th Percentile (Median)	3.560	4.340	3.510
75th Percentile (Q3)	4.020	4.725	4.170
Maximum	4.800	4.980	4.930
95.0% CI Mean	3.0911 to 3.696	4.0117 to 4.3988	3.3759 to 3.9063
95.0% CI Sigma	0.65901 to 1.1023	0.51113 to 0.79213	0.52801 to 0.91884
Anderson-Darling Normality Test	0.312776	0.826259	0.389291
P-Value (A-D Test)	0.5306	0.0302	0.3600
Skewness	-0.235169	-0.967994	0.139571
P-Value (Skewness)	0.5557	0.0121	0.7411
Kurtosis	-0.671690	0.679609	-0.313701
P-Value (Kurtosis)	0.3705	0.2865	0.8435

We are given the 95% confidence interval for each sample Mean (95% CI Mean) as well as the 95% confidence interval for the Standard Deviation (95% CI Sigma – do not confuse this with Process Sigma Quality Level).

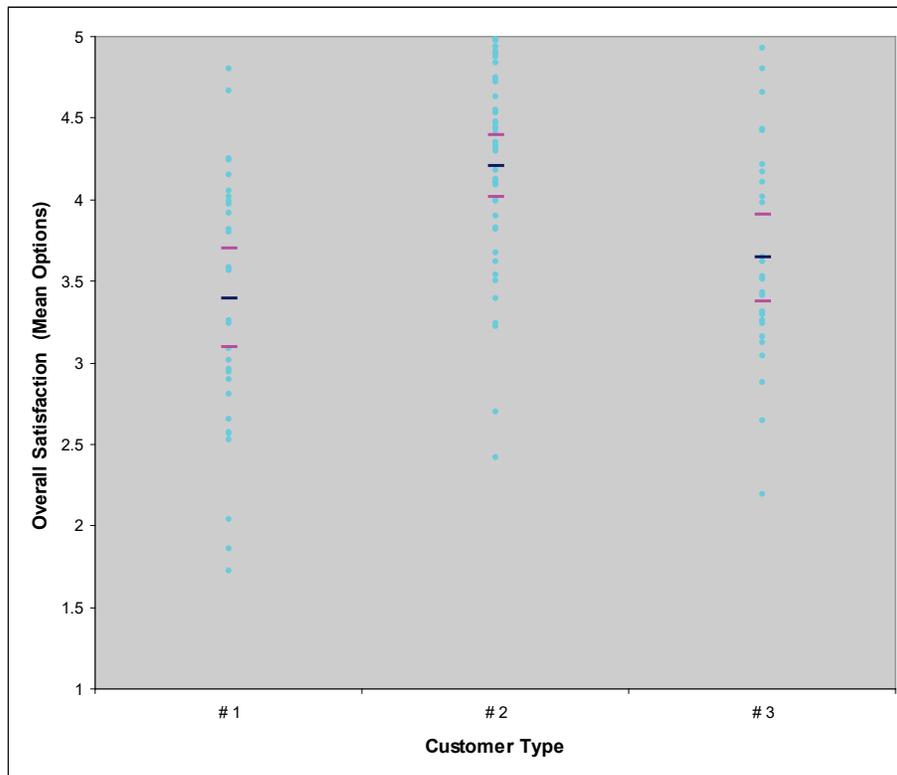
These confidence intervals are very important in understanding our data and making decisions from the data. How often are we driven by sample estimates only and fail to consider the confidence interval or margin of error? For example, newspapers will often fail to take into account the confidence interval when reporting opinion poll results. (To calculate confidence intervals for discrete proportion data, use **SigmaXL > Templates and Calculators > Basic Statistical Templates > 1 Proportion Test and Confidence Interval**).

Note that a confidence interval of 95% implies that, on average, the true population parameter (Mean, Median, Standard Deviation, or Proportion) will lie within the interval 19 times out of 20.

A confidence interval or margin of error does not take into account measurement error or survey bias, so the actual uncertainty may be greater than stated. This should be addressed with good data collection, reliable measurement systems, and good survey design.

Confidence Intervals for the Mean can be obtained in several ways with SigmaXL: Descriptive Statistics, Histograms & Descriptive Statistics, 1-Sample t-test & Confidence Intervals, One-Way ANOVA, and Multi-Vari Charts.

To illustrate confidence intervals for the mean of *Overall Satisfaction* graphically, we have generated a Multi-Vari Chart (with 95% CI Mean Options) using the **Customer Data.xlsx** data. This chart type will be covered later (Part Q).



The dots correspond to individual data points. The tick marks show the 95% upper confidence limit, mean, and 95% lower confidence limit. Clearly we can see that Customer Type 2 has a significantly higher level of mean satisfaction; the lower limit does not overlap with the upper limit for Types 1 and 3. On the other hand, we see overlap of the CI's when comparing types 1 and 3. Hypothesis testing will now be used to compare the mean satisfaction scores more precisely and determine statistical significance for the results.

Part D – Hypothesis Testing – One Sample Z and t-Test

One Sample Z, t and Equivalence Test Templates

Click **SigmaXL > Templates and Calculators > Basic Statistical Templates >**

1 Sample Z-Test and Confidence Interval for Mean or

1 Sample t-Test and Confidence Interval for Mean or

1 Sample Equivalence Test for Mean

These templates are also located at **SigmaXL > Statistical Tools > Basic Statistical Templates.**

See **Measure Phase Part B – Templates and Calculators** for One Sample Z, t and Equivalence Test template examples:

[Basic Statistical Templates – 1 Sample Z-Test and Confidence Interval for Mean](#)

[Basic Statistical Templates – 1 Sample t-Test and Confidence Interval for Mean](#)

[Basic Statistical Templates – 1 Sample Equivalence Test for Mean](#)

One Sample t-Test

1. Open **Customer Data.xlsx**, select **Sheet 1** tab (or press **F4** to activate last worksheet). Click **SigmaXL > Statistical Tools > 1 Sample t-Test & Confidence Intervals**. If necessary, check **Use Entire Data Table**, click **Next**.
2. Ensure that **Stacked Column Format** is selected. Select *Overall Satisfaction*, click **Numeric Data Variable (Y) >>**, select *Customer Type*, click **Optional Group Category (X) >>**.
3. Historically, our average customer satisfaction score has been 3.5. We would like to see if this has changed, with the results grouped by customer type. Null Hypothesis $H_0: \mu=3.5$; Alternative Hypothesis $H_a: \mu \neq 3.5$
4. Enter 3.5 for the **Null Hypothesis H_0 : Mean** value. Set H_a as *Not Equal To*, **Confidence Level** = 95.0%, and check **Display Test Assumptions Report**.

The screenshot shows the '1 Sample t-Test' dialog box. On the left, a list of variables includes 'Avg No. of orders per mo' which is highlighted. The main area has two radio buttons: 'Stacked Column Format (1 Numeric Data Column & 1 Optional Group Category)' is selected, and 'Unstacked Column Format (1 or more Numeric Data Columns)' is unselected. Below these, there are two input fields: 'Numeric Data Variable (Y) >>' containing 'Overall Satisfaction' and 'Optional Group Category (X) >>' containing 'Customer Type'. A '<< Remove' button is positioned below these fields. At the bottom left, there are three input fields: 'H0: Mean =' with the value '3.5', 'Ha:' with a dropdown menu set to 'Not Equal To', and 'Confidence Level:' with the value '95.0'. At the bottom right, there is a checked checkbox for 'Display Test Assumptions Report'. On the far right, there are three buttons: 'OK >>', 'Cancel', and 'Help'.

5. Click **OK**. Results:

1 Sample t-Test - Overall Satisfaction			
Test Information			
H ₀ : Mean (Mu) = 3.5			
H _a : Mean (Mu) Not Equal To 3.5			
Customer Type	1	2	3
Count	31	42	27
Mean	3.394	4.205	3.641
StDev	0.824680	0.621200	0.670478
SE Mean	0.148117	0.095853182	0.129034
t	-0.718700	7.357	1.093599814
P-Value (2-sided)	0.4779	0.0000	0.2842
UC (2-sided, 95%)	3.696	4.399	3.906
LC (2-sided, 95%)	3.091	4.012	3.376

- Note the P-Values. Customer Type 2 shows a significant change (increase) in Satisfaction Mean (P-Value < .05), whereas Customer Types 1 and 3 show no change (P-Value ≥ .05). Also note the confidence intervals around each mean match the results from Descriptive Statistics.
- In the Measure Phase we determined that Overall Satisfaction for Customer Type 2 has nonnormal data but this does not imply that the P-Value for the 1 Sample t-test is wrong. The Central Limit Theorem applies here: the distribution of averages tends to be normal, even if the individual observations are not-normal. With a sample size of 42, the t-test is fairly robust against skewed data.
- 1 Sample t-Test Assumptions Report:

1 Sample t-Test Assumptions Report			
Normality:	Anderson Darling P-Value = 0.531. Fail to reject null hypothesis: "data are sampled from a normal distribution," so conclude that the assumption of normality is not violated.	Anderson Darling P-Value = 0.030. Reject null hypothesis: "data are sampled from a normal distribution," so conclude that the assumption of normality is violated (at 95% confidence level). Skewness value = -0.9680 and Kurtosis value = 0.6796. See robustness and outliers.	Anderson Darling P-Value = 0.360. Fail to reject null hypothesis: "data are sampled from a normal distribution," so conclude that the assumption of normality is not violated.
Robustness:	Not applicable for normal data.	Minimum sample size for a robust t-test = 22. Since the sample size is greater than this, the t-test is robust to nonnormality.	Not applicable for normal data.
Outliers (Boxplot Rules):	No outliers found.	Potential (1.5*IQR) outlier lower count = 1. It is recommended to review the data with graphical tools: Boxplot, Normal Probability Plot, Histogram and Run Chart / Control Chart. Consider using a Nonparametric Test.	No outliers found.
Randomness (Independence):	Nonparametric Runs Test (Exact) P-Value = 0.066. Fail to reject null hypothesis: "data are random," so conclude that the assumption of randomness (independence) is not violated.	Nonparametric Runs Test (Exact) P-Value = 1.000. Fail to reject null hypothesis: "data are random," so conclude that the assumption of randomness (independence) is not violated.	Nonparametric Runs Test (Exact) P-Value = 1.000. Fail to reject null hypothesis: "data are random," so conclude that the assumption of randomness (independence) is not violated.

This is a text report with color highlight: Green (OK), Yellow (Warning) and Red (Serious Violation).

Each sample is tested for Normality using the Anderson Darling test. If not normal, the minimum sample size for robustness of the 1 sample t-Test is determined utilizing Monte Carlo

regression equations (see [Basic Statistical Templates – Minimum Sample Size for Robust t-Tests and ANOVA](#)). If the sample size is inadequate, a warning is given and a suitable Nonparametric Test is recommended (Wilcoxon if symmetric, Sign Test if not symmetric).

Each sample is tested for Outliers defined as: Potential: Tukey's Boxplot ($> Q3 + 1.5 * IQR$ or $< Q1 - 1.5 * IQR$); Likely: Tukey's Boxplot $2.2 * IQR$; Extreme: Tukey's Boxplot $3 * IQR$. If outliers are present, a warning is given and recommendation to review the data with a Boxplot and Normal Probability Plot. Here we have a potential outlier for Customer Type 2.

Tip: If the removal of outlier(s) result in an Anderson Darling P-Value that is > 0.1 , a notice is given that excluding the outlier(s), the sample data are inherently normal.

Each sample is also tested for Randomness using the Exact Nonparametric Runs Test. If the sample data is not random, a warning is given and recommendation to review the data with a Run Chart.

See Appendix [Hypothesis Test Assumptions Report](#) for further details.

Part E – Power and Sample Size

Power and Sample Size – One Sample t-Test – Customer Data

Using the One Sample t-Test (above, Part D), we determined that Customer Types 1 and 3 resulted in “Fail to reject $H_0: \mu=3.5$ ”. A failure to reject H_0 does not mean that we have proven the null to be true. The question that we want to consider here is “What was the power of the test?” Restated, “What was the likelihood that given $H_a: \mu \neq 3.5$ was true, we would have rejected H_0 and accepted H_a ?” To answer this, we will use the Power and Sample Size Calculator.

Tip: Typical sample size rules of thumb address confidence interval size and robustness to normality (e.g. $n=30$). Computing power is more difficult because it involves the magnitude of change in mean to be detected, so one needs to use the power and sample size calculator.

Please use the following guidelines when using the power and sample size calculator:

Power $\geq .99$ (Beta Risk is $\leq .01$) is considered Very High Power

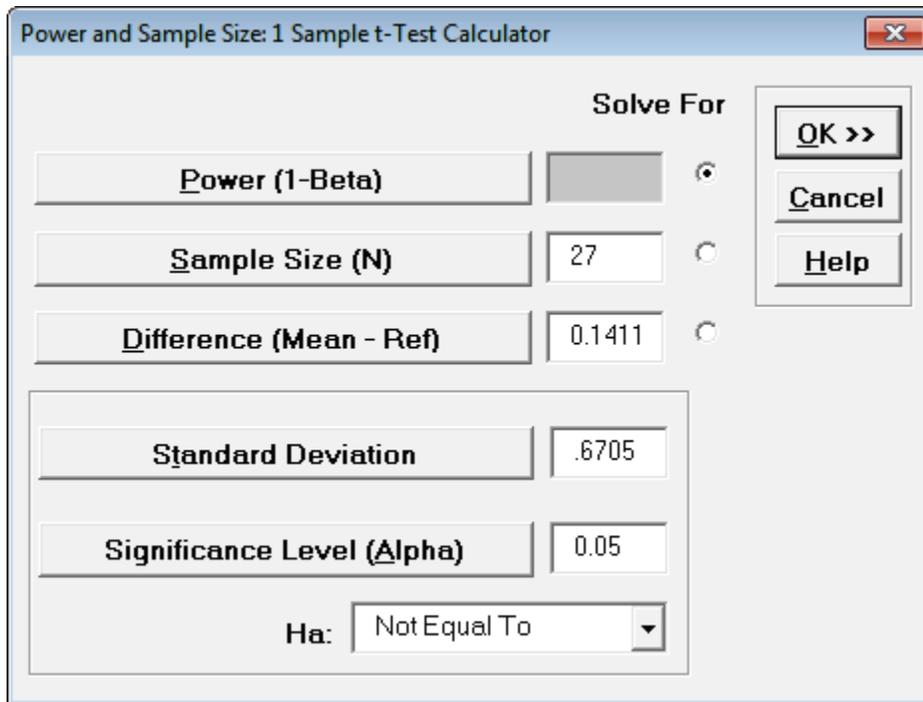
Power $\geq .95$ and $< .99$ (Beta Risk is $\leq .05$) is High Power

Power $\geq .8$ and $< .95$ (Beta Risk is $\leq .2$) is Medium Power. Typically, a power value of .9 to detect a difference of 1 standard deviation is considered adequate for most applications. If the data collection is difficult and/or expensive, then .8 might be used.

Power $\geq .5$ and $< .8$ (Beta Risk is $\leq .5$) is Low Power – not recommended.

Power $< .5$ (Beta Risk is $> .5$) is Very Low Power – do not use!

1. Click **SigmaXL > Statistical Tools > Power and Sample Size Calculators > 1 Sample t-Test Calculator**. We will only consider the statistics from Customer Type 3 here. We will treat the problem as a two-sided test with **$H_a: Not Equal To$** to be consistent with the original test.
2. Enter 27 in **Sample Size (N)**. The difference to be detected in this case would be the difference between the sample mean and the hypothesized value i.e. $3.6411 - 3.5 = 0.1411$. Enter 0.1411 in **Difference**. Leave **Power** value blank, with **Solve For Power** selected (default). Given any two values of Power, Sample size, and Difference, SigmaXL will solve for the remaining selected third value. Enter the sample standard deviation value of 0.6705 in **Standard Deviation**. Keep **Alpha** and **H_a** at the default values as shown:



- Click **OK**. The resulting report is shown:

Power and Sample Size: 1 Sample t Test				
H0: Mean (μ) = Reference				
Ha: Mean (μ) \neq Reference				
Solve For: Power (1 - Beta)				
Sample Size (N)	Difference	Standard Deviation	Significance Level (Alpha)	Power (1 - Beta)
27	0.1411	0.6705	0.05	0.183614

- A power value of 0.1836 is very poor. It is the probability of detecting the specified difference. Alternatively, the associated Beta risk is $1 - 0.1836 = 0.8164$ which is the probability of failure to detect such a difference. Typically, we would like to see Power > 0.9 or Beta < 0.1. In order to detect a difference this small we would need to increase the sample size. We could also set the difference to be detected as a larger value.
- First, we will determine what sample size would be required in order to obtain a Power value of 0.9. Click **Recall SigmaXL Dialog** menu or press **F3** to recall last dialog. Select the **Solve For Sample Size** button as shown. It is not necessary to delete the entered sample size of 27 – it will be ignored. Enter a **Power Value** of .9:

Power and Sample Size: 1 Sample t-Test Calculator

Solve For

Power (1-Beta) 0.9

Sample Size (N) 27

Difference (Mean - Ref) .1411

Standard Deviation .6705

Significance Level (Alpha) 0.05

Ha: Not Equal To

OK >>
Cancel
Help

6. Click **OK**. The resulting report is shown:

Power and Sample Size: 1 Sample t Test					
H0: Mean (μ) = Reference					
Ha: Mean (μ) \neq Reference					
Solve For: Sample Size (N)					
Power (1 - Beta)	Difference	Standard Deviation	Significance Level (Alpha)	Sample Size (N)	Actual Power
0.9	0.1411	0.6705	0.05	240	0.900958197

7. A sample size of 240 would be required to obtain a power value of 0.9. The actual power is rarely the same as the desired power due to the restriction that the sample size must be an integer. The actual power will always be greater than or equal to the desired power.
8. Now we will determine what the difference would have to be to obtain a Power value of 0.9, given the original sample size of 27. Click **Recall SigmaXL Dialog** menu or press **F3** to recall last dialog. Select the **Solve For Difference** button as shown:

9. Click **OK**. The resulting report is shown:

Power and Sample Size: 1 Sample t Test				
H0: Mean (μ) = Reference				
Ha: Mean (μ) \neq Reference				
Solve For: Difference (Mean - Reference)				
Power (1 - Beta)	Sample Size (N)	Standard Deviation	Significance Level (Alpha)	Difference
0.9	27	0.6705	0.05	0.434637

10. A difference of 0.435 would be required to obtain a Power value of 0.9, given the sample size of 27.

Power and Sample Size – One Sample t-Test – Graphing the Relationships between Power, Sample Size, and Difference

In order to provide a graphical view of the relationship between Power, Sample Size, and Difference, SigmaXL provides a tool called **Power and Sample Size with Worksheets**. Similar to the Calculators, **Power and Sample Size with Worksheets** allows you to solve for Power ($1 - \text{Beta}$), Sample Size, or Difference (specify two, solve for the third). You must have a worksheet with Power, Sample Size, or Difference values. Other inputs such as Standard Deviation and Alpha can be included in the worksheet or manually entered.

1. Open the file **Sample Size and Difference Worksheet.xlsx**, select the **Sample Size & Diff** sheet tab. Click **SigmaXL > Statistical Tools > Power & Sample Size with Worksheets > 1 Sample t-Test**. If necessary, check **Use Entire Data Table**. Click **Next**.
2. Ensure that **Solve For Power (1 – Beta)** is selected. Select *Sample Size (N)* and *Difference* columns as shown. Enter the **Standard Deviation** value of 1. Enter .05 as the **Significance Level** value:

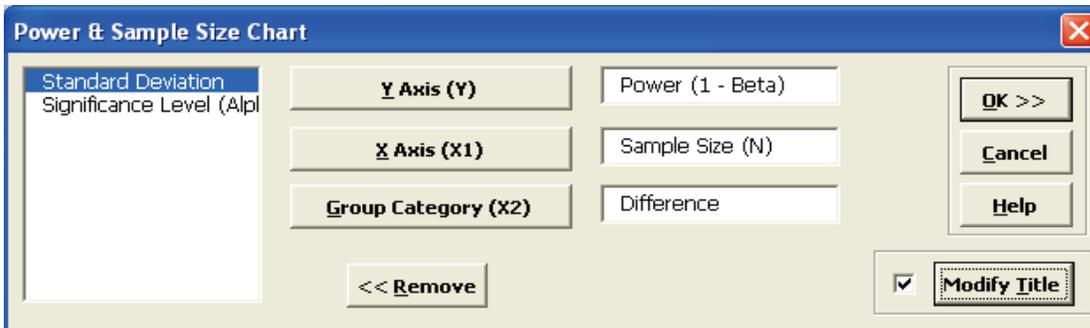
The screenshot shows the 'Power and Sample Size: 1 Sample t-Test' dialog box. The 'Solve For' section has three radio buttons: 'Power (1-Beta) >>' (selected), 'Sample Size (N) >>', and 'Difference (Mean - Ref) >>'. Below this, there are two input fields: 'Standard Deviation >>' with the value '1' and 'Significance Level (Alpha) >>' with the value '.05'. There is also a dropdown menu for 'Ha:' set to 'Not Equal To'. On the right side, there are three buttons: 'OK >>', 'Cancel', and 'Help'. At the bottom, there is a '<< Remove' button.

Note: By setting Standard Deviation to 1, the Difference values will be a multiple of Standard Deviation.

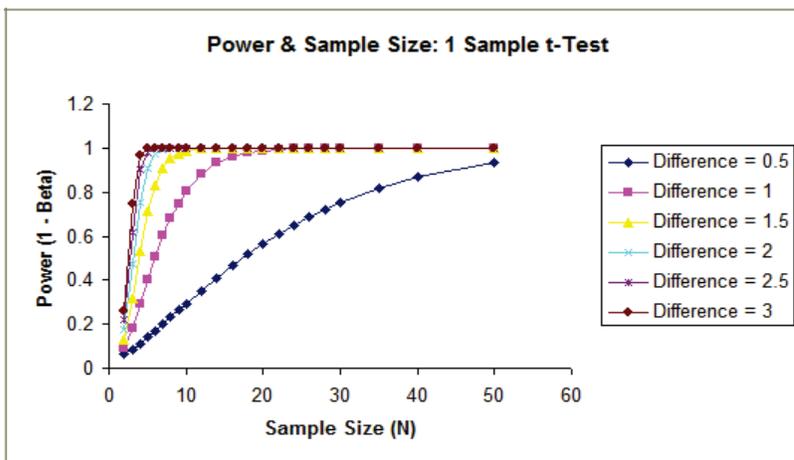
Click **OK**. The output report is shown below:

Power and Sample Size: 1 Sample t Test					
H0: Mean (μ) = Reference					
Ha: Mean (μ) \neq Reference					
Solve For: Power (1 - Beta)					
Sample Size (N)	Difference	Standard Deviation	Significance Level (Alpha)	Power (1 - Beta)	
2	0.5	1	0.05	0.061948607	
3	0.5	1	0.05	0.084107056	
4	0.5	1	0.05	0.111274916	
5	0.5	1	0.05	0.140516692	
6	0.5	1	0.05	0.17070708	
7	0.5	1	0.05	0.201327803	
8	0.5	1	0.05	0.232077006	
9	0.5	1	0.05	0.262746095	
10	0.5	1	0.05	0.293175607	

- To create a graph showing the relationship between Power, Sample Size and Difference, click **SigmaXL > Statistical Tools > Power & Sample Size Chart**. Check **Use Entire Data Table**. Click **Next**.
- Select *Power (1 – Beta)*, click **Y Axis (Y)**; select *Sample Size (N)*, click **X Axis (X1)**; select *Difference*, click **Group Category (X2)**. Click **Add Title**. Enter *Power & Sample Size: 1 Sample t-Test*:



- Click **OK**. Click **OK**. The resulting Power & Sample Size Chart is displayed:



Since the Standard Deviation was specified as 1, the Difference values are multiples of the Standard Deviation.

Part F – One Sample Nonparametric Tests

Introduction to Nonparametric Tests

Nonparametric tests make fewer assumptions about the distribution of the data compared to parametric tests like the t-Test. Nonparametric tests do not rely on the estimation of parameters such as the mean or the standard deviation. They are sometimes called distribution-free tests.

Nonparametric tests use Medians and Ranks, thus they are robust to outliers in the data. If, however, the data are normal and free of outliers, nonparametric tests are less powerful than normal based tests to detect a real difference when one exists.

Nonparametric tests should be used when the data are nonnormal, the data cannot be readily transformed to normality, and the sample size is too small for robustness due to the Central Limit Theorem.

Exact Nonparametric Tests

Nonparametric tests do not assume that the sample data are normally distributed, but they do assume that the test statistic follows a Normal or Chi-Square distribution when computing the “large sample” or “asymptotic” P-Value. The One-Sample Sign Test, Wilcoxon Signed Rank, Two Sample Mann-Whitney and Runs Test assume a Normal approximation for the test statistic. Kruskal-Wallis and Mood’s Median use Chi-Square to compute the P-Value. With very small sample sizes (see recommendations below), these approximations may be invalid, so exact methods should be used. SigmaXL computes the exact P-Values utilizing permutations and fast network algorithms.

It is important to note that while exact P-Values are “correct,” they do not increase (or decrease) the power of a small sample test, so they are not a solution to the problem of failure to detect a change due to inadequate sample size.

For data that require more computation time than specified, Monte Carlo P-Values provide an approximate (but unbiased) P-Value that typically matches exact to two decimal places using 10,000 replications. One million replications give a P-Value that is typically accurate to three decimal places. A confidence interval (99% default) is given for the Monte Carlo P-Values. Note that the Monte Carlo confidence interval for P-Value is **not** the same as a confidence interval on the test statistic due to data sampling error. The confidence level for the hypothesis test statistic is still 95%, so **all reported P-Values less than .05 will be highlighted in red** to indicate significance. The 99% Monte Carlo P-Value confidence interval is due to the uncertainty in Monte Carlo sampling, and it becomes smaller as the number of replications increases (irrespective of the data

sample size). The Exact P-Value will lie within the stated Monte Carlo confidence interval 99% of the time.

Recommended Sample Sizes for use of Exact Nonparametric Tests

Sign Test: $N \leq 50$

Wilcoxon Signed Rank: $N \leq 15$

Mann-Whitney: Each sample $N \leq 10$

Kruskal-Wallis: Each sample $N \leq 5$

Mood's Median: Each sample $N \leq 10$

Runs Test (Above/Below) or Runs Test (Up/Down) Test: $N \leq 50$

These are sample size guidelines for when exact nonparametric tests should be used rather than “large sample” asymptotic based on the Normal or Chi-Square approximation. It is always acceptable to use an exact test, but computation time can become an issue especially for tests with two or more samples. In those cases, one can always use a Monte Carlo P-Value with 99% confidence interval.

Validation of SigmaXL Exact P-Values

SigmaXL Exact P-Values are validated by comparison to textbook examples, published paper examples and other exact software such as StatXact, SPSS Exact, SAS, and Matlab.

Monte Carlo P-Values are validated using 1e6 replications and compared against exact. Repeated simulations are used to validate the confidence intervals.

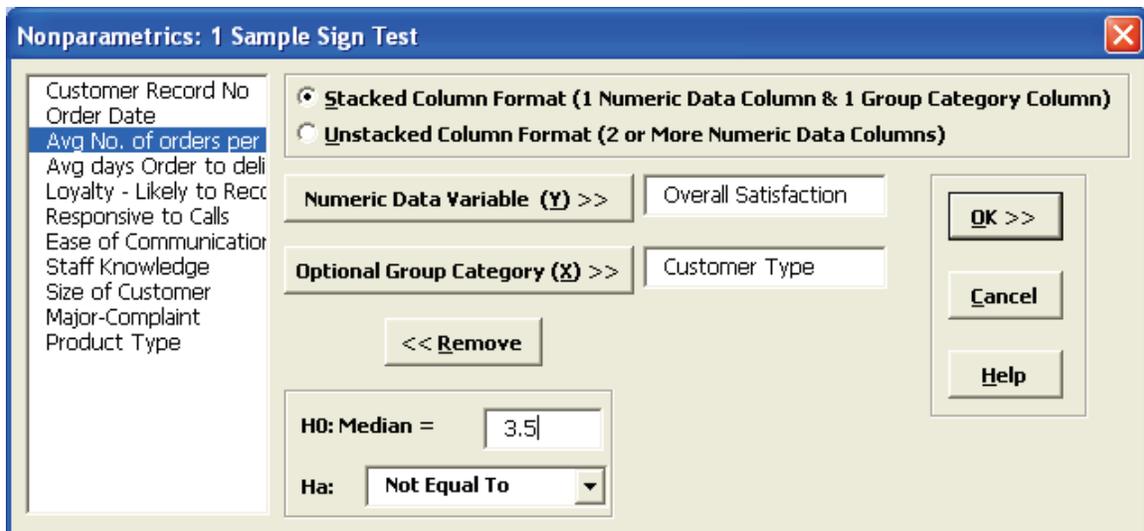
For further details and references refer to the Appendix [Exact and Monte Carlo P-Values for Nonparametric and Contingency Tests.](#)

One Sample Sign Test

The sign test is the simplest of the nonparametric tests, and is similar to testing if a two-sided coin is fair. Count the number of positive values (larger than hypothesized median), the number of negative values (smaller than the hypothesized median), and test whether there are significantly more positives (or negatives) than expected. The One Sample Sign Test is a nonparametric equivalent to the parametric One Sample t-Test.

Historically, our Median customer satisfaction score has been 3.5. We would like to see if this has changed, with the results grouped by customer type (H_0 : Median=3.5, H_a : Median \neq 3.5, $\alpha = 0.05$).

1. Open **Customer Data.xlsx**, select **Sheet 1** tab. Click **SigmaXL > Statistical Tools > Nonparametric Tests>1 Sample Sign**. If necessary, check **Use Entire Data Table**, click **Next**.
2. Ensure that **Stacked Column Format** is selected. Select *Overall Satisfaction*, click **Numeric Data Variable (Y) >>**; select *Customer Type*, click **Optional Group Category (X) >>**.
3. Enter 3.5 for the **Null Hypothesis H_0 : Median** value. Set **H_a** as *Not Equal To*.



4. Click **OK**. Results:

1 Sample Sign Test for Medians: Overall Satisfaction			
Test Information			
H ₀ : Median = 3.5			
H _a : Median Not Equal To 3.5			
Customer Type			
	1	2	3
Count (N)	31	42	27
Median	3.560	4.340	3.510
Points Below 3.5	15	5	13
Points Equal To 3.5	0	1	0
Points Above 3.5	16	36	14
P-Value (2-sided)	1.0000	0.0000	1.0000

Note the P-Values. Customer Type 2 shows a significant change (increase) in Satisfaction Median (P-Value < .05), whereas Customer Types 1 and 3 show no change (P-Value ≥ .05). While the P-Values are not the same as those given by the 1 sample t-Test, the conclusions do match.

If Count (N) is less than or equal to 50, the Sign Test computes an exact P-Value using the binomial distribution. For N > 50, the P-Value is estimated using a normal approximation. Since this is always done automatically and is very fast, the Sign Test is not included in the separate Nonparametric Exact menu.

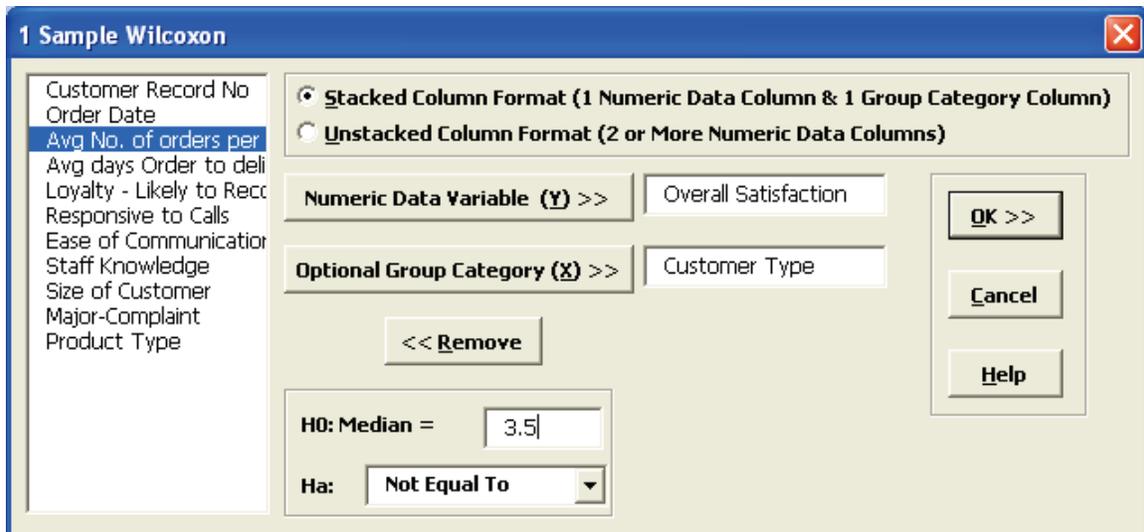
One Sample Wilcoxon Signed Rank Test

The Wilcoxon Signed Rank test is a more powerful nonparametric test than the Sign Test, but it adds an assumption that the distribution of values is symmetric around the median. An example of a symmetric distribution is the uniform distribution. Symmetry can be observed with a histogram, or by checking to see if the Skewness is large (> .5 or < - .5).

The One Sample Wilcoxon Test is a nonparametric equivalent to the parametric One Sample t-Test.

Historically, our Median customer satisfaction score has been 3.5. We would like to see if this has changed, with the results grouped by customer type (H₀: Median=3.5, H_a: Median≠3.5, α = 0.05).

1. Open **Customer Data.xlsx**, select Sheet 1 tab (or press F4 to activate last worksheet). Click **SigmaXL > Statistical Tools > Nonparametric Tests > 1 Sample Wilcoxon**. If necessary, check Use Entire Data Table, click Next.
2. Ensure that **Stacked Column Format** is selected. Select *Overall Satisfaction*, click **Numeric Data Variable (Y) >>**; select *Customer Type*, click **Optional Group Category (X) >>**.
3. Enter 3.5 for the **Null Hypothesis H₀: Median** value. Keep H_a as *Not Equal To*.



4. Click **OK**. Results:

1 Sample Wilcoxon Test: Overall Satisfaction			
Test Information			
H ₀ : Median = 3.5			
H _a : Median Not Equal To 3.5			
Customer Type			
	1	2	3
Count (N)	31	42	27
Count for Test	31	41	27
Median	3.56	4.34	3.51
Wilcoxon Statistic	217.50	802.50	222.00
P-Value (2-sided)	0.5566	0.0000	0.4349

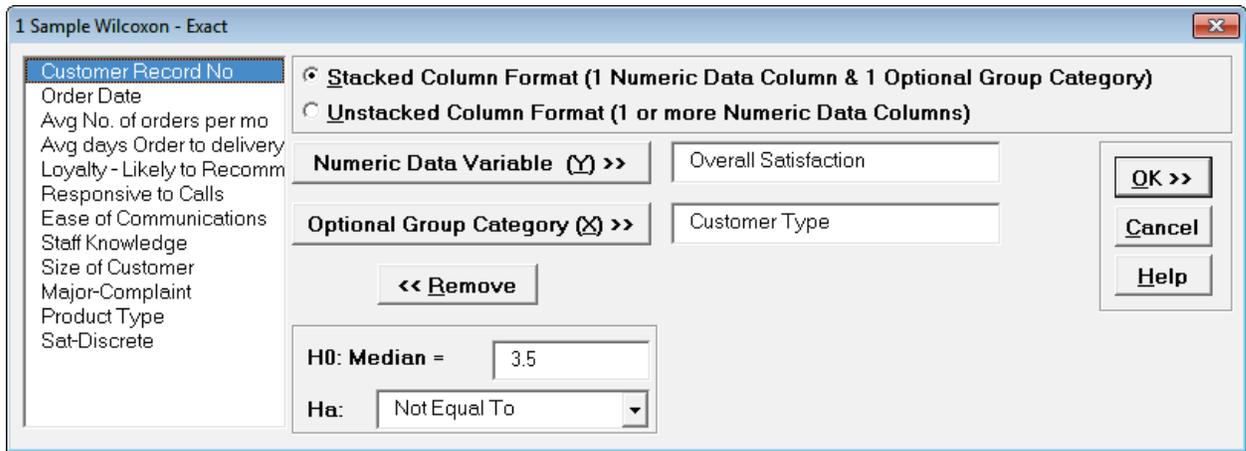
Note the P-Values. Customer Type 2 shows a significant change (increase) in Satisfaction Median (P-Value < .05), whereas Customer Types 1 and 3 show no change (P-Value ≥ .05). Although the P-Values are not identical to the sign test and t-Test, the conclusions match. (Note, in the case of Customer Type 2, the Sign Test is preferred since the data is not symmetrical but skewed).

One Sample Wilcoxon Signed Rank Test - Exact

We will now redo the above example to compute exact P-Values. Typically, this would not be necessary unless the sample sizes were smaller ($N \leq 15$ for Wilcoxon), but this gives us continuity on the example. We will consider a small sample problem later.

1. Open **Customer Data.xlsx**, select Sheet 1 tab (or press F4 to activate last worksheet). Click **SigmaXL > Statistical Tools > Nonparametric Tests – Exact > 1 Sample Wilcoxon - Exact**. If necessary, check **Use Entire Data Table**, click **Next**.

2. Ensure that **Stacked Column Format** is selected. Select *Overall Satisfaction*, click **Numeric Data Variable (Y) >>**; select *Customer Type*, click **Optional Group Category (X) >>**.
3. Enter 3.5 for the **Null Hypothesis H0: Median** value. Keep **Ha** as *Not Equal To*.



4. Click **OK**. Results:

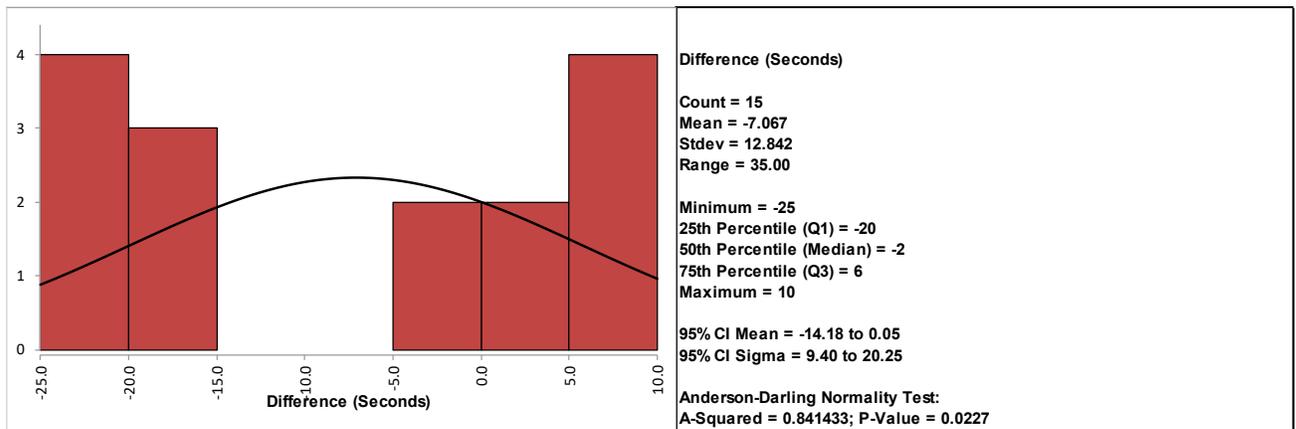
1 Sample Wilcoxon Test - Exact: Overall Satisfaction			
Test Information			
H ₀ : Median = 3.5			
H _a : Median Not Equal To 3.5			
Customer Type	1	2	3
Count (N)	31	42	27
Count for Test	31	41	27
Median	3.56	4.34	3.51
Wilcoxon Statistic	217.50	802.50	222.00
Exact P-Value (2-sided)	0.5584	0.0000	0.4410

The Wilcoxon Statistics are identical to the above “large sample” or “asymptotic” results. The Exact P-Values are close but slightly different. This was expected because the sample sizes are reasonable (N > 15), so the “large sample” P-Values are valid using a normal approximation for the Wilcoxon Statistic.

Note, if Count (N) is greater than 1000, the Exact P-Value is estimated using a continuity - corrected normal approximation. Since the Wilcoxon Exact P-Value is computed very quickly for sample sizes as large as 1000, Monte Carlo P-Values are not required.

5. Now we will consider a small sample problem. Open **Nonnormal Task Time Difference – Small Sample.xlsx**. A study was performed to determine the effectiveness of training to reduce the time required to complete a short but repetitive process task. Fifteen operators were randomly selected and the difference in task time was recorded in seconds (after training – before training). A negative value denotes that the operator completed the task in less time after training than before.

6. Click **SigmaXL > Graphical Tools > Histograms & Descriptive Statistics**. If necessary, check **Use Entire Data Table**. Click **Next**. Select *Difference (Seconds)*, click **Numeric Data Variables (Y) >>**. Click **OK**.



This small sample data fails the Anderson Darling Normality Test (P-Value = .023). Note that this is due to the data being uniform or possibly bimodal, not due to a skewed distribution. Now we will perform a 1 Sample t-Test and review the assumptions.

7. Select **Task Time Difference** tab (or press **F4** to activate last worksheet). Click **SigmaXL > Statistical Tools > 1 Sample t-Test & Confidence Intervals**. If necessary, check **Use Entire Data Table**, click **Next**.
8. Ensure that **Stacked Column Format** is selected. Select *Difference (Seconds)*, click **Numeric Data Variable (Y) >>**.
9. This is a one-sided test because we have no reason to expect that the task time will increase, so the Null Hypothesis $H_0: \mu = 0$; and the Alternative Hypothesis $H_a: \mu < 0$. Enter 0 for the **H0: Mean** value. Select **H_a** as *Less Than*, **Confidence Level** = 95.0%, and check **Display Test Assumptions Report**.

The screenshot shows the '1 Sample t-Test' dialog box with the following settings:

- Stacked Column Format (1 Numeric Data Column & 1 Optional Group Category)**
- Unstacked Column Format (1 or more Numeric Data Columns)**
- Numeric Data Variable (Y) >>**: Difference (Seconds)
- Optional Group Category (X) >>**: (Empty)
- H0: Mean =**: 0
- Ha:** Less Than
- Confidence Level:** 95.0
- Display Test Assumptions Report**

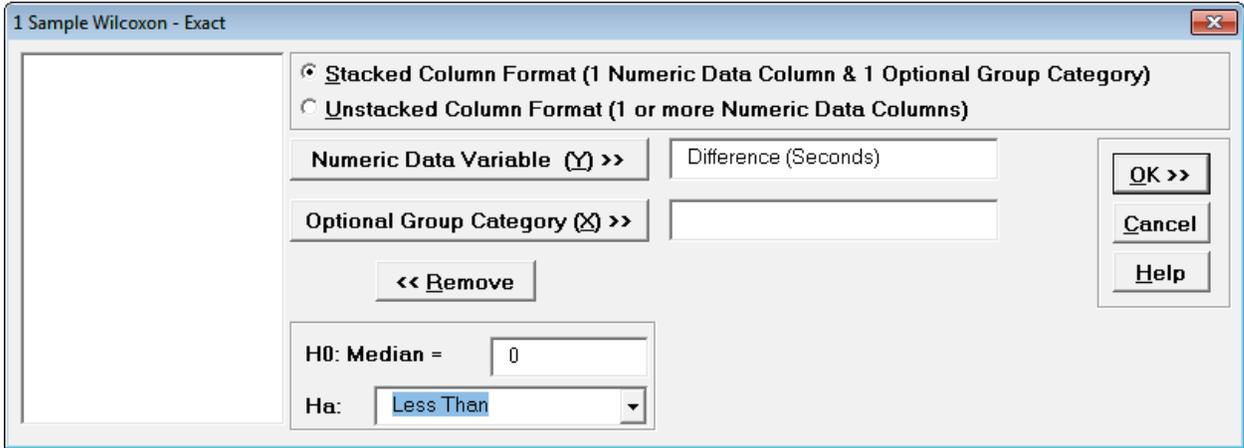
10. Click OK.

1 Sample t-Test	
Test Information	
H ₀ : Mean (Mu) = 0	
H _a : Mean (Mu) Less Than 0	
Results: Difference (Seconds)	
Count	15
Mean	-7.067
StDev	12.842
SE Mean	3.316
t	-2.131
P-Value (1-sided)	0.0256
UC (1-sided, 95%)	-1.226
1 Sample t-Test Assumptions Report	
Normality:	Anderson Darling P-Value = 0.023. Reject null hypothesis: "data are sampled from a normal distribution," so conclude that the assumption of normality is violated (at 95% confidence level). Skewness value = -0.1369 and Kurtosis value = -1.8385. See robustness and outliers.
Robustness:	Minimum sample size for a robust t-test = 16. It is recommended that the Exact Nonparametric One Sample Wilcoxon Test be used (SigmaXL > Statistical Tools > Nonparametric Tests - Exact > 1 Sample Wilcoxon - Exact).
Outliers (Boxplot Rules):	No outliers found.
Randomness (Independence):	Nonparametric Runs Test (Exact) P-Value = 1.000. Fail to reject null hypothesis: "data are random," so conclude that the assumption of randomness (independence) is not violated.

The **1 Sample t-Test Assumptions Report** highlights that the data are not normal, but note that Kurtosis equal to -1.84 is the issue here, not Skewness. This was observed in the Histogram above with the data being uniform or possibly bimodal.

The sample size is too small for a robust t-Test, so the Exact One Sample Wilcoxon Test is recommended. The Wilcoxon Test is recommended over the Sign Test because it is a more powerful test and meets the requirement that the data be symmetrical. The Exact test is recommended because the sample size is very small (N <= 15).

11. Select **Task Time Difference** tab (or press **F4** to activate last worksheet). Click **SigmaXL > Statistical Tools > Nonparametric Tests – Exact > 1 Sample Wilcoxon - Exact**. If necessary, check **Use Entire Data Table**, click **Next**.
12. Ensure that **Stacked Column Format** is selected. Select *Difference (Seconds)*, click **Numeric Data Variable (Y) >>**. Enter 0 for the **H0: Median** value. Select **Ha** as *Less Than*.



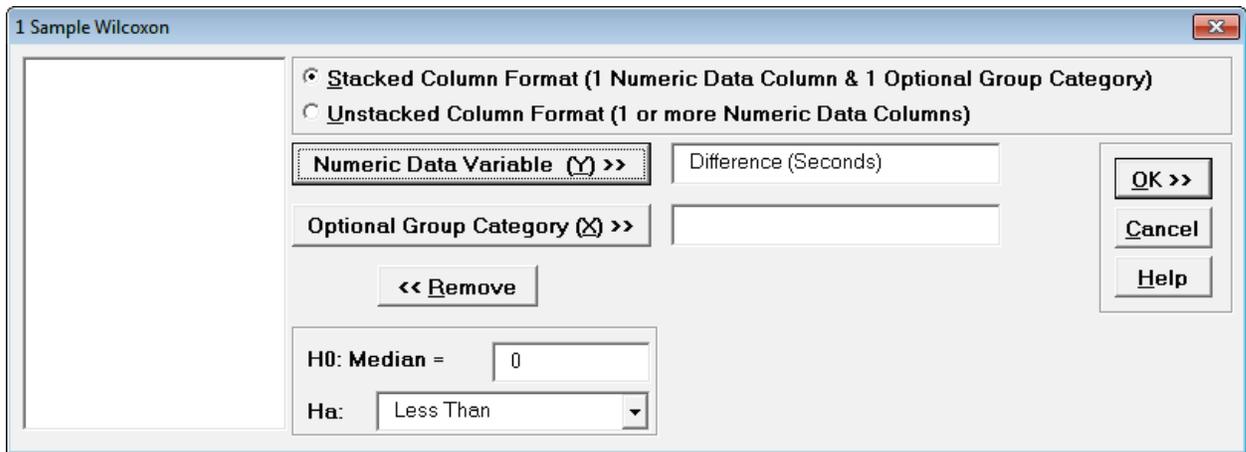
13. Click **OK**. Results:

1 Sample Wilcoxon Test - Exact	
Test Information	
H ₀ : Median = 0	
H _a : Median Less Than 0	
Results: Difference (Seconds)	
Count (N)	15
Count for Test	14
Median	-2
Wilcoxon Statistic	26.00
Exact P-Value (1-sided)	0.0497

With the P-Value = .0497 we reject H₀ and conclude that the Median Task Time Difference is significantly less than 0, so the training is effective.

By way of comparison we will now rerun the analysis using the large sample (asymptotic) Wilcoxon test.

14. Select **Task Time Difference** tab (or press **F4** to activate last worksheet). Click **SigmaXL > Statistical Tools > Nonparametric Tests > 1 Sample Wilcoxon**. If necessary, check **Use Entire Data Table**, click **Next**.
15. Ensure that **Stacked Column Format** is selected. Select *Difference (Seconds)*, click **Numeric Data Variable (Y) >>**. Enter 0 for the **H0: Median** value. Select **Ha** as *Less Than*.



16. Click **OK**. Results:

1 Sample Wilcoxon Test	
Test Information	
H ₀ : Median = 0	
H _a : Median Less Than 0	
Results: Difference (Seconds)	
Count (N)	15
Count for Test	14
Median	-2
Wilcoxon Statistic	26.00
P-Value (1-sided)	0.0513

Now with the P-Value = .0513 we **incorrectly fail to reject H₀**.

The difference between exact and large sample P-Value is small but it was enough to lead us to falsely conclude that the training is ineffective.

In conclusion, whenever you have a small sample size and are performing a nonparametric test, always use the Exact option.

Part G – Two Sample t-Test

Two Sample t and Equivalence Test Templates

Click **SigmaXL > Templates and Calculators > Basic Statistical Templates > 2 Sample t-Test and Confidence Interval (Compare 2 Means)** or **2 Sample Equivalence Test (Compare 2 Means)**.

These templates are also located at **SigmaXL > Statistical Tools > Basic Statistical Templates**.

See **Measure Phase Part B – Templates and Calculators** for Two Sample t-Test template examples:

[Basic Statistical Templates – 2 Sample t-Test and Confidence Interval \(Compare 2 Means\)](#)

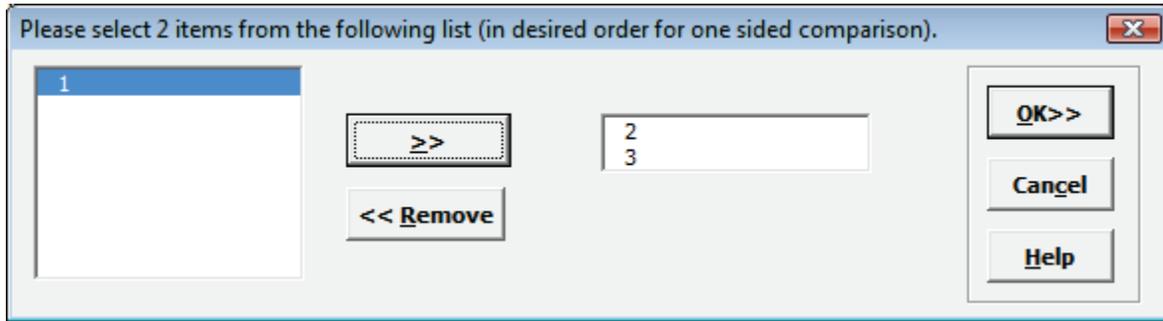
[Basic Statistical Templates – 2 Sample Equivalence Test \(Compare 2 Means\)](#)

Two Sample t-Test

1. Open **Customer Data.xlsx**, click **Sheet 1** tab (or press **F4** to activate last worksheet). We will look at comparing means of Customer Satisfaction by Customer Type (2 vs. 3), using the Two Sample t-test. $H_0: \mu_2 = \mu_3$, $H_a: \mu_2 \neq \mu_3$
2. Click **SigmaXL > Statistical Tools > 2 Sample t-test**. If necessary, check **Use Entire Data Table**, click **Next**.
3. With stacked column format checked, select *Overall Satisfaction*, click **Numeric Data Variable Y >>**; select *Customer Type*, click **Group Category X >>**; **H0: Mean Diff = 0**; **Ha: Not Equal To**; **Confidence Level: 95%**; ensure that **Assume Equal Variances** and **Display Test Assumption Report** are checked:

The screenshot shows the '2 Sample t-Test' dialog box in SigmaXL. On the left, a list of variables is shown, with 'Avg No. of orders per mo' selected. The main area has two radio buttons: 'Stacked Column Format (1 Numeric Data Column & 1 Group Category Column)' is selected, and 'Unstacked Column Format (2 Numeric Data Columns)' is unselected. Below these, there are two input fields: 'Numeric Data Variable (Y) >>' containing 'Overall Satisfaction' and 'Group Category (X) >>' containing 'Customer Type'. A '<< Remove' button is below the group category field. Further down, 'H0: Mean Diff =' is set to '0', 'Ha:' is set to 'Not Equal To', and 'Confidence Level:' is set to '95.0'. At the bottom, there are two checked checkboxes: 'Assume Equal Variances' and 'Display Test Assumptions Report'. On the right side, there are three buttons: 'OK >>', 'Cancel', and 'Help'.

4. Click **OK**. Select *Customer Type 2* and 3.



5. Click **OK**. Resulting output:

2 Sample t-Test: Overall Satisfaction		
Test Information		
H_0 : Mean Difference = 0		
H_a : Mean Difference Not Equal To 0		
Assume Equal Variance		
Customer Type	2	3
Count	42	27
Mean	4.205	3.641
Standard Deviation	0.621200	0.670478
Mean Difference	0.564127	
Std Error Difference	0.158060	
DF	67	
t	3.569	
P-Value (2-sided)	0.0007	
UC (2-sided, 95%)	0.879616	
LC (2-sided, 95%)	0.248638	

Given the P-Value of .0007 we reject H_0 and conclude that Mean Customer Satisfaction is significantly different between Customer type 2 and 3. This confirms previous findings.

6. 2 Sample t-Test Assumptions Report:

2 Sample t-Test Assumptions Report		
Normality:	Anderson Darling P-Value = 0.030. Reject null hypothesis: "data are sampled from a normal distribution," so conclude that the assumption of normality is violated (at 95% confidence level). Skewness value = -0.9680 and Kurtosis value = 0.6796. See robustness and outliers.	Anderson Darling P-Value = 0.360. Fail to reject null hypothesis: "data are sampled from a normal distribution," so conclude that the assumption of normality is not violated.
Robustness:	Minimum sample size for a robust two sample t-test = 2. Since each sample size is greater than this, the two sample t-test is robust to nonnormality.	
Outliers (Boxplot Rules):	Potential (1.5*IQR) outlier lower count = 1. It is recommended to review the data with graphical tools: Boxplot, Normal Probability Plot, Histogram and Run Chart / Control Chart. Consider using a Nonparametric Test.	No outliers found.
Randomness (Independence):	Nonparametric Runs Test (Exact) P-Value = 1.000. Fail to reject null hypothesis: "data are random," so conclude that the assumption of randomness (serial independence) is not violated.	(Exact) P-Value = 1.000. Fail to reject null hypothesis: "data are random," so conclude that the assumption of randomness (serial independence) is not violated.
Equal Variance:	Levene's test for Equal Variances P-Value = 0.622. Fail to reject null hypothesis: "variances are equal," so conclude that the assumption of equal variances (or standard deviations) is not violated.	

This is a text report with color highlight: Green (OK), Yellow (Warning) and Red (Serious Violation).

Each sample is tested for Normality using the Anderson-Darling test. If not normal, the minimum sample size for robustness of the 2 sample t-Test is determined utilizing Monte Carlo regression equations (see [Basic Statistical Templates – Minimum Sample Size for Robust t-Tests and ANOVA](#)). If the sample size is inadequate, a warning is given and the Nonparametric Mann-Whitney Test is recommended.

Each sample is tested for Outliers defined as: Potential: Tukey's Boxplot ($> Q3 + 1.5 \cdot IQR$ or $< Q1 - 1.5 \cdot IQR$); Likely: Tukey's Boxplot $2.2 \cdot IQR$; Extreme: Tukey's Boxplot $3 \cdot IQR$. If outliers are present, a warning is given and recommendation to review the data with a Boxplot and Normal Probability Plot. Here we have a potential outlier for Customer Type 2.

Tip: If the removal of outlier(s) result in an Anderson Darling P-Value that is > 0.1 , a notice is given that excluding the outlier(s), the sample data are inherently normal.

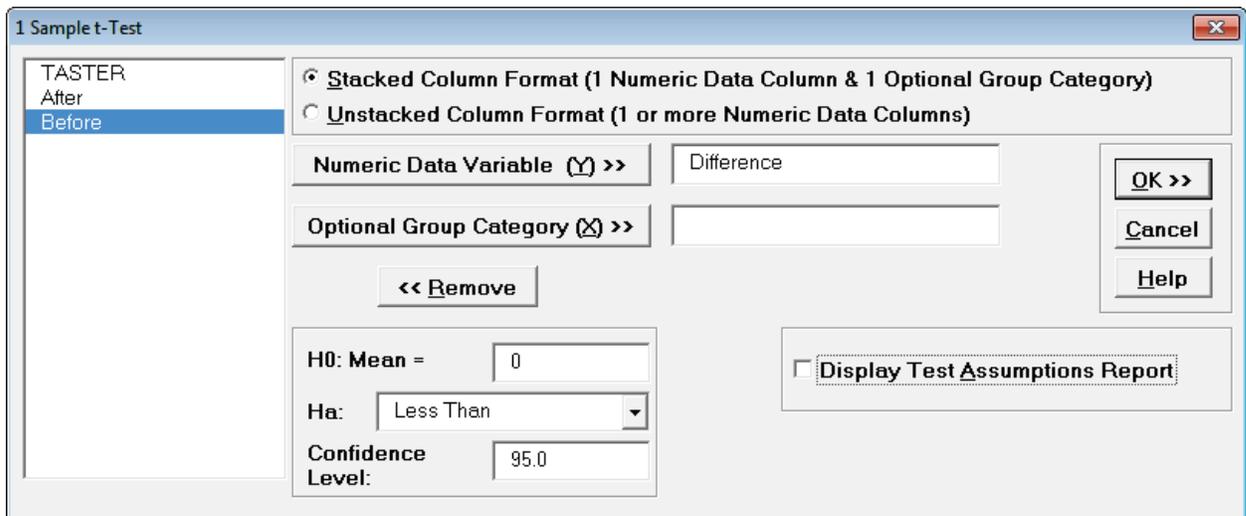
Each sample is tested for Randomness using the Exact Nonparametric Runs Test. If the sample data is not random, a warning is given and recommendation to review the data with a Run Chart.

A test for Equal Variances is also applied. If all sample data are normal, the F-Test is utilized, otherwise Levene's Test is used. If the variances are unequal and the test being used is the equal variance option, then a warning is given and Welch's test is recommended.

See Appendix [Hypothesis Test Assumptions Report](#) for further details.

Paired t-Test

1. Open the file **Dietcola.xlsx**. These are the results of a Before and After taste test on sweetness for diet cola. Ten tasters were used and one month elapsed with the cola in warm storage between the before and after results. Do a one sample t-test on the column of differences.
2. Click **SigmaXL > Statistical Tools > 1-Sample t-test & Confidence Intervals**. Ensure that entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**. Select *Difference*, click **Numeric Data Variable (Y) >>**, set **H0: Mean = 0**, **Ha: Less Than** (this is a one-sided or one-tail test – sweetness cannot increase). Uncheck **Display Test Assumptions Report** for now:



3. Click **OK**. Result:

1 Sample t-Test	
Test Information	
H ₀ : Mean (Mu) = 0	
H _a : Mean (Mu) Less Than 0	
Results:	
Count	Difference 10
Mean	-1.02
StDev	1.196
SE Mean	0.378242
t	-2.697
P-Value (1-sided)	0.0123
UC (1-sided, 95%)	-0.326641

Given the P-Value of .012, we reject H₀ and conclude that the sweetness has in fact decreased.

4. Now redo the analysis using the paired t-test: Click **Sheet 1 Tab**; Click **SigmaXL > Statistical Tools > Paired t-Test**; Click **Next**; Select *After*, click **Numeric Data Variable 1 >>**, select *Before* as

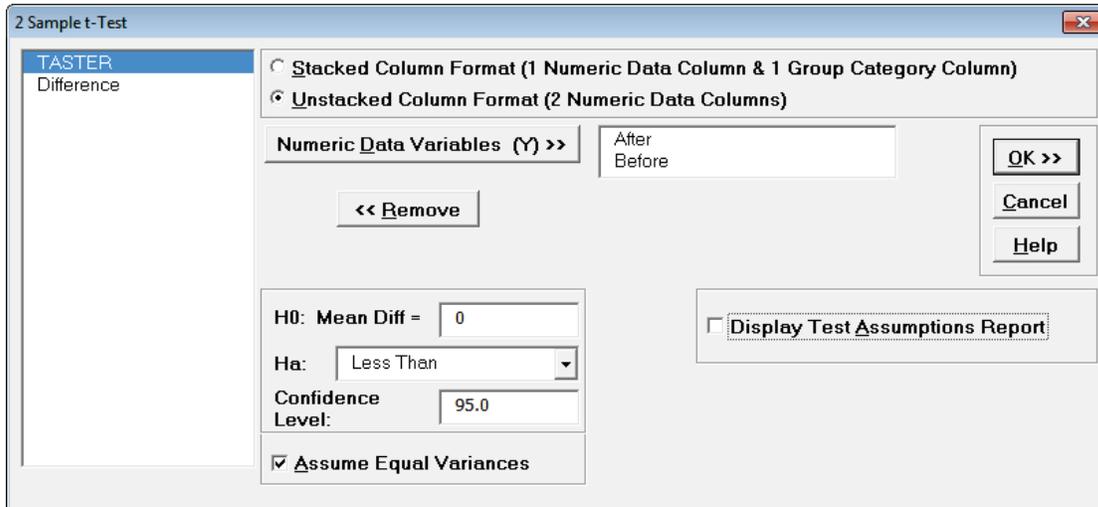
Numeric Data Variable 2 >>, H₀: Mean Diff = 0, H_a: Less Than. Check Display Test Assumptions Report.

5. Click **OK**. Results are identical to the One sample t-test analysis of difference column, with the added assumptions report showing that all assumptions are met:

Paired t-Test	
Test Information	
H ₀ : Mean Difference = 0	
H _a : Mean Difference Less Than 0	
Results:	
	After - Before
Count	10
Mean	-1.02
StDev	1.196
SE Mean	0.378242
t	-2.697
P-Value (1-sided)	0.0123
UC (1-sided, 95%)	-0.326641
Paired t-Test Assumptions Report	
Normality:	Anderson Darling P-Value = 0.369. Fail to reject null hypothesis: "data are sampled from a normal distribution," so conclude that the assumption of normality is not violated.
Robustness:	Not applicable for normal data.
Outliers (Boxplot Rules):	No outliers found.
Randomness (Independence):	Nonparametric Runs Test (Exact) P-Value = 0.079. Fail to reject null hypothesis: "data are random," so conclude that the assumption of randomness (independence) is not violated.

Unpaired 2 Sample t-Test vs. Paired t-Test

1. Open the **Dietcola.xlsx** file, click the **Sheet 1** tab (or press **F4** to activate last worksheet).
2. Click **SigmaXL > Statistical Tools > 2 Sample t-Test**. Ensure that entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
3. Check **Unstacked Column Format**. Select *After, Before* and click **Numeric Data Variables (Y) >>**. **H0: Mean Diff = 0, Ha: Less Than**; check **Assume Equal Variances**. Uncheck **Display Test Assumptions Report**:



4. Click **OK**. Results:

2 Sample t-Test		
Test Information		
H ₀ : Mean Difference = 0		
H _a : Mean Difference Less Than 0		
Assume Equal Variance		
Results:		
	After	Before
Count	10	10
Mean	6.580	7.600
Standard Deviation	1.871	0.951023
Mean Difference	-1.02	
Std Error Difference	0.663626	
DF	18	
t	1.537	
P-Value (1-sided)	0.0708	
UC (1-sided, 95%)	0.130770	

Now the P-Value is .07, indicating a fail to reject H₀. What changed? Hint: Compare the SE Mean of the Paired t-test to the Std Error Difference of the unpaired two-sample t-test. Where

does the additional variability come from in the two-sample t-test? The paired t-test is the appropriate test to use here.

Power & Sample Size for 2 Sample T-Test

To determine Power & Sample Size for a 2 Sample t-Test, you can use the Power & Sample Size Calculator or Power & Sample Size with Worksheet.

1. Click **SigmaXL > Statistical Tools > Power & Sample Size Calculators > 2 Sample t-Test Calculator**.
2. Select **Solve For Power (1 – Beta)**. Enter **Sample Size** and **Difference** as shown:

The screenshot shows a dialog box titled "Power and Sample Size: 2 Sample t-Test Calculator". It has a "Solve For" section with three radio buttons: "Power (1-Beta)" (selected), "Sample Size (N) for Each Group", and "Difference (Mean1 - Mean2)". Below this are input fields for "Sample Size (N) for Each Group" (30), "Difference (Mean1 - Mean2)" (1), "Standard Deviation" (1.0), "Significance Level (Alpha)" (0.05), and "Ha:" (Not Equal To). On the right side, there are buttons for "OK >>", "Cancel", and "Help".

Note that we are calculating the power or likelihood of detection given that Mean1 – Mean2 = 1, with sample size for each group = 30, standard deviation = 1, significance level = .05, and Ha: Not Equal To (two-sided test).

3. Click **OK**. The resulting report is displayed:

Power and Sample Size: 2 Sample t Test				
H0: Mean 1 = Mean 2				
Ha: Mean 1 ≠ Mean 2				
Solve For: Power (1 - Beta)				
Sample Size (N)	Difference	Standard Deviation	Significance Level (Alpha)	Power (1 - Beta)
30	1	1	0.05	0.967708259

A power value of 0.97 is good, hence we have the basis for the “minimum sample size n=30” rule of thumb used for continuous data.

4. To determine Power & Sample Size using a Worksheet, click **SigmaXL > Statistical Tools > Power & Sample Size with Worksheet > 2 Sample t-Test**.
5. A graph showing the relationship between Power, Sample Size and Difference can then be created using **SigmaXL > Statistical Tools > Power & Sample Size Chart**. See Part E for an example using the 1 Sample t-Test.

Part H – Two Sample F and Comparison Test

Two Sample F-Test Template

Click SigmaXL > Templates and Calculators > Basic Statistical Templates >

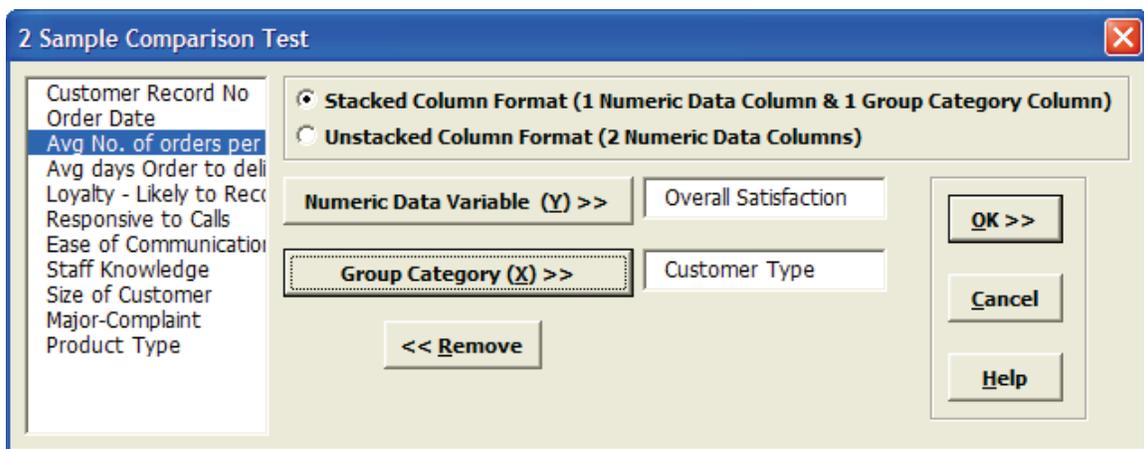
2 Sample F-Test and CI (Compare 2 Standard Deviations) or SigmaXL > Statistical Tools > Basic Statistical Templates > **2 Sample F-Test and CI (Compare 2 Standard Deviations)** to access the 2 Sample F-Test calculator. The template gives the following default example.

 2 Sample F-Test and Ratio Confidence Interval (Compare 2 Standard Deviations)			
Sample Data (user inputs):		Sample 1	Sample 2
Sample Size	n	10	10
Sample Standard Deviation	s	2.0000	1.0000
Null Hypothesis (hypothesized ratio of StDevs)	$H_0: \sigma_1/\sigma_2 =$	1.0	
Alternative Hypothesis	$H_a: \sigma_1/\sigma_2$	Not Equal To	
Confidence Level (enter .95 for 95%)		95.0%	
Results:			
Ratio of Standard Deviations		2.0000	
Sample Variance		4.0000	1.0000
Ratio of Variances		4.0000	
df		9	9
F-statistic		4.0000	
alpha		0.0500	
P-Value (2-sided)		0.0510	
Upper Confidence Limit (2 Sided)		4.0130	
Lower Confidence Limit (2 Sided)		0.9968	

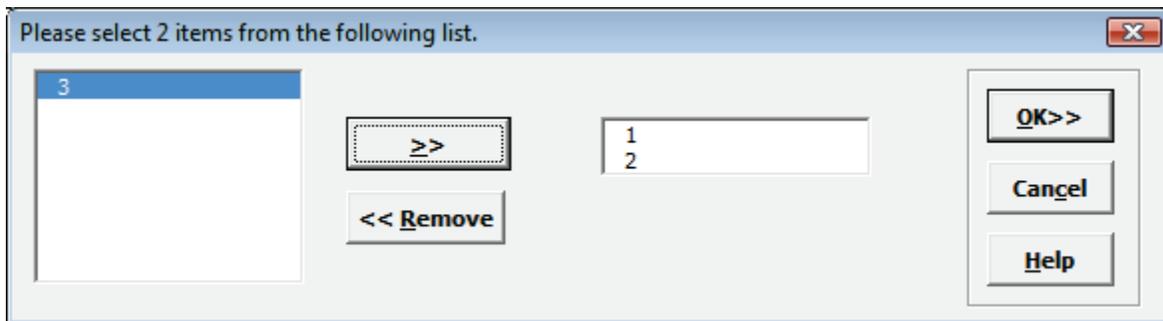
Two Sample Comparison Test

We will now do a full comparison test of Customer Satisfaction for Customer Type 1 and 2. This test checks each sample for normality, equal variance (F-test and Levene's), 2 sample t-test (assuming equal and unequal variance), and Mann Whitney test for equal Medians. Depending on the normality, variance, and sample size results, the appropriate P-Values are highlighted in yellow.

1. Open **Customer Data.xlsx**, click on **Sheet 1** Tab.
2. Click **SigmaXL > Statistical Tools > 2 Sample Comparison Tests**. Click **Next**. Check **Stacked Column Format**. Select *Overall Satisfaction*, click **Numeric Data Variable (Y) >>**, select *Customer Type*, click **Group Category (X) >>**.



3. Click **OK**. Select *Customer Type 1* and *2*.



4. Click OK.

2 Sample Comparison - Overall Satisfaction		
Customer Type	1	2
Count	31	42
Mean	3.394	4.205
Median	3.560	4.340
Standard Deviation	0.824680	0.621200
AD Normality Test p-value	0.5306	0.0302
Test for Equal Variances:		
F-test (use with normal data):		
F (test statistic)	1.762	
p-value (2-sided)	0.0916	
Levene's test (use with non-normal data):		
p-value (2-sided)	0.0443	
2 Sample t-Test for means:		
Assume Equal Variance:		
t (test statistic)	-4.799	
p-value (2-sided)	0.0000	
p-value (1-sided)	0.0000	
Assume Unequal Variance:		
t (test statistic)	-4.601	
p-value (2-sided)	0.0000	
p-value (1-sided)	0.0000	
2 Sample Mann-Whitney test for medians:		
p-value (2-sided)	0.0000	
p-value (1-sided)	0.0000	

Test 1: Anderson Darling
H0: Data are Normal
Ha: Data are Not Normal

Test 2: Equal Variances
H0: $\sigma_1^2 = \sigma_2^2$
Ha: $\sigma_1^2 \neq \sigma_2^2$

Test 3: Equal Means
H0: $\mu_1 = \mu_2$
Ha: $\mu_1 \neq \mu_2$

Test 4: Equal Medians
H0: Median₁ = Median₂
Ha: Median₁ \neq Median₂

Customer Type 2 has nonnormal data. This makes Levene’s test the appropriate test for unequal variance. Levene’s test indicates that Customer type 2 has a significantly lower variance, or standard deviation. The lower standard deviation translates to a **consistent** level of satisfaction.

Since Levene’s test indicates unequal variance, the appropriate t-test assumes unequal variance. The t-test indicates that Customer Type 2 has a significantly higher mean satisfaction.

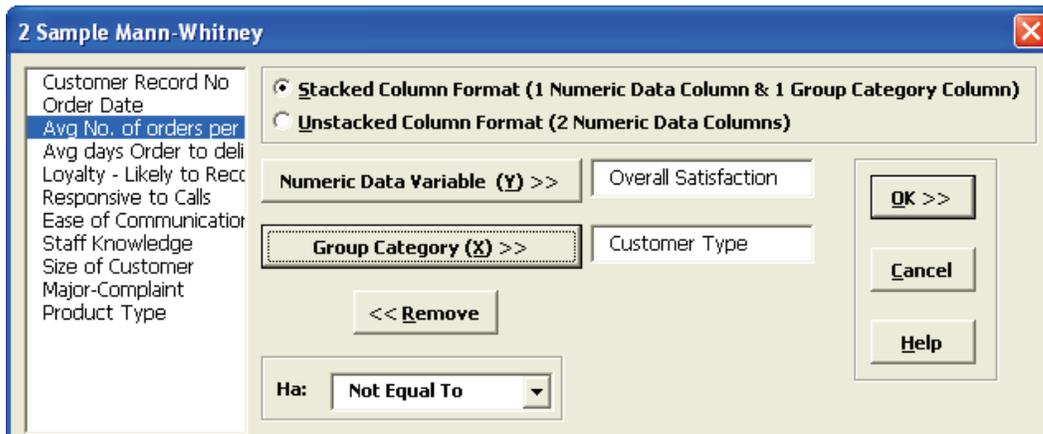
Clearly the next step would be to determine a root cause or best practices to reduce the variability in overall satisfaction and increase the mean for all customer types.

Part I – Two Sample Nonparametric Test: Mann-Whitney

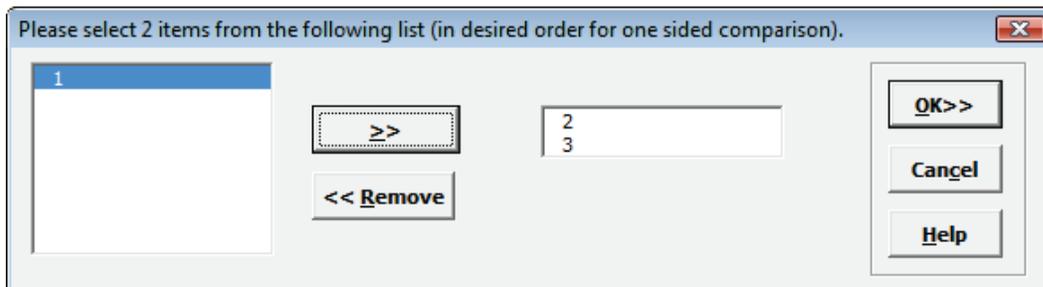
Two Sample Mann-Whitney Test

We will look at comparing medians of Customer Satisfaction by Customer Type, using the Two Sample Mann-Whitney test. (H_0 : Median2 = Median3, H_a : Median2 \neq Median3.) The Two Sample Mann-Whitney Test is the nonparametric equivalent to the parametric Two Sample t-Test. The test does not assume sample normality but does assume that the samples have equal shapes. If the shapes are different, the null hypothesis is that the distributions are the same.

1. Open **Customer Data.xlsx**, click **Sheet 1** tab (or press **F4** to activate last worksheet).
2. Click **SigmaXL > Statistical Tools > Nonparametric Tests > 2 Sample Mann-Whitney**. If necessary, check **Use Entire Data Table**, click **Next**.
3. With **Stacked Column Format** checked, select *Overall Satisfaction*, click **Numeric Data Variable (Y) >>**; select *Customer Type*, click **Group Category (X) >>**; and **Ha: Not Equal To**:



4. Click **OK**. Select *Customer Type 2* and 3.



5. Click **OK**. Resulting output:

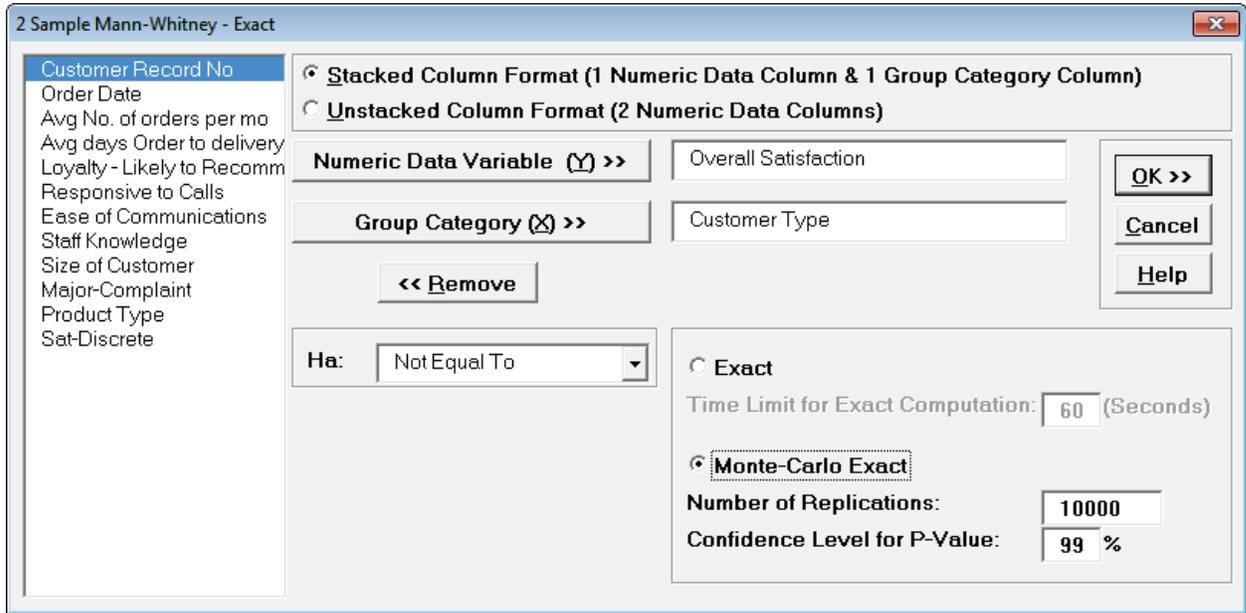
2 Sample Mann-Whitney - Overall Satisfaction		
Test Information		
H ₀ : Median Difference = 0		
H _a : Median Difference ≠ 0		
Customer Type	2	3
Count	42	27
Median	4.340	3.510
Mann-Whitney Statistic	1744.00	
P-Value (2-sided, adjusted for ties)	0.0008	

Given the P-Value of .0008 we reject H₀ and conclude that Median Customer Satisfaction is significantly different between Customer types 2 and 3. This confirms previous findings and matches the results of the 2 Sample t-Test.

Two Sample Mann-Whitney Test - Exact

We will now redo the above example to compute exact P-Values. Typically, this would not be necessary unless the sample sizes were smaller (each sample $N \leq 10$ for Mann-Whitney), but this gives us continuity on the example. Due to the large number of permutations, we will use Monte Carlo Exact for this analysis and will consider a small sample problem later.

1. Open **Customer Data.xlsx**, click **Sheet 1** tab (or press **F4** to activate last worksheet).
2. Click **SigmaXL > Statistical Tools > Nonparametric Tests – Exact > 2 Sample Mann-Whitney – Exact**. If necessary, check **Use Entire Data Table**, click **Next**.
3. With **Stacked Column Format** checked, select *Overall Satisfaction*, click **Numeric Data Variable (Y) >>**; select *Customer Type*, click **Group Category (X) >>** and **Ha: Not Equal To**. Select **Monte Carlo Exact** with the default **Number of Replications = 10000** and **Confidence Level for P-Value = 99%**.



Tip 1: If **Exact** is selected and the computation time limit is exceeded, a dialog will prompt you to use Monte Carlo or to increase the computation time.

Tip 2: 10,000 replications will result in a Monte Carlo P-Value that is typically correct to two decimal places. One million (1e6) replications will result in three decimal places of accuracy and typically require less than 60 seconds to solve for any data set.

Tip 3: The Monte Carlo 99% confidence interval for P-Value is **not** the same as a confidence interval on the test statistic due to data sampling error. The confidence level for the hypothesis test statistic is still 95%, so **all reported P-Values less than .05 will be highlighted in red** to indicate significance. The 99% Monte Carlo P-Value confidence interval is due to the uncertainty in Monte Carlo sampling, and it becomes smaller as the number of replications increases (irrespective of the data sample size). The Exact P-Value will lie within the stated Monte Carlo confidence interval 99% of the time.

4. Click **OK**. Select *Customer Type 2* and *3*. Click **OK**. Results:

2 Sample Mann-Whitney - Monte Carlo: Overall Satisfaction		
Test Information		
H ₀ : Median Difference = 0		
H _a : Median Difference ≠ 0		
Customer Type	2	3
Count	42	27
Median	4.340	3.510
Mann-Whitney Statistic	1744.00	
Monte Carlo P-Value (2-sided)	0.0004	
Monte Carlo P-Value 99% CI Upper	0.0009	
Monte Carlo P-Value 99% CI Lower	0.0000	

Given the Monte Carlo P-Value of .0004 we reject H₀ and conclude that Median Customer Satisfaction is significantly different between Customer types 2 and 3. The Mann Whitney Statistic is very close to the above “large sample” or “asymptotic” result. This was expected because the sample size is reasonable (each sample N > 10), so the “large sample” P-Values are valid using a normal approximation for the Mann-Whitney Statistic.

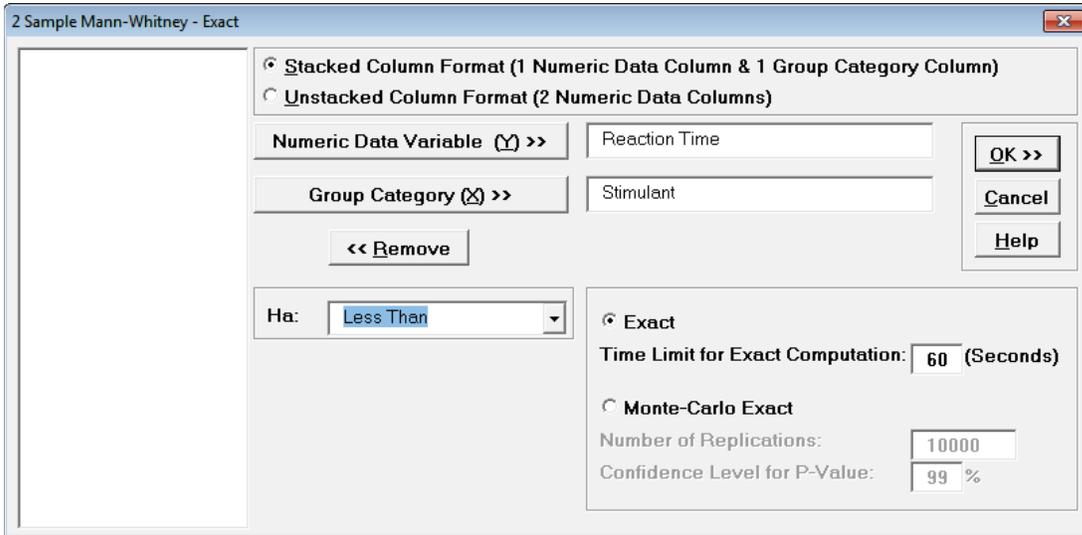
The Monte Carlo P-Value 99% confidence interval is 0.0000 to 0.0009. Note that the Monte Carlo P-Value will be slightly different every time it is run (the Monte Carlo seed value is derived from the system clock). This was demonstrated using 10,000 replications, but with a P-Value this low, it is recommended that the number of replications be increased to 1e5 or 1e6 to get a better estimate.

5. Now we will consider a small sample problem. Open **Stimulant Test.xlsx**. This data is from:

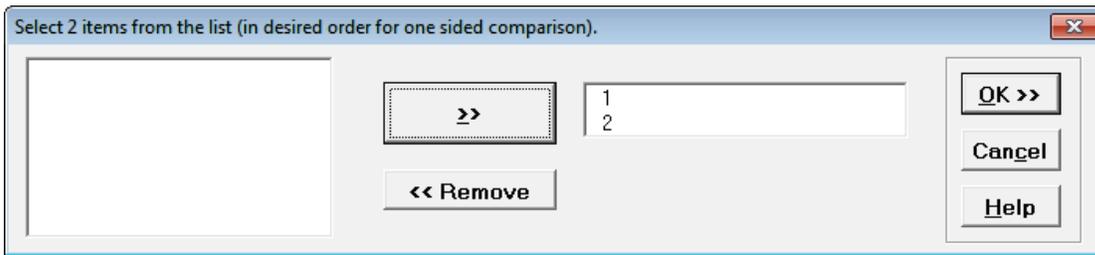
Narayanan, A. and Watts, D. “Exact Methods in the NPAR1WAY Procedure,” SAS Institute Inc., Cary, NC. <http://support.sas.com/rnd/app/stat/papers/exact.pdf>

Researchers conducted an experiment to compare the effects of two stimulants. Thirteen randomly selected subjects received the first stimulant, and six randomly selected subjects received the second stimulant. The reaction times are in minutes. We will test the null hypothesis of no difference between the medians of the two stimulants against the alternative that stimulant 1 has smaller median reaction time than stimulant 2.

6. Select **Reaction Time** tab. Click **SigmaXL > Statistical Tools > Nonparametric Tests – Exact > 2 Sample Mann-Whitney - Exact**. If necessary, check **Use Entire Data Table**, click **Next**.
7. With **Stacked Column Format** checked, select *Reaction Time*, click **Numeric Data Variable (Y) >>**; select *Stimulant*, click **Group Category (X) >>**; and **Ha: Less Than**. Select **Exact** with the default **Time Limit for Exact Computation** = 60 seconds.



8. Click **OK**. Select *Stimulant 1* and *2*.



This sets the order for the one-sided test, so the alternative hypothesis H_a is Median 1 < Median 2.

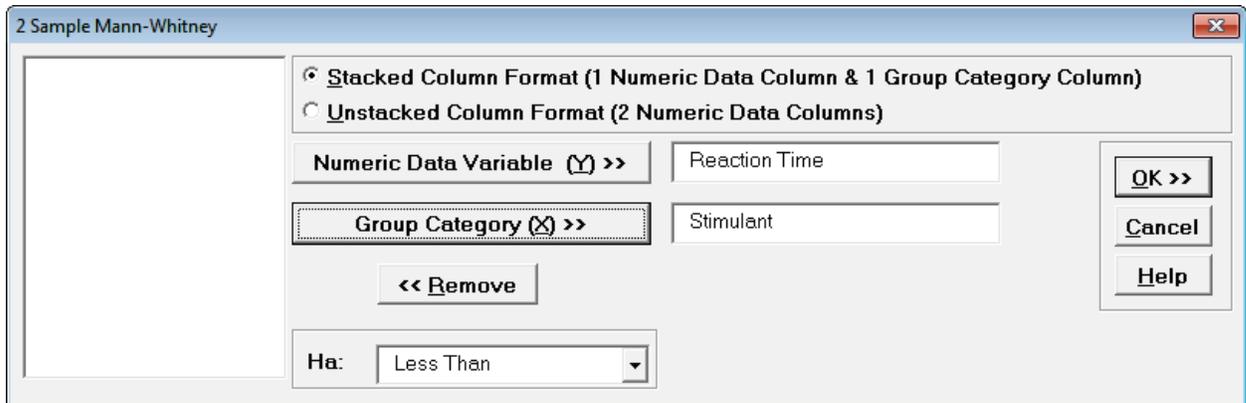
9. Click **OK**. Results:

2 Sample Mann-Whitney - Exact - Reaction Time		
Test Information		
H_0 : Median Difference = 0		
H_a : Median Difference Less Than 0		
Stimulant	1	2
Count	13	6
Median	3.270	3.485
Mann-Whitney Statistic	110.50	
Exact P-Value (1-sided)	0.0527	

With the P-Value = .0527 we fail to reject H_0 , so cannot conclude that there is a difference in median reaction times. This exact P-Value matches that given in the reference paper.

By way of comparison we will now rerun the analysis using the “large sample” or “asymptotic” Mann-Whitney test.

10. Select **Reaction Time** tab (or press **F4** to activate last worksheet). Click **SigmaXL > Statistical Tools > Nonparametric Tests > 2 Sample Mann-Whitney**. If necessary, check **Use Entire Data Table**, click **Next**.
11. With **Stacked Column Format** checked, select *Reaction Time*, click **Numeric Data Variable (Y) >>**; select *Stimulant*, click **Group Category (X) >>**; and **Ha: Less Than**.



12. Click **OK**. Select *Stimulant 1* and *2*. Click **OK**. Results:

2 Sample Mann-Whitney - Reaction Time		
Test Information		
H ₀ : Median Difference = 0		
H _a : Median Difference Less Than 0		
Stimulant	1	2
Count	13	6
Median	3.270	3.485
Mann-Whitney Statistic	110.50	
P-Value (1-sided, adjusted for ties)	0.0421	

Now with the P-Value = .0421 we **incorrectly reject H₀**.

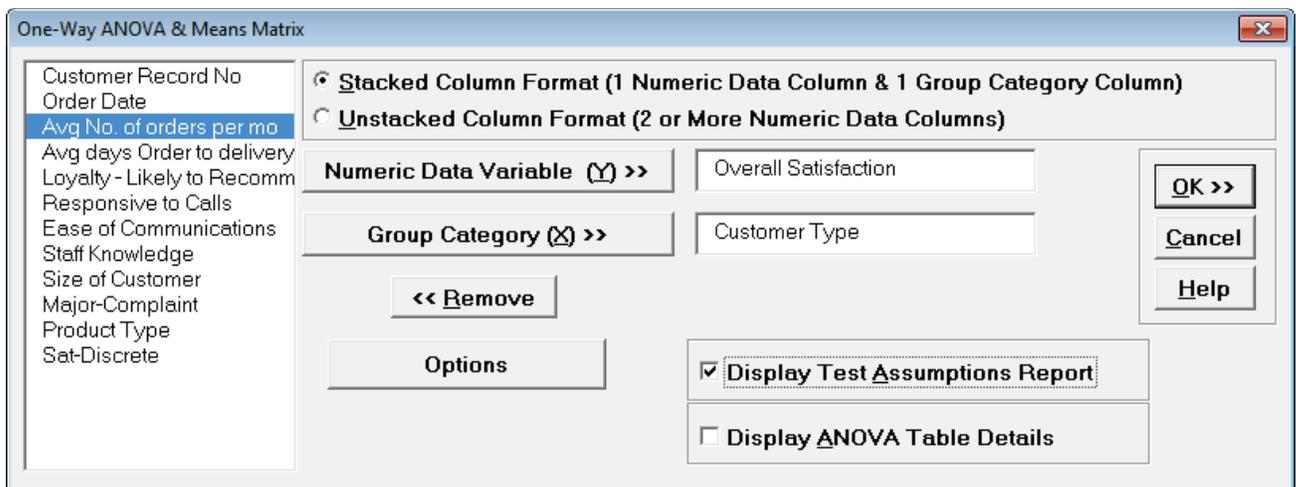
The difference between exact and large sample P-Value is small but it was enough to lead us to falsely conclude that stimulant 1 resulted in a reduced median reaction time.

In conclusion, whenever you have a small sample size and are performing a nonparametric test, always use the Exact option.

Part J – One-Way ANOVA & Means Matrix

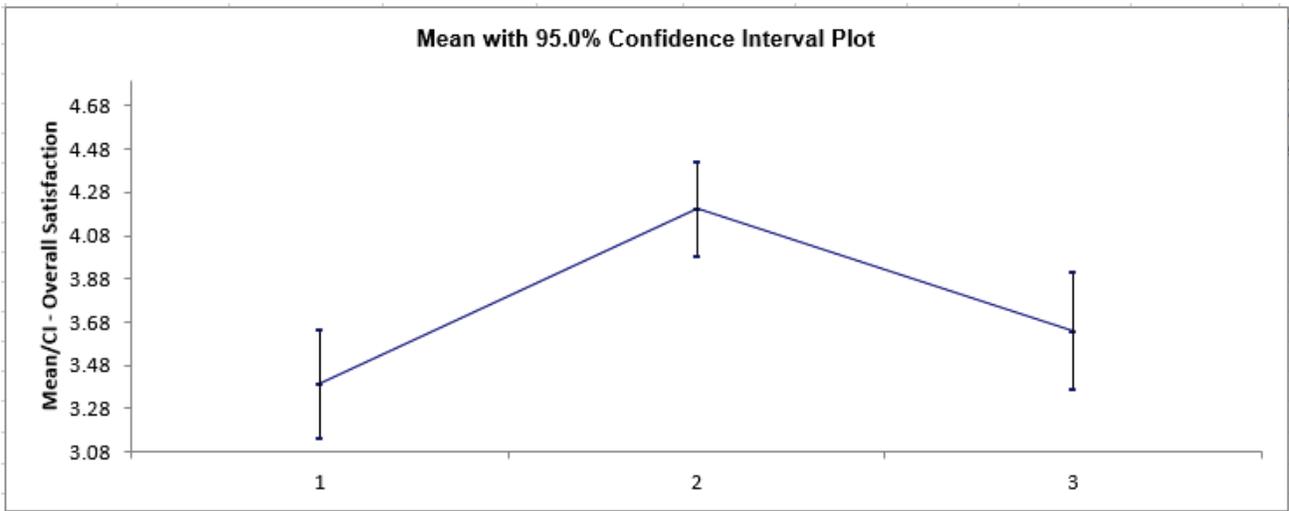
One-Way ANOVA & Means Matrix

1. One-Way ANOVA and Means Matrix allows you to quickly do multiple pairwise comparisons. The One-Way ANOVA tests $H_0: \mu_1 = \mu_2 = \mu_3$; H_a : at least one pairwise set of means are not equal.
2. Open **Customer Data.xlsx**, click on **Sheet 1** tab (or press **F4** to activate last worksheet).
3. Click **SigmaXL > Statistical Tools > One-Way ANOVA & Means Matrix**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**.
4. Click **Next**. Ensure that **Stacked Column Format** is checked. Select *Overall Satisfaction*, click **Numeric Data Variable (Y) >>**; select *Customer Type*, click **Group Category (X) >>**. Uncheck **Display ANOVA Table Details**. Check **Display Test Assumptions Report**.



5. Click **OK**. The results are shown below:

One-Way ANOVA & Means Matrix: Overall Satisfaction			
H ₀ : Mean 1 = Mean 2 = ... = Mean k			
H _a : At least one pair Mean i ≠ Mean j			
Customer Type	1	2	3
Count	31	42	27
Mean	3.394	4.205	3.641
Standard Deviation	0.824680	0.621200	0.670478
UC (2-sided, 95.0%, pooled)	3.644	4.420	3.910
LC (2-sided, 95.0%, pooled)	3.143	3.990	3.373
Pooled Standard Deviation =		0.702810	R-Sq = 20.95%
DF =		97	R-Sq adj. = 19.32%
F =		12.856	
P-Value =		0.0000	
Pairwise Mean Difference (row - column)	1	2	3
1			
2	0		
3			
		-0.811690	-0.247563
		0	0.564127
			0
Fisher Pairwise Probabilities	1	2	3
1			
2			
3			
		0.0000	0.1840
			0.0016



- The ANOVA P-Value of 0.0000 tells us that at least one pairwise set of means are not equal. From the Pairwise Mean Difference (Means Matrix), we conclude that Mean Overall Satisfaction is significantly different between Customer Type 2 and 3, as well as 1 and 2. Note that the default probabilities are Fisher Pairwise. The P-Values will be slightly different than the previous 2 Sample t-Test results because the variances from all 3 customer types are “pooled” here. This also results in slightly different confidence intervals. See below for more details on the multiple comparison options.
- A graphical view of the Overall Satisfaction Mean and 95% Confidence Intervals are given to complement the Means Matrix. The fact that the CI’s for Customer Type 2 do not overlap those of Type 1 or 3, clearly shows that Customer Type 2 has a significantly higher mean satisfaction score. The overlap of CI’s for Type 1 and 3 shows that the mean satisfaction scores for 1 and 3 are not significantly different.
- Note that the Confidence Level shown in the graph can be modified by clicking the **Options** button and setting **Confidence Level**. This also changes the alpha value used to highlight the P-Values in red:

$$\alpha = (100 - \text{Confidence Level})/100$$

so the default Confidence Level = 95.0%, gives an alpha = 0.05.

9. The R-Square (R-Sq) value of 20.95% indicates that Customer Type “explains” approximately 21% of the variation in Overall Satisfaction. We need to “drill down” to understand the root causes and best practices associated with Customer Type 2.

10. One-Way ANOVA Assumptions Report:

One-Way ANOVA Assumptions Report			
Normality:	Anderson Darling P-Value = 0.531. Fail to reject null hypothesis: "data are sampled from a normal distribution," so conclude that the assumption of normality is not violated.	Anderson Darling P-Value = 0.030. Reject null hypothesis: "data are sampled from a normal distribution," so conclude that the assumption of normality is violated (at 95% confidence level). Skewness value = -0.9680 and Kurtosis value = 0.6796. See robustness and outliers.	Anderson Darling P-Value = 0.360. Fail to reject null hypothesis: "data are sampled from a normal distribution," so conclude that the assumption of normality is not violated.
Robustness:	Minimum sample size for a robust ANOVA test = 2. Since each sample size is greater than this, the ANOVA test is robust to nonnormality.		
Outliers (Boxplot Rules):	No outliers found.	Potential (1.5*IQR) outlier lower count = 1. It is recommended to review the data with graphical tools: Boxplot, Normal Probability Plot, Histogram and Run Chart / Control Chart. Consider using a Nonparametric Test.	No outliers found.
Randomness (Independence):	Nonparametric Runs Test (Exact) P-Value = 0.066. Fail to reject null hypothesis: "data are random," so conclude that the assumption of randomness (serial independence) is not violated.	Nonparametric Runs Test (Exact) P-Value = 1.000. Fail to reject null hypothesis: "data are random," so conclude that the assumption of randomness (serial independence) is not violated.	Nonparametric Runs Test (Exact) P-Value = 1.000. Fail to reject null hypothesis: "data are random," so conclude that the assumption of randomness (serial independence) is not violated.
Equal Variance:	Levene's test for Equal Variances P-Value = 0.115. Fail to reject null hypothesis: "variances are equal," so conclude that the assumption of equal variances (or standard deviations) is not violated.		

This is a text report with color highlight: Green (OK), Yellow (Warning) and Red (Serious Violation).

Each sample is tested for Normality using the Anderson Darling test. If not normal, the minimum sample size for robustness of the ANOVA Test is determined utilizing Monte Carlo regression equations (see [Basic Statistical Templates – Minimum Sample Size for Robust t-Tests and ANOVA](#)). If the sample size is inadequate, a warning is given and the appropriate Nonparametric test is recommended (Kruskal-Wallis if there are no extreme outliers, Mood’s Median if there are extreme outliers).

Each sample is tested for Outliers defined as: Potential: Tukey's Boxplot (> Q3 + 1.5*IQR or < Q1 – 1.5*IQR); Likely: Tukey's Boxplot 2.2*IQR; Extreme: Tukey's Boxplot 3*IQR. If outliers are present, a warning is given and recommendation to review the data with a Boxplot and Normal Probability Plot. Here we have a potential outlier for Customer Type 2.

Tip: If the removal of outlier(s) result in an Anderson Darling P-Value that is > 0.1, a notice is given that excluding the outlier(s), the sample data are inherently normal.

Each sample is tested for Randomness using the Exact Nonparametric Runs Test. If the sample data is not random, a warning is given and recommendation to review the data with a Run Chart.

A test for Equal Variances is also applied. If all sample data are normal, Bartlett's Test is utilized, otherwise Levene's Test is used. If the variances are unequal, then a warning is given and Welch's ANOVA is recommended.

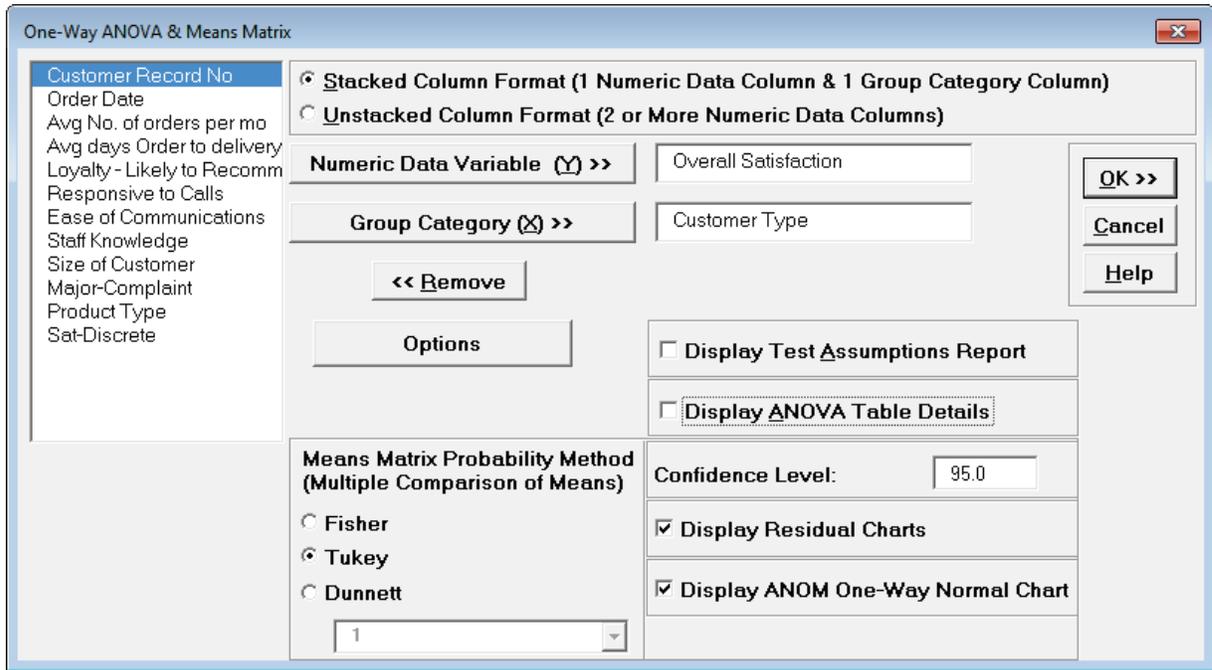
See Appendix [Hypothesis Test Assumptions Report](#) for further details.

One-Way ANOVA & Means Matrix - Options

Notes on Means Matrix Probability Method (Multiple Comparison of Means) and ANOM Chart:

- Fisher:
 - Also known as Fisher's Least Significant Difference (LSD)
 - Pairwise 2 sample t-tests with pooled standard deviation
 - Does not correct for family-wise error rate, so should only be used for $k = 3$ means and in the restricted case where the ANOVA P-Value is $< \alpha$ (this is also known as Protected Fisher LSD). For $k = 3$ means, Protected Fisher LSD is more powerful than Tukey.
- Tukey:
 - Similar to LSD, uses pairwise tests with pooled standard deviation, but is a studentized range statistic that corrects for family-wise error rate. Recommended for $k > 3$.
- Dunnett with Control:
 - If one of the groups are a control reference group, Dunnett with Control is more powerful than Tukey because it is doing fewer pairwise comparisons (only considers those pairwise against the control group).
 - Uses pooled standard deviation and a multivariate t distribution that corrects for family-wise error rate.
- See Appendix [Multiple Comparison of Means and Variances \(a.k.a. Post-Hoc Tests\)](#) for further details and references.
- Display ANOM Normal One-Way Chart:
 - The ANOM alpha = $(100 - \text{Confidence Level})/100$.
 - This chart is also available at **SigmaXL > Graphical Tools > Analysis of Means (ANOM)**. See [Part P – Analysis of Means \(ANOM\) Charts](#).

1. Press **F3** or click **Recall SigmaXL Dialog** to recall last dialog. Uncheck **Display Test Assumptions Report**. Click the **Options** button. Select **Tukey**. Check **Display Residual Charts** and **Display ANOM Normal One-Way Chart**; **Confidence Level** = 95.0 as shown:



Note: The **Confidence Level** is used to set the level in the Mean/Confidence Interval Plot, the alpha level ($\alpha = (100 - CI)/100$) used to highlight the P-Values and the alpha level for the ANOM chart. However, the confidence level used in the Residuals Normal Probability Plot is always 95%.

2. Click **OK**. The Pairwise Means Difference (Means Matrix) and Tukey Probability results are:

Pairwise Mean Difference (row - column)		1	2	3
1	1	0	-0.811690	-0.247563
	2		0	0.564127
	3			0

Tukey Probabilities		1	2	3
1	1		0.0000	0.3777
	2			0.0044
	3			

3. The significant Tukey Probability values above have not changed as compared to Fisher, but note that they are larger than Fisher to compensate for the family-wise error rate:

Pairwise Mean Difference (row - column)		1	2	3
1		0	-0.811690	-0.247563
2			0	0.564127
3				0

Fisher Pairwise Probabilities		1	2	3
1			0.0000	0.1840
2				0.0016
3				

4. Press **F3** or click **Recall SigmaXL Dialog** to recall last dialog. Click **Options**. Uncheck all **Display** options and select **Dunnett** with Control Level = 2 (i.e. we are treating Customer Type 2 as the Control Group):

The screenshot shows the 'One-Way ANOVA & Means Matrix' dialog box. The 'Means Matrix Probability Method' is set to 'Dunnett' with a control level of 2. The 'Display' options are all unchecked.

One-Way ANOVA & Means Matrix

Customer Record No
Order Date
Avg No. of orders per mo
Avg days Order to delivery
Loyalty - Likely to Recomm
Responsive to Calls
Ease of Communications
Staff Knowledge
Size of Customer
Major-Complaint
Product Type
Sat-Discrete

Stacked Column Format (1 Numeric Data Column & 1 Group Category Column)
 Unstacked Column Format (2 or More Numeric Data Columns)

Numeric Data Variable (Y) >> Overall Satisfaction

Group Category (X) >> Customer Type

<< **Remove**

Options

Display Test Assumptions Report
 Display ANOVA Table Details
 Display Residual Charts
 Display ANOM One-Way Normal Chart

Means Matrix Probability Method (Multiple Comparison of Means)
 Fisher
 Tukey
 Dunnett
2

Confidence Level: 95.0

OK >>
Cancel
Help

5. Click **OK**. The Pairwise Means Difference (Means Matrix) and Dunnett Probability results are:

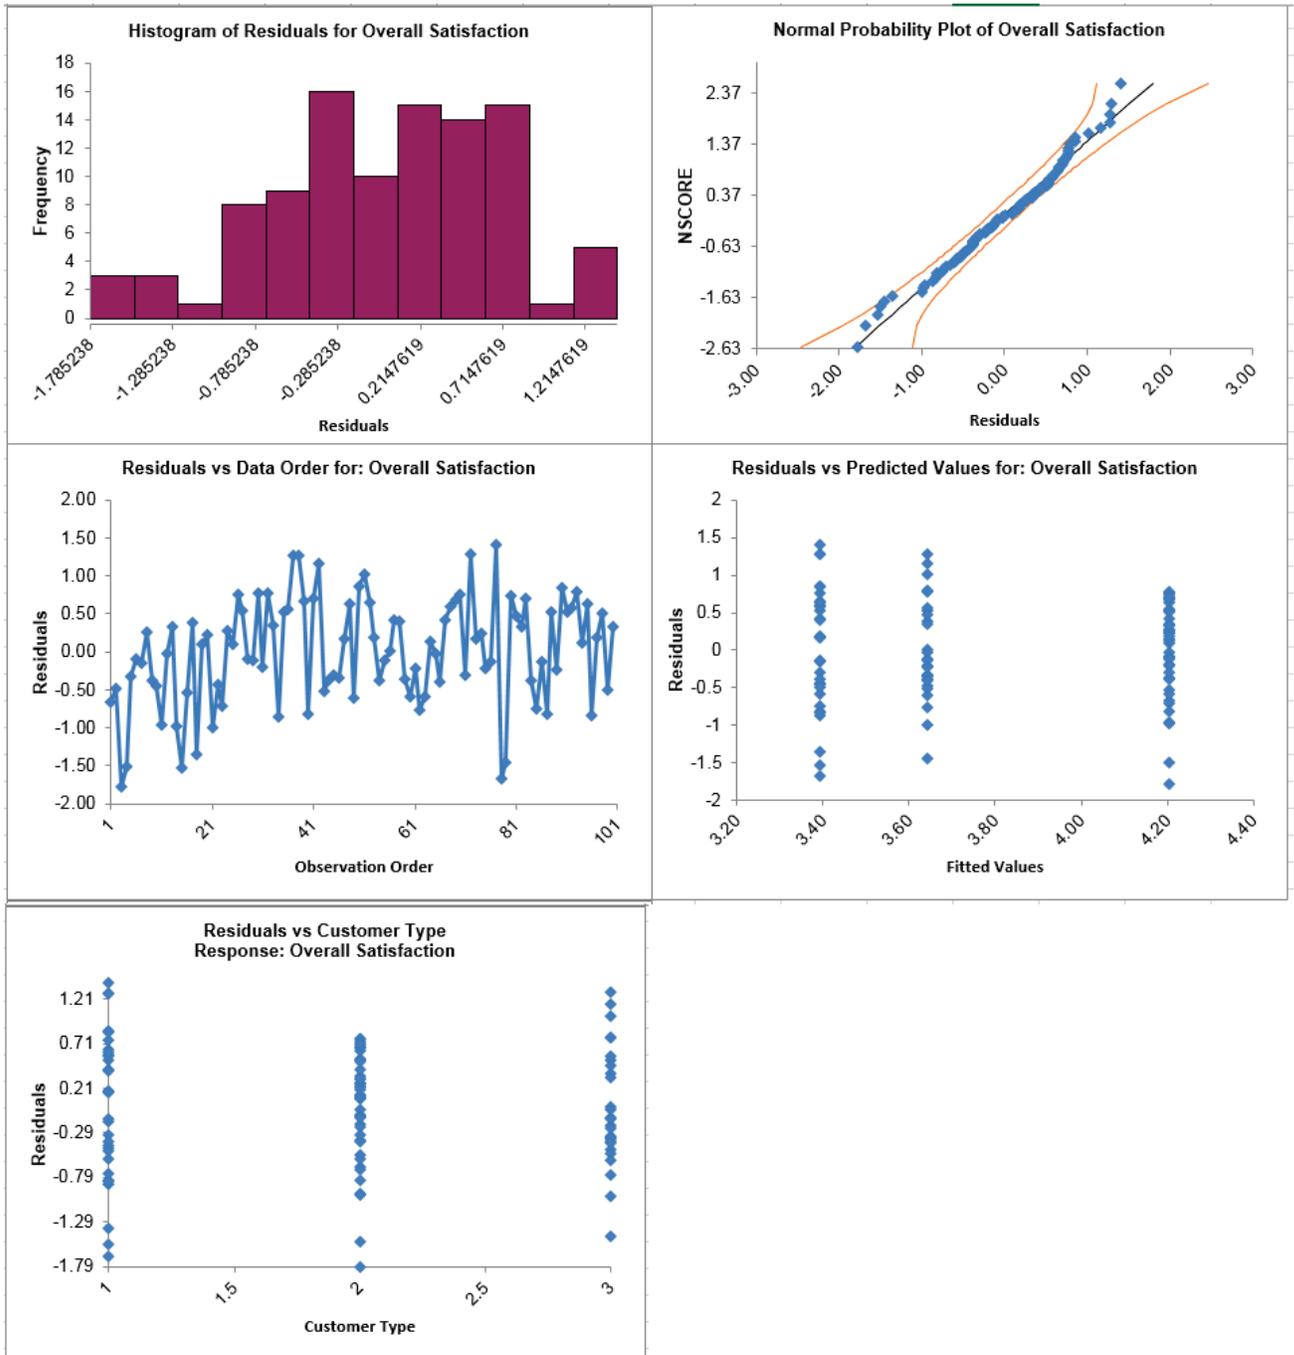
Pairwise Mean Difference (row - column)		1	2	3
1		0	-0.811690	-0.247563
2			0	0.564127
3				0

Dunnett Probabilities (Control = 2)		1	2	3
1			0.0000	
2				0.0031
3				

All three mean differences are shown but only two probability values are calculated: Customer Type 1 – Control Type 2 and Control Type 2 – Type 3. (Typically, one would show 1 – 2 and 3 – 2 for Dunnett, but we are displaying the upper triangle for consistency with the other options).

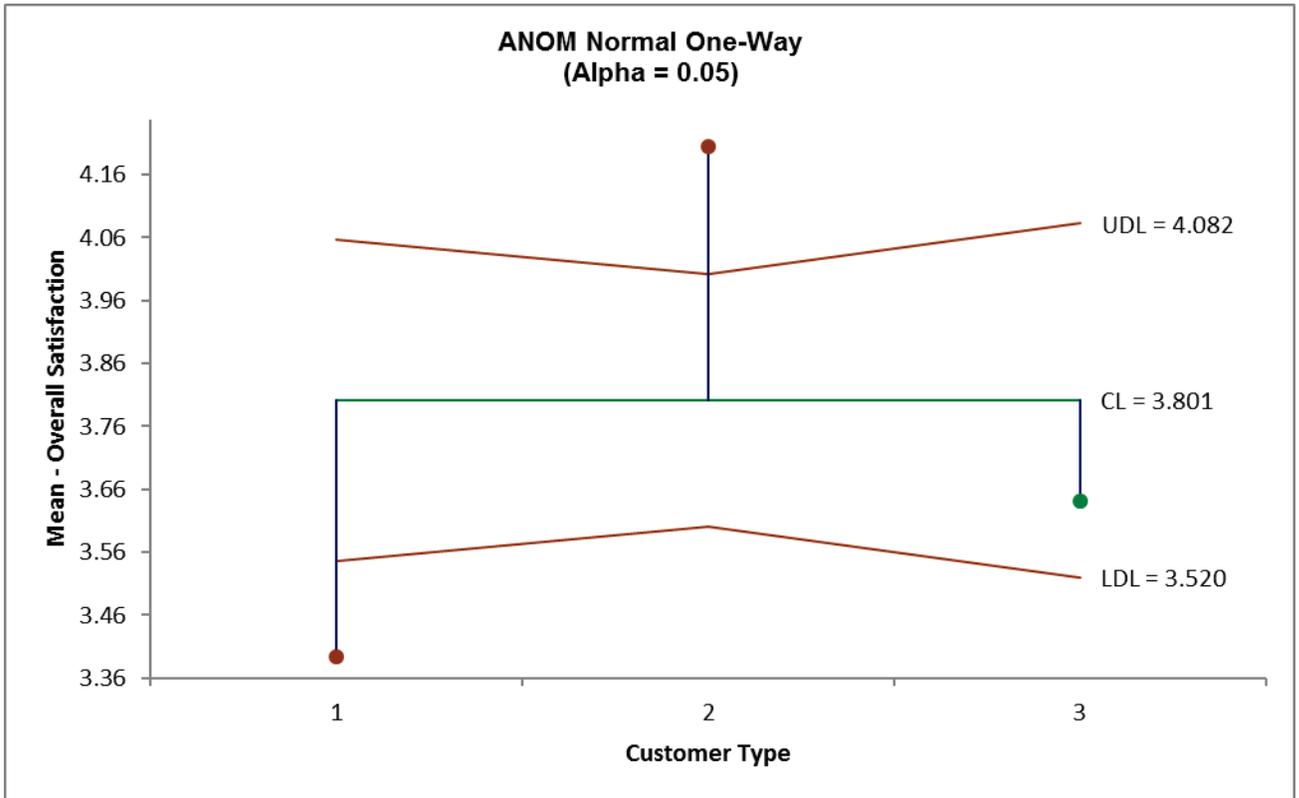
Note that the Dunnett 2 – 3 P-Value is .003 which is smaller than the Tukey .004 but larger than the Fisher .002. Dunnett is more powerful than Tukey (due to having fewer tests), but less powerful than Fisher. It does, however, have an advantage over Fisher because it protects the family-wise error rate for the comparisons being made.

6. Click the **ANOVA Residuals** sheet tab to display the Residual Plots (created at steps 1, 2):



Residuals are the unexplained variation from the ANOVA model (Actual – Predicted/Fitted values). We expect to see the residuals approximately normally distributed with no obvious patterns in the above graphs, which is the case here.

7. Click **ANOM_Normal_One_Way** sheet tab to display the ANOM Chart (created at steps 1, 2):



Here we see that Customer Type 1 mean satisfaction score is significantly below the overall mean and Customer Type 2 is significantly higher. This is consistent with the results that are observed in the Pairwise Means Difference (Means Matrix), but is easier to interpret.

The varying decision limits are due to the varying sample sizes for each Customer Type, with smaller sample size giving wider limits in a manner similar to a control chart. If the data are balanced the decision limit lines will be constant.

Power & Sample Size for One-Way ANOVA

To determine Power & Sample Size for a One-Way ANOVA, you can use the Power & Sample Size Calculator or Power & Sample Size with Worksheet.

1. Click **SigmaXL > Statistical Tools > Power & Sample Size Calculators > One-Way ANOVA Calculator**.
2. Select **Solve For Power (1 – Beta)**. Enter **Sample Size** and **Maximum Difference** as shown:

The screenshot shows the 'Power and Sample Size: One-Way Anova Calculator' dialog box. The 'Solve For' section has three radio buttons: 'Power (1-Beta)' (selected), 'Sample Size (N) for Each Group', and 'Maximum Difference'. The 'Sample Size (N) for Each Group' is set to 30, and 'Maximum Difference' is set to 1. Below this, there are input fields for 'Groups (Levels)' (3), 'Standard Deviation' (1.0), and 'Significance Level (Alpha)' (.05). On the right side, there are three buttons: 'OK >>', 'Cancel', and 'Help'.

Note that we are calculating the power or likelihood of detection given that the maximum difference between group means = 1, with sample size for each group = 30, 3 groups, standard deviation = 1, significance level = .05, and H_a : Not Equal To (two-sided test).

3. Click **OK**. The resulting report is displayed:

Power and Sample Size: One-Way ANOVA					
H0: Mean 1 = Mean 2 = ... = Mean k					
Ha: At least one pair Mean $i \neq$ Mean j					
Solve For: Power (1 - Beta)					
Sample Size (N)	Maximum Difference	Groups	Standard Deviation	Significance Level (Alpha)	Power (1 - Beta)
30	1	3	1	0.05	0.936276828

A power value of 0.94 is acceptable. Note that this value is less than the power value of 0.97 obtained with the two-sample t-Test.

4. Press **F3** or click **Recall SigmaXL Dialog** to recall last dialog. Change the number of groups to 4. Click **OK**. Note that the power value is 0.907. If the number of groups (levels) increases, you will have to increase the sample size in order to maintain statistical power.
5. To determine Power & Sample Size using a Worksheet, click **SigmaXL > Statistical Tools > Power & Sample Size with Worksheet > One-Way ANOVA**.
6. A graph showing the relationship between Power, Sample Size and Maximum Difference can then be created using **SigmaXL > Statistical Tools > Power & Sample Size Chart**. See Part E for an example using the 1 Sample t-Test.

Part K – Two-Way ANOVA

Two-Way ANOVA

Two-Way ANOVA tests the following:

H0 (Factor X1): $\mu_1 = \mu_2 = \dots = \mu_k$

Ha (Factor X1): at least one pairwise set of means are not equal ($\mu_i \neq \mu_j$);

H0 (Factor X2): $\mu_1 = \mu_2 = \dots = \mu_k$

Ha (Factor X2): at least one pairwise set of means are not equal ($\mu_i \neq \mu_j$);

H0 (Interaction): There is **no** interaction between factors X1 and X2

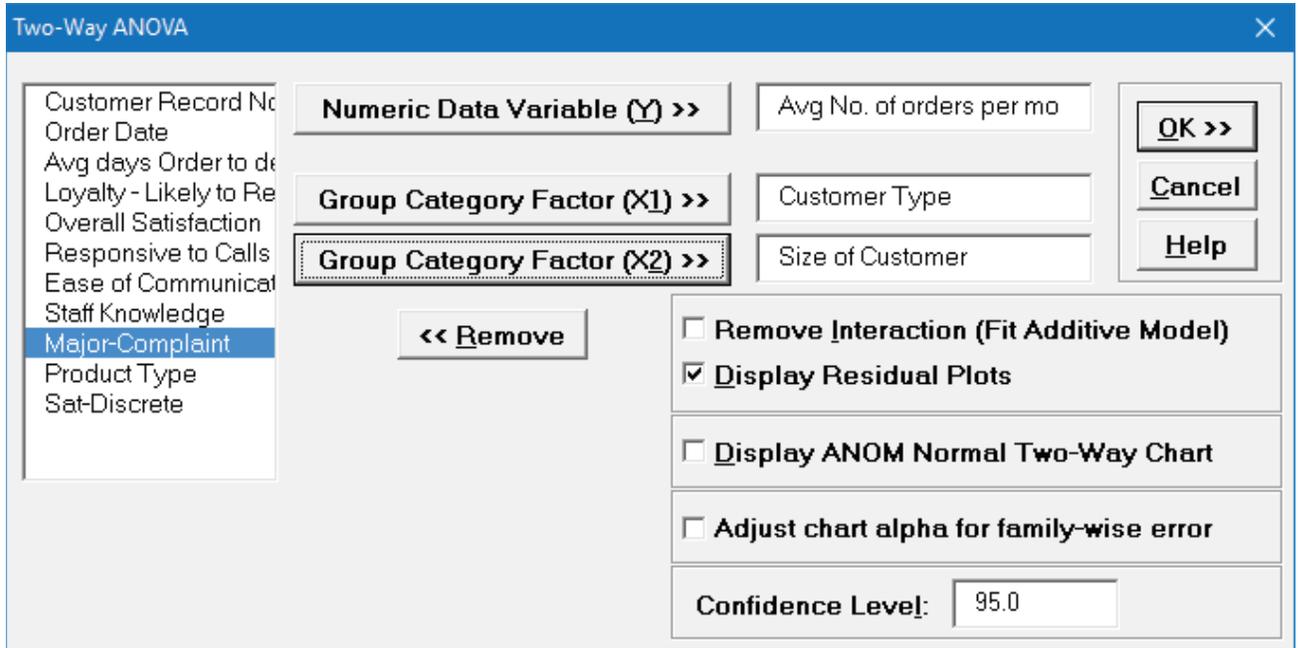
Ha (Interaction): There is an interaction between factors X1 and X2.

A Two-Way ANOVA analysis will typically have balanced data from a designed experiment, with an equal number of observations for each combination level of X1 and X2. SigmaXL will also accommodate unbalanced data. The minimum requirement is one observation per combination level of X1 and X2. An error message will be produced if this minimum requirement is not met.

1. Open **Customer Data.xlsx**, click on **Sheet 1** tab (or press **F4** to activate last worksheet).
2. Click **SigmaXL > Statistical Tools > Two-Way ANOVA**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**.
3. Click **Next**. Select *Avg No of Orders per Mo*, click **Numeric Data Variable (Y) >>**; select *Customer Type*, click **Group Category Factor (X1) >>**; select *Size of Customer*, click **Group Category Factor (X2) >>**. **Remove Interaction (Fit Additive Model)** should remain unchecked.

Tip: If the Two-Way ANOVA report gives a P-Value for the interaction that is high (> 0.1), you should then press **F3** or click **Recall SigmaXL Dialog** to recall this dialog, check **Remove Interaction (Fit Additive Model)** and rerun the analysis.

Display ANOM Normal Two-Way Chart should be unchecked. We will discuss this tool later in [Part P – Analysis of Means \(ANOM\) Charts](#). Note that this option is only available when **Remove Interaction (Fit Additive Model)** is unchecked because ANOM Normal Two-Way always includes the interaction term in the model. The ANOM alpha is $(100 - \text{Confidence Level})/100$.



4. Click **OK**. The results are shown below:

Two-Way ANOVA: Avg No. of orders per mo

H0 (Factor Customer Type): Mean 1 = Mean 2 = ... = Mean k
 Ha (Factor Customer Type): At least one pair Mean i ≠ Mean j

H0 (Factor Size of Customer): Mean 1 = Mean 2 = ... = Mean k
 Ha (Factor Size of Customer): At least one pair Mean i ≠ Mean j

H0 (Interaction): There is no interaction between factors X1 and X2
 Ha (Interaction): There is an interaction between factors X1 and X2

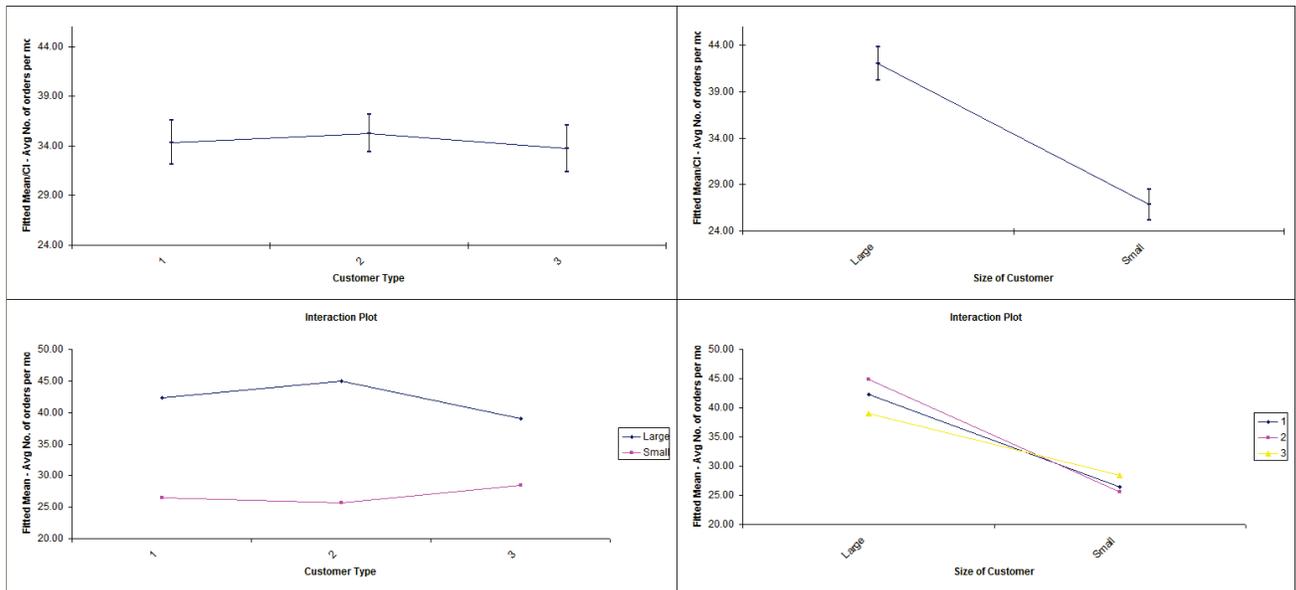
Factor and Model Summary:

Number of Levels - Customer Type	3
Number of Levels - Size of Customer	2
Number of Replicates	N/A
Design Type:	Unbalanced
Confidence Level	95.00%
R-Square	65.31%
R-Square Adjusted	63.46%
S (Pooled Standard Deviation)	6.162

Analysis of Variance:

Source	DF	SS	MS	F	P
Customer Type	2	39.443	19.721	0.519375	0.5966
Size of Customer	1	5298.9	5298.9	139.55	0.0000
Interaction	2	294.72	147.36	3.881	0.0240
Error	94	3569.3	37.971		
Total	99	10288	103.92		

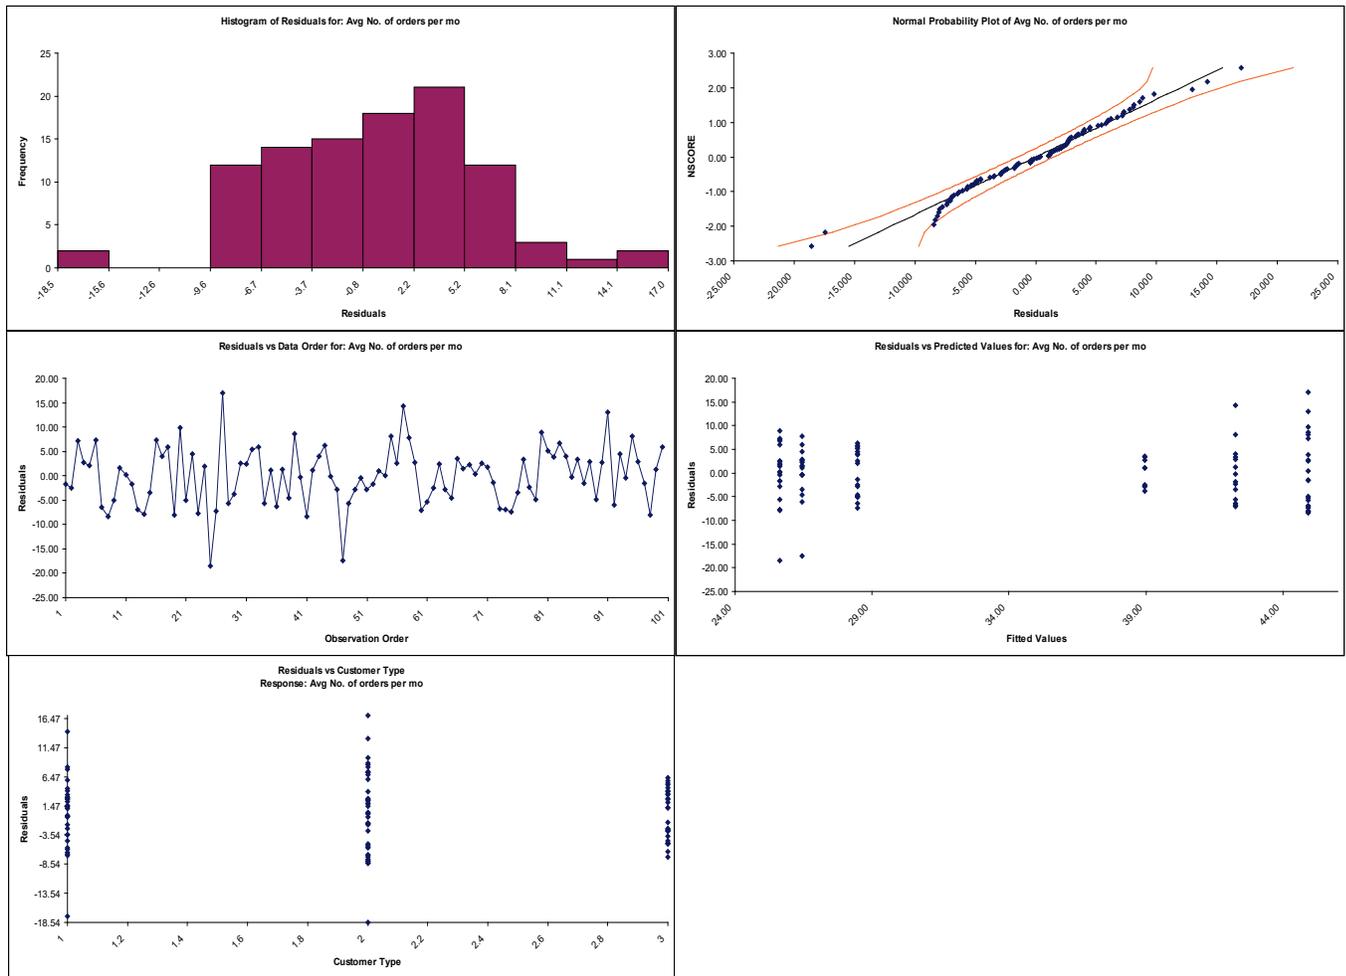
5. Scroll down to view the Main Effects (with confidence intervals) and Interaction Plots:



Note that the mean values shown are fitted (predicted least squares) means not data means. This compensates for unbalanced data.

6. From the ANOVA table, we can see that the Size of Customer term is significant with a P-Value less than .05. Customer Type is not significant by itself, but the Interaction term is significant. This indicates that the effect of Size of Customer on Average Number of Orders per Month depends on Customer Type (we could also equivalently say that the effect of Customer Type depends on Size of Customer).
7. This is also confirmed looking at the Main Effects and Interaction plots. Customer Type by itself is not significant. Size of Customer is obviously significant. Looking at the Interaction plot, the different slopes illustrate that the change in Average Number of Orders per Month across Customer Types depends on Customer Size, albeit this is a relatively small effect.

8. Click the **Two-Way Residuals** sheet to view the residual graphs:



Residuals are the unexplained variation from the ANOVA model (Actual – Predicted/Fitted values). We expect to see the residuals approximately normally distributed with no obvious patterns in the above graphs, which is the case here.

The Residuals versus Size of Customer graph is not shown because Size of Customer is text. In order to display this plot, Size could be coded numerically with 1 = Small and 2 = Large. Simply create a new column called Size-Coded and use the following Excel formula to create the coded values for the first record:

`=IF(K2="Small",1,IF(K2="Large",2))`

Copy and Paste this formula to obtain coded values for all 100 records. Rerun the Two-Way ANOVA analysis to create the residual graphs

Part L – Tests for Equal Variance & Welch’s ANOVA

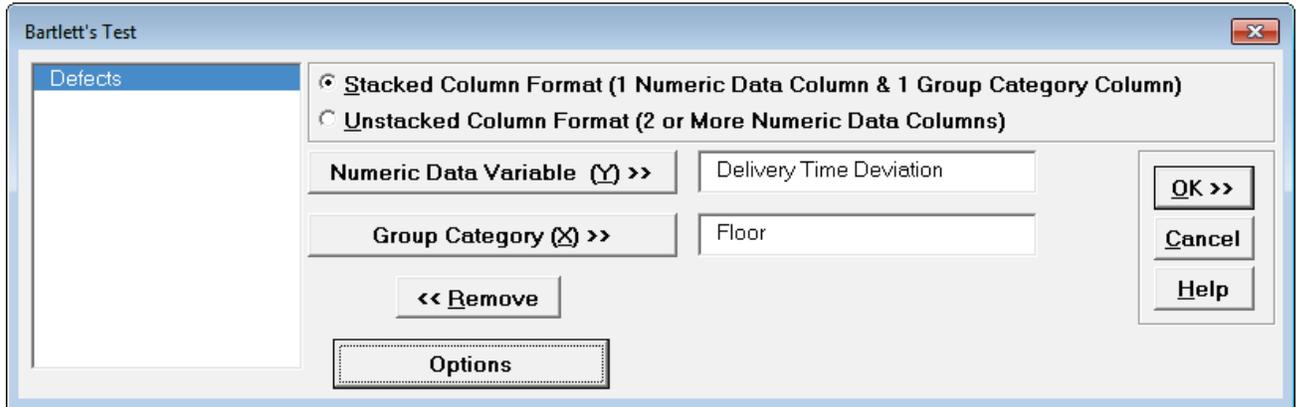
Bartlett’s Test

Bartlett’s Test is similar to the 2 sample F-Test (**SigmaXL > Statistical Tools > 2 Sample Comparison Test**) but allows for multiple group comparison of variances (or standard deviations). Like the F-Test, Bartlett’s requires that the data from each group be normally distributed but is more powerful than Levene’s Test.

Notes on Multiple Comparison of Variances Probability Method and ANOM Chart:

- F-Test Pairwise
 - Pairwise 2 sample F-tests
 - Does not correct for family-wise error rate, so should only be used for $k = 3$ groups and in the restricted case where the Bartlett P-Value is $< \alpha$.
- F-Test with Bonferroni Correction
 - Pairwise 2 sample F-tests with Bonferroni correction
 - Recommended for $k > 3$
 - Bonferroni P-Value' = P-Value * m
 - $m = \text{number of pairwise comparisons } k(k-1)/2$
- See Appendix [Multiple Comparison of Means and Variances \(a.k.a. Post-Hoc Tests\)](#) for further details and references.
- Display ANOM Variances Chart:
 - The ANOM alpha = $(100 - \text{Confidence Level})/100$.
 - This chart is also available at **SigmaXL > Graphical Tools > Analysis of Means (ANOM)**. See [Part P – Analysis of Means \(ANOM\) Charts](#).

1. Open **Delivery Times.xlsx**, click on **Sheet 1** tab.
2. Click **SigmaXL > Statistical Tools > Equal Variance Tests > Bartlett**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**.
3. Click **Next**. Ensure that **Stacked Column Format** is checked. Select *Delivery Time Deviation*, click **Numeric Data Variable (Y) >>**; select *Floor*, click **Group Category (X) >>**.



4. Click **OK**. The results are shown below:

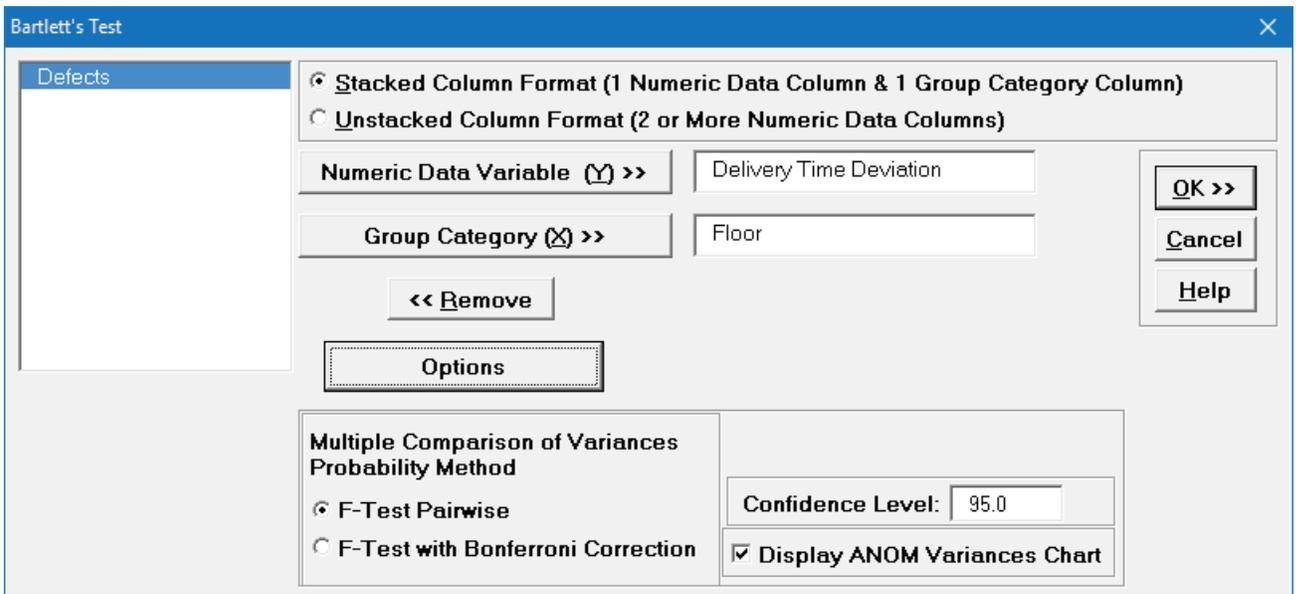
Bartlett's Test For Equal Variance: Delivery Time Deviation										
(Use with normal data)										
Test Information										
H ₀ : Variance 1 = Variance 2 = ... = Variance k										
H _a : At least one pair Variance i ≠ Variance j										
Floor	1	2	3	4	5	6	7	8	9	10
Count	73	73	73	73	73	72	72	72	72	72
Mean	4.989	6.064	6.242	5.581	6.250	8.004	6.647	3.649	6.539	6.084
Median	4.847	6.411	6.476	6.061	7.148	7.261	5.906	3.619	6.799	6.393
StDev	7.019	7.612	7.782	6.827	7.608	6.223	7.015	6.682	7.634	6.653
AD Normality Test P-Value	0.8084	0.9515	0.4024	0.7642	0.0693	0.4844	0.1543	0.4230	0.4014	0.1475
Bartlett's Test Statistic	7.073									
P-Value	0.6295									

5. All 10 Anderson-Darling Test P-Values are > .05 indicating that all group data are normal. Since the assumption of normality is met, Bartlett's is the appropriate test to use. If any one of the groups have a low P-Value for the Normality test, then Levene's test should be used.
6. With the P-Value = 0.63 we fail to reject H₀; we do not have evidence to show that the group variances are unequal (practically speaking we will assume that the variances are equal).
7. If the Equal Variance test is only being used to test the assumption for use in ANOVA then it is not necessary to examine the Multiple Comparison of Variances. However, in the context of a process improvement project, we often do want to know which groups are significantly different. This can give us important clues to identify opportunities for variance reduction.
8. With a "Fail-to-Reject H₀" it is unnecessary to review the Multiple Comparison of Variances or the ANOM Variances Chart, but we will do so here for demonstration purposes.

9. The default Multiple Comparison of Variances is a matrix of F-Test Pairwise Probabilities:

Multiple Comparison of Variances										
F-Test Pairwise Probabilities										
	1	2	3	4	5	6	7	8	9	10
1		0.4934	0.3832	0.8147	0.4958	0.3110	0.9967	0.6789	0.4790	0.6518
2			0.8513	0.3583	0.9970	0.0909	0.4925	0.2735	0.9798	0.2575
3				0.2689	0.8483	0.0607	0.3828	0.2003	0.8717	0.1875
4					0.3602	0.4352	0.8185	0.8568	0.3467	0.8275
5						0.0916	0.4948	0.2751	0.9768	0.2591
6							0.3146	0.5498	0.0871	0.5749
7								0.6830	0.4781	0.6558
8									0.2642	0.9703
9										0.2486
10										

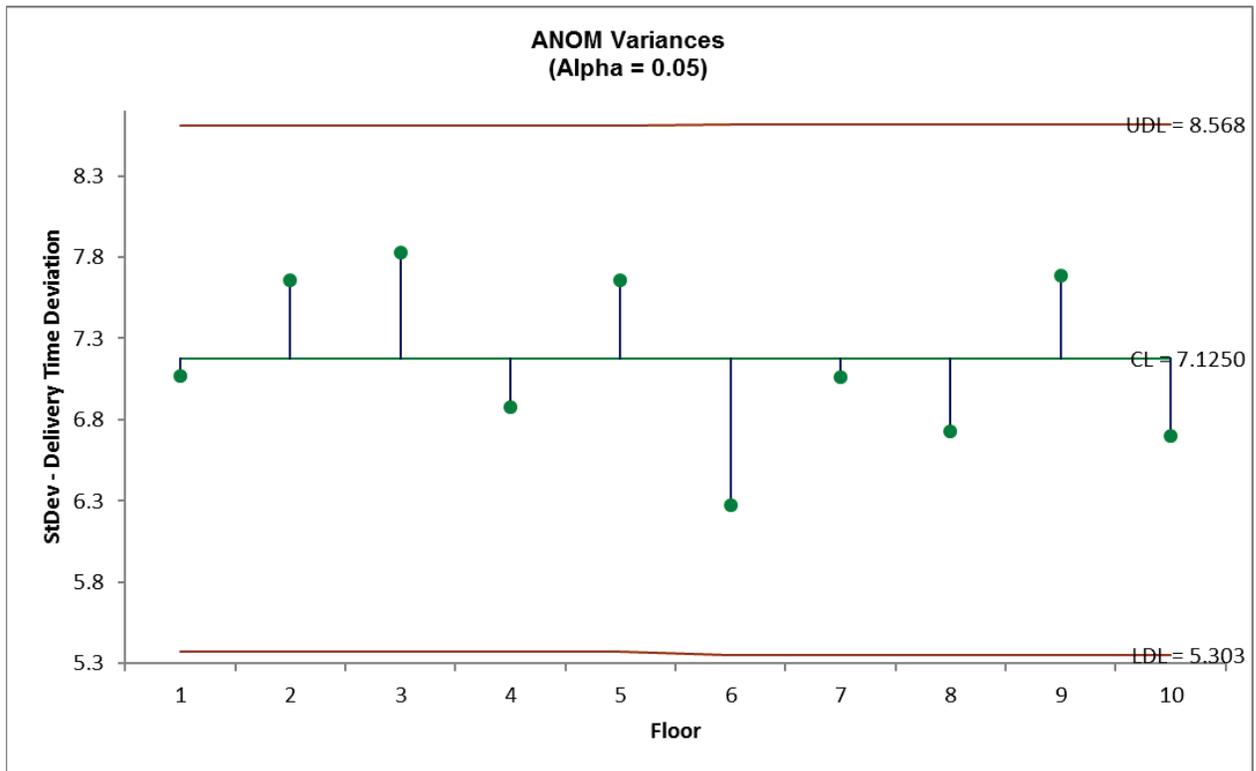
10. Press **F3** or click **Recall SigmaXL Dialog** to recall last dialog. Click **Options**. Check **Display ANOM Variances Chart**.



Note: The **Confidence Level** determines the alpha level ($\alpha = (100 - CI)/100$) used to highlight the P-Values and the alpha level for the ANOM chart. However, the alpha level used to highlight P-Values in the Anderson-Darling Normality Test is always 0.05.

11. We will not run **F-Test with Bonferroni Correction** in this example, but typically that would be used when there are more than 3 groups.

12. Click **OK**. Click **ANOM_Variiances** sheet tab to display the ANOM Chart:



The ANOM Variance chart visually shows that none of the group standard deviations are significantly different from the grand mean of all the standard deviations. It is called an ANOM Variance Chart but displays Standard Deviations for ease of interpretation (similar to a Standard Deviation S Control Chart). This does, however, result in non-symmetrical decision limits. The ANOM Variance chart in **SigmaXL > Graphical Tools > Analysis of Means (ANOM) > ANOM Variance** has an option to display Variances.

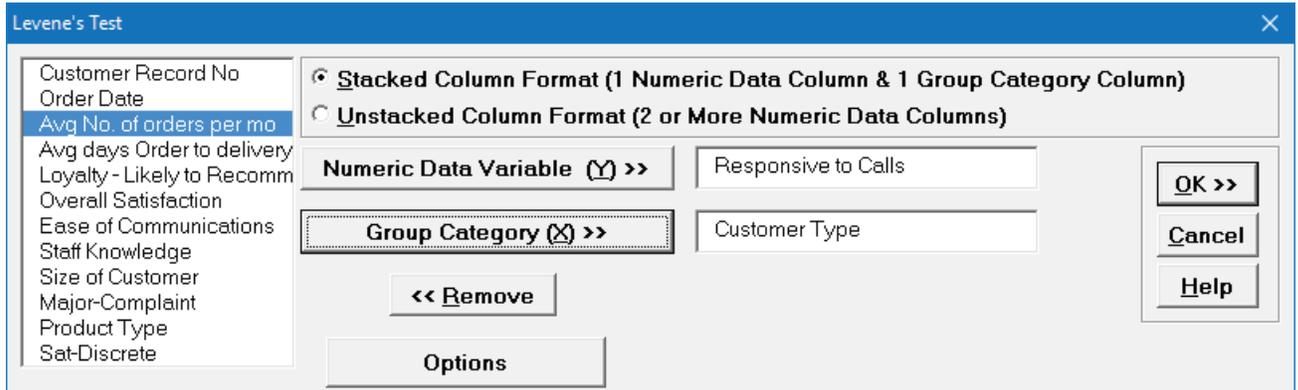
Levene's Test

Levene's Test for multiple group comparison of variances is less powerful than Bartlett's Test, but is robust to the assumption of normality. (This is a modification of the original Levene's Test, sometimes referred to as the Browne-Forsythe Test that uses Absolute Deviations from the Median or ADM).

Notes on Multiple Comparison of Variances Probability Method and ANOM Chart:

- Levene Pairwise
 - Pairwise 2 sample Levene tests
 - Does not correct for family-wise error rate, so should only be used for $k = 3$ groups and in the restricted case where the Levene P-Value is $< \alpha$.
- Tukey ADM (Absolute Deviations from the Median)
 - Application of Tukey on ADM
 - Recommended for $k > 3$
 - This post-hoc test is unique to SigmaXL, inspired by the method used in ANOM Levene Variances.
- See Appendix [Multiple Comparison of Means and Variances \(a.k.a. Post-Hoc Tests\)](#) for further details and references.
- Display ANOM Levene Robust Variances Chart:
 - The ANOM alpha = $(100 - \text{Confidence Level})/100$.
 - This chart is also available at **SigmaXL > Graphical Tools > Analysis of Means (ANOM)**. See [Part P – Analysis of Means \(ANOM\) Charts](#).

1. Open **Customer Data.xlsx**, click on **Sheet 1** tab.
2. Click **SigmaXL > Statistical Tools > Equal Variance Tests > Levene**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**.
3. Click **Next**. Ensure that **Stacked Column Format** is checked. Select *Responsive to Calls*, click **Numeric Data Variable (Y) >>**; select *Customer Type*, click **Group Category (X) >>**.



4. Click **OK**. The results are shown below:

Levene's Test For Equal Variance: Responsive to Calls			
(Use with non-normal data)			
Test Information			
H ₀ : Variance 1 = Variance 2 = ... = Variance k			
H _a : At least one pair Variance i ≠ Variance j			
Customer Type	1	2	3
Count	31	42	27
Mean	3.412	4.226	3.821
Median	3.500	4.720	4.180
StDev	1.304	0.921232	1.091543381
AD Normality Test P-Value	0.0021	0.0000	0.0190
Levene's Test Statistic	4.433		
DF Num	2		
DF Den	97		
P-Value	0.0144		

5. The Levene's Test P-Value of 0.0144 tells us that we reject H₀. At least one pairwise set of variances are not equal. The normality test P-Values indicate that all 3 groups have nonnormal data (P-Values < .05). Since Levene's Test is robust to the assumption of normality, it is the correct test for equal variances (rather than Bartlett's Test).
6. If the Equal Variances test is only being used to test the assumption for use in ANOVA then it is not necessary to examine the Multiple Comparison of Variances. However, in the context of a

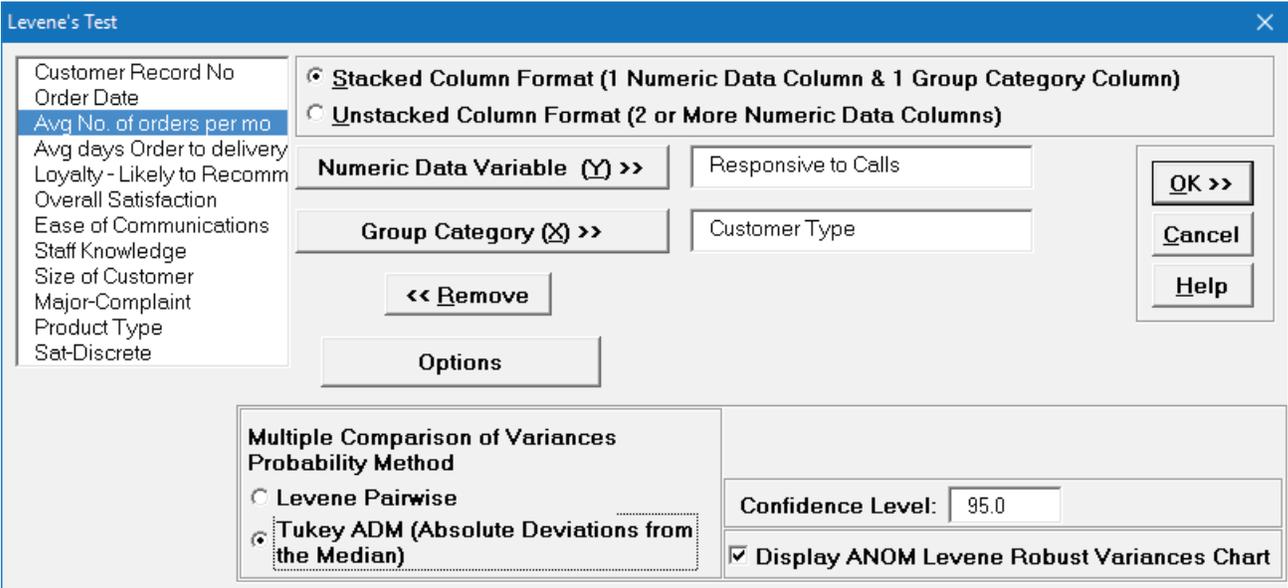
process improvement project, we often do want to know which groups are significantly different. This can give us important clues to identify opportunities for variance reduction.

7. The default Multiple Comparison of Variances is a matrix of Levene Pairwise Probabilities:

Multiple Comparison of Variances			
Levene Pairwise Probabilities	1	2	3
1		0.0036	0.0789
2			0.2991
3			

Customer Type 1 versus Customer Type 2 shows a significant difference in variance.

Press **F3** or click **Recall SigmaXL Dialog** to recall last dialog. Click **Options**. Select **Tukey ADM (Absolute Deviations from Median)**. Check **Display ANOM Levene Robust Variances Chart**.



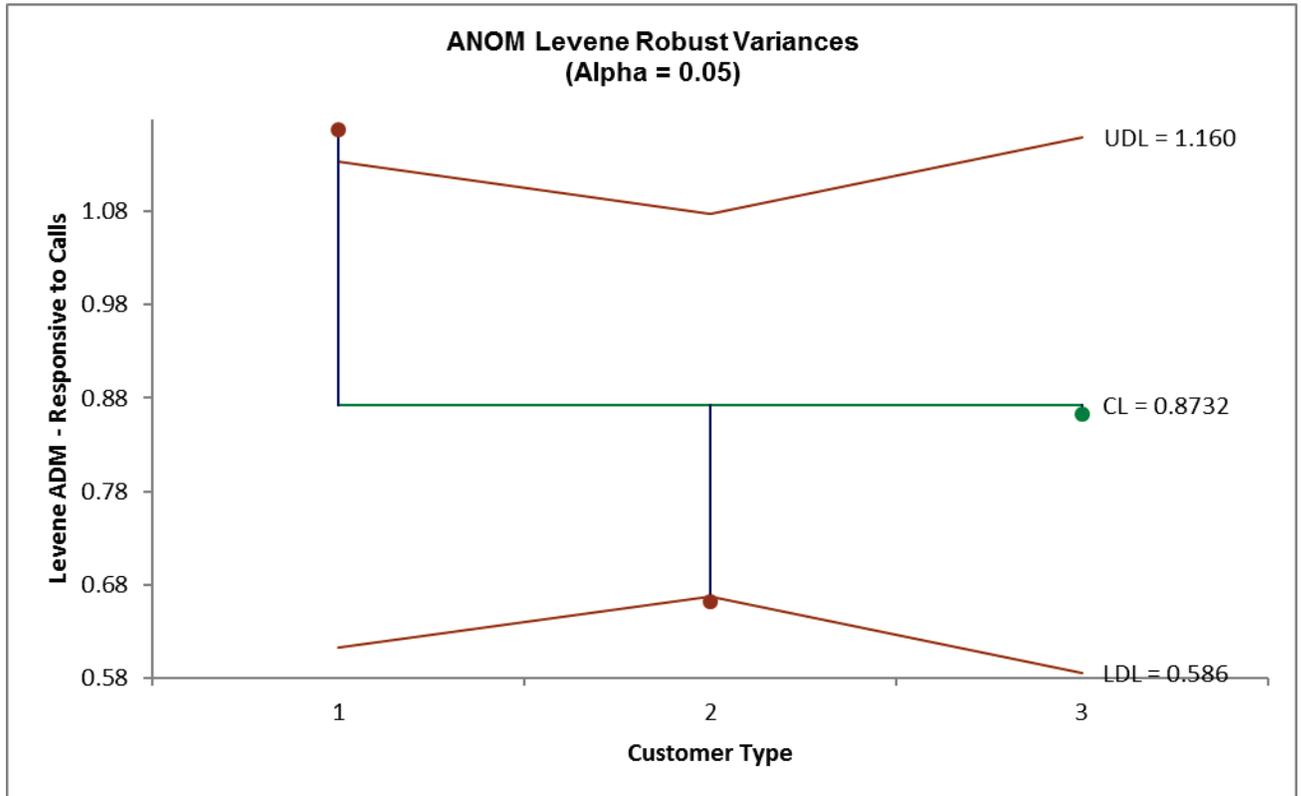
Note: The **Confidence Level** determines the alpha level ($\alpha = (100 - CI)/100$) used to highlight the P-Values and the alpha level for the ANOM chart. However, the alpha level used to highlight P-Values in the Anderson-Darling Normality Test is always 0.05.

8. Click **OK**. The Multiple Comparison of Variances is a matrix of Tukey ADM Probabilities:

Multiple Comparison of Variances			
Tukey ADM Probabilities	1	2	3
1		0.0102	0.2466
2			0.4920
3			

The 1 – 2 P-Value is significant but larger than the Levene Pairwise because it adjusts for the family-wise error rate. Note that it is smaller than the Bonferroni corrected value = $.0036 * 3 = .011$, so more powerful than Bonferroni. The difference in power between Tukey and Bonferroni becomes more prominent with a larger number of groups, so Bonferroni is not included as an option.

9. Click the **ANOM_Levene** sheet tab to display the ANOM Chart:



10. The ANOM chart clearly shows Customer Type 1 has significantly higher variance (ADM) than overall and Customer Type 2 has significantly lower variance.
11. The varying decision limits are due to the varying sample sizes for each Customer Type, with smaller sample size giving wider limits in a manner similar to a control chart. If the data are balanced, the decision limit lines will be constant.
12. Now that we have determined that the variances are not equal, we are presented with a problem if we want to test for equal group means. Classical ANOVA assumes that the group variances are equal, so should not be used. A modified ANOVA called Welch's ANOVA is robust to the assumption of equal variances and will be demonstrated next.

Welch's ANOVA Test (Assume Unequal Variance)

Welch's ANOVA is a test for multiple comparison of means. It is a modified One-Way ANOVA that is robust to the assumption of equal variances. Welch's ANOVA is an extension of the 2 sample t-test for means, assuming unequal variance. Nonparametric methods could also be used here but they are not as powerful as Welch's ANOVA.

Notes on Means Matrix Probability Method (Multiple Comparison of Means):

- Welch Pairwise:
 - Pairwise 2 sample t-tests with unpooled standard deviation and weighted degrees of freedom (2 sample t-test for unequal variance)
 - Does not correct for family-wise error rate, so should only be used for $k = 3$ means and in the restricted case where the Welch ANOVA P-Value is $< \alpha$.
- Games-Howell:
 - Similar to Welch Pairwise, uses unpooled standard deviation and weighted degrees of freedom, but is a studentized range statistic that corrects for family-wise error rate. Recommended for $k > 3$.
 - It is an extension of the Tukey test, but does not assume equal variance.
- See Appendix [Multiple Comparison of Means and Variances \(a.k.a. Post-Hoc Tests\)](#) for further details and references.
- ANOM Chart for Welch ANOVA is not included in SigmaXL since it requires two-stage sampling (see Nelson, Wludyka, Copeland, 2005, Chapter 8, *Heteroscedastic Data*) [References for Analysis of Means \(ANOM\) Charts](#).

SigmaXL: Analyze Phase Tools

1. Open **Customer Data.xlsx**, click on Sheet 1 tab (or press **F4** to activate last worksheet).
2. Click **SigmaXL > Statistical Tools > Equal Variance Tests > Welch's ANOVA (Assume Unequal Variance)** or **SigmaXL > Statistical Tools > Welch's ANOVA (Assume Unequal Variance)**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**.
3. Click **Next**. Ensure that **Stacked Column Format** is checked. Select *Responsive to Calls*, click **Numeric Data Variable (Y) >>**; select *Customer Type*, click **Group Category (X) >>**. Check **Display Test Assumptions Report**.

Welch's ANOVA

Customer Record No
Order Date
Avg No. of orders per mo
Avg days Order to delivery
Loyalty - Likely to Recomm
Overall Satisfaction
Ease of Communications
Staff Knowledge
Size of Customer
Major-Complaint
Product Type
Sat-Discrete

Stacked Column Format (1 Numeric Data Column & 1 Group Category Column)
 Unstacked Column Format (2 or More Numeric Data Columns)

Numeric Data Variable (Y) >> Responsive to Calls

Group Category (X) >> Customer Type

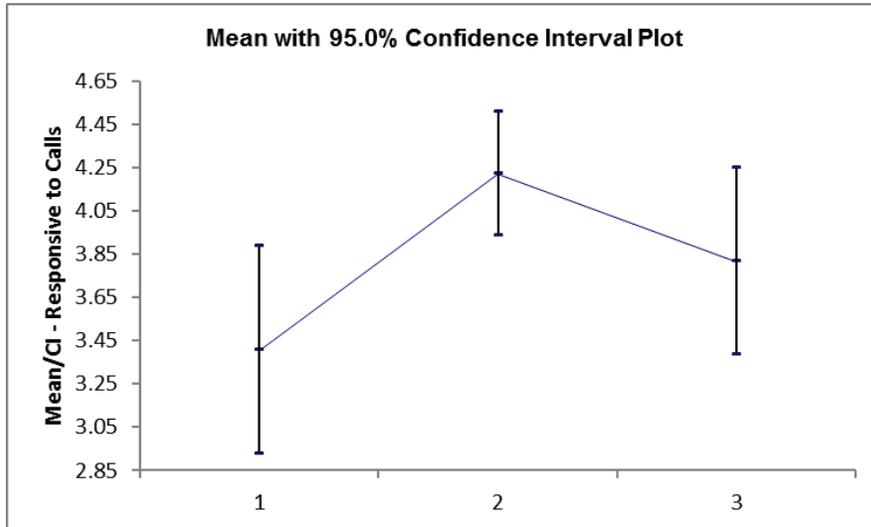
Display Test Assumptions Report

OK >>
Cancel
Help

<< Remove
Options

4. Click **OK**. The results are shown below:

Welch's ANOVA: Responsive to Calls (Use if groups have unequal variance)			
Test Information			
H ₀ : Mean 1 = Mean 2 = ... = Mean k			
H _a : At least one pair Mean i ≠ Mean j			
Customer Type			
Count	1	2	3
	31	42	27
Mean	3.412	4.226	3.821
Median	3.500	4.720	4.180
StDev	1.304	0.921232	1.091543381
AD Normality Test P-Value	0.0021	0.0000	0.0190
Welch's F Ratio			
DF Num	4.654		
DF Den	2		
P-Value	55.753		
	0.0135		
Pairwise Mean Difference (row - column)			
	1	2	3
1		-0.814578	-0.409869
2	0		0.404709
3		0	
Welch Pairwise Probabilities			
	1	2	3
1		0.0045	0.1981
2			0.1170
3			



5. The P-Value for Welch's ANOVA is 0.0135, therefore we reject H_0 and conclude that the group means for Responsive to Calls are not equal.
6. From the Pairwise Mean Difference (Means Matrix), we conclude that Mean Responsive to Calls is significantly different between Customer Type 1 and 2. Note that the default probabilities are Welch Pairwise. See below for more details on the multiple comparison options.
8. A graphical view of the Responsive to Calls Mean and 95% Confidence Intervals are given to complement the Means Matrix. Note that the standard deviations are unpooled, resulting in different CI widths for each group. The fact that the CI's for Customer Type 1 do not overlap those of Type 2, visually shows that there is a significant difference in mean Responsive to Calls. The overlap of CI's for Type 2 and 3 shows that the mean scores for 2 and 3 are not significantly different.
9. Later, we will explore the relationship between Responsive to Calls and Overall Satisfaction.

10. Welch’s ANOVA Assumptions Report:

Welch's ANOVA Assumptions Report			
Normality:	Anderson Darling P-Value = 0.002. Reject null hypothesis: "data are sampled from a normal distribution," so conclude that the assumption of normality is violated (at 95% confidence level). Skewness value = -0.3047 and Kurtosis value = -1.2945. See robustness and outliers.	Anderson Darling P-Value = 0.000. Reject null hypothesis: "data are sampled from a normal distribution," so conclude that the assumption of normality is violated (at 95% confidence level). Skewness value = -1.2693 and Kurtosis value = 0.5291. See robustness and outliers.	Anderson Darling P-Value = 0.019. Reject null hypothesis: "data are sampled from a normal distribution," so conclude that the assumption of normality is violated (at 95% confidence level). Skewness value = -0.9436 and Kurtosis value = 0.1684. See robustness and outliers.
Robustness:	Minimum sample size for a robust ANOVA test = 4. Since each sample size is greater than this, the ANOVA test is robust to nonnormality.		
Outliers (Boxplot Rules):	No outliers found.	No outliers found.	No outliers found.
Randomness (Independence):	Nonparametric Runs Test (Exact) P-Value = 0.465. Fail to reject null hypothesis: "data are random," so conclude that the assumption of randomness (serial independence) is not violated.	Nonparametric Runs Test (Exact) P-Value = 0.158. Fail to reject null hypothesis: "data are random," so conclude that the assumption of randomness (serial independence) is not violated.	Nonparametric Runs Test (Exact) P-Value = 0.568. Fail to reject null hypothesis: "data are random," so conclude that the assumption of randomness (serial independence) is not violated.
Equal Variance:	Levene's Test for Equal Variances P-Value = 0.014. Reject null hypothesis: "variances are equal," so conclude that the assumption of equal variances (or standard deviations) is violated (at 95% confidence level). Welch's ANOVA is appropriate.		

This is a text report with color highlight: Green (OK), Yellow (Warning) and Red (Serious Violation).

Each sample is tested for Normality using the Anderson-Darling test. If not normal, the minimum sample size for robustness of the ANOVA Test is determined utilizing Monte Carlo regression equations (see [Basic Statistical Templates – Minimum Sample Size for Robust t-Tests and ANOVA](#)). If the sample size is inadequate, a warning is given and the appropriate Nonparametric test is recommended (Kruskal-Wallis if there are no extreme outliers, Mood’s Median if there are extreme outliers).

Each sample is tested for Outliers defined as: Potential: Tukey's Boxplot ($> Q3 + 1.5 \cdot IQR$ or $< Q1 - 1.5 \cdot IQR$); Likely: Tukey's Boxplot $2.2 \cdot IQR$; Extreme: Tukey's Boxplot $3 \cdot IQR$. If outliers are present, a warning is given and recommendation to review the data with a Boxplot and Normal Probability Plot

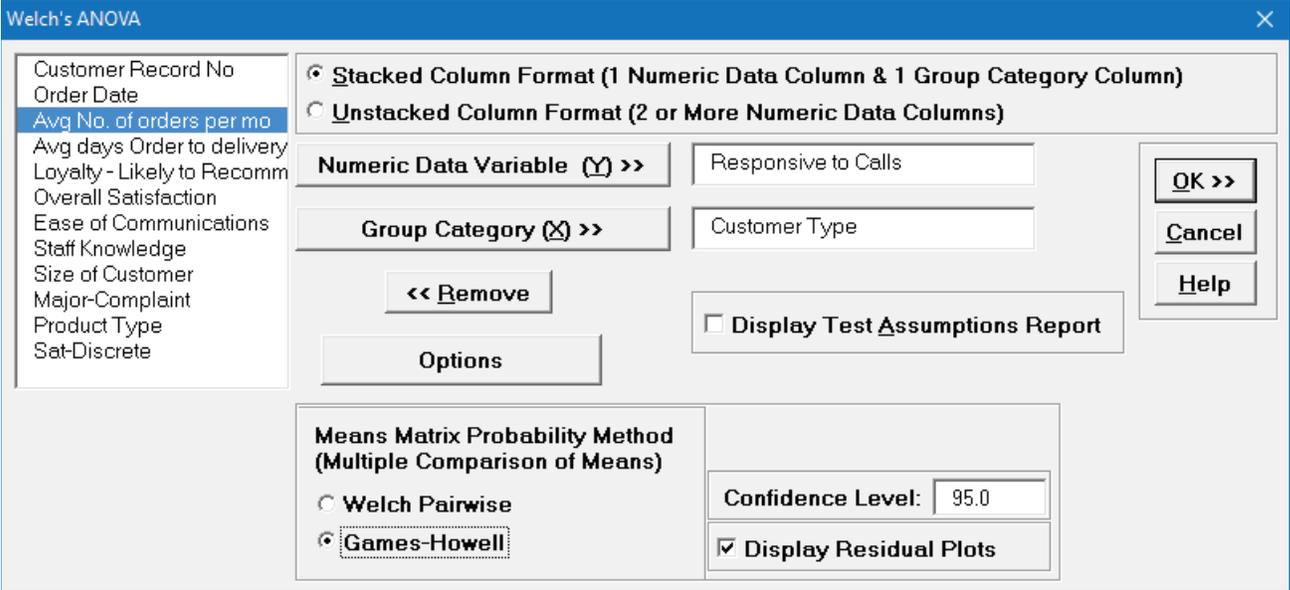
Tip: If the removal of outlier(s) result in an Anderson-Darling P-Value that is > 0.1 , a notice is given that excluding the outlier(s), the sample data are inherently normal.

Each sample is tested for Randomness using the Exact Nonparametric Runs Test. If the sample data is not random, a warning is given and recommendation to review the data with a Run Chart.

A test for Equal Variances is also applied. If all sample data are normal, Bartlett’s Test is utilized, otherwise Levene’s Test is used. Since we are using Welch’s ANOVA, it is confirmed as appropriate.

See Appendix [Hypothesis Test Assumptions Report](#) for further details.

11. Press **F3** or click **Recall SigmaXL Dialog** to recall last dialog. Uncheck **Display Test Assumptions Report**. Click the **Options** button. Select **Games-Howell**. Check **Display Residual Charts** as shown:



Note: The **Confidence Level** is used to set the level in the Mean/Confidence Interval Plot and the alpha level ($\alpha = (100 - CI)/100$) used to highlight the P-Values. However, the confidence level used in the Residuals Normal Probability Plot is always 95%.

12. Click **OK**. The Pairwise Means Difference (Means Matrix) and Games-Howell Probability results are:

Pairwise Mean Difference (row - column)	1	2	3
1	0	-0.814578	-0.409869
2		0	0.404709
3			0
Games-Howell Probabilities	1	2	3
1		0.0123	0.3998
2			0.2572
3			

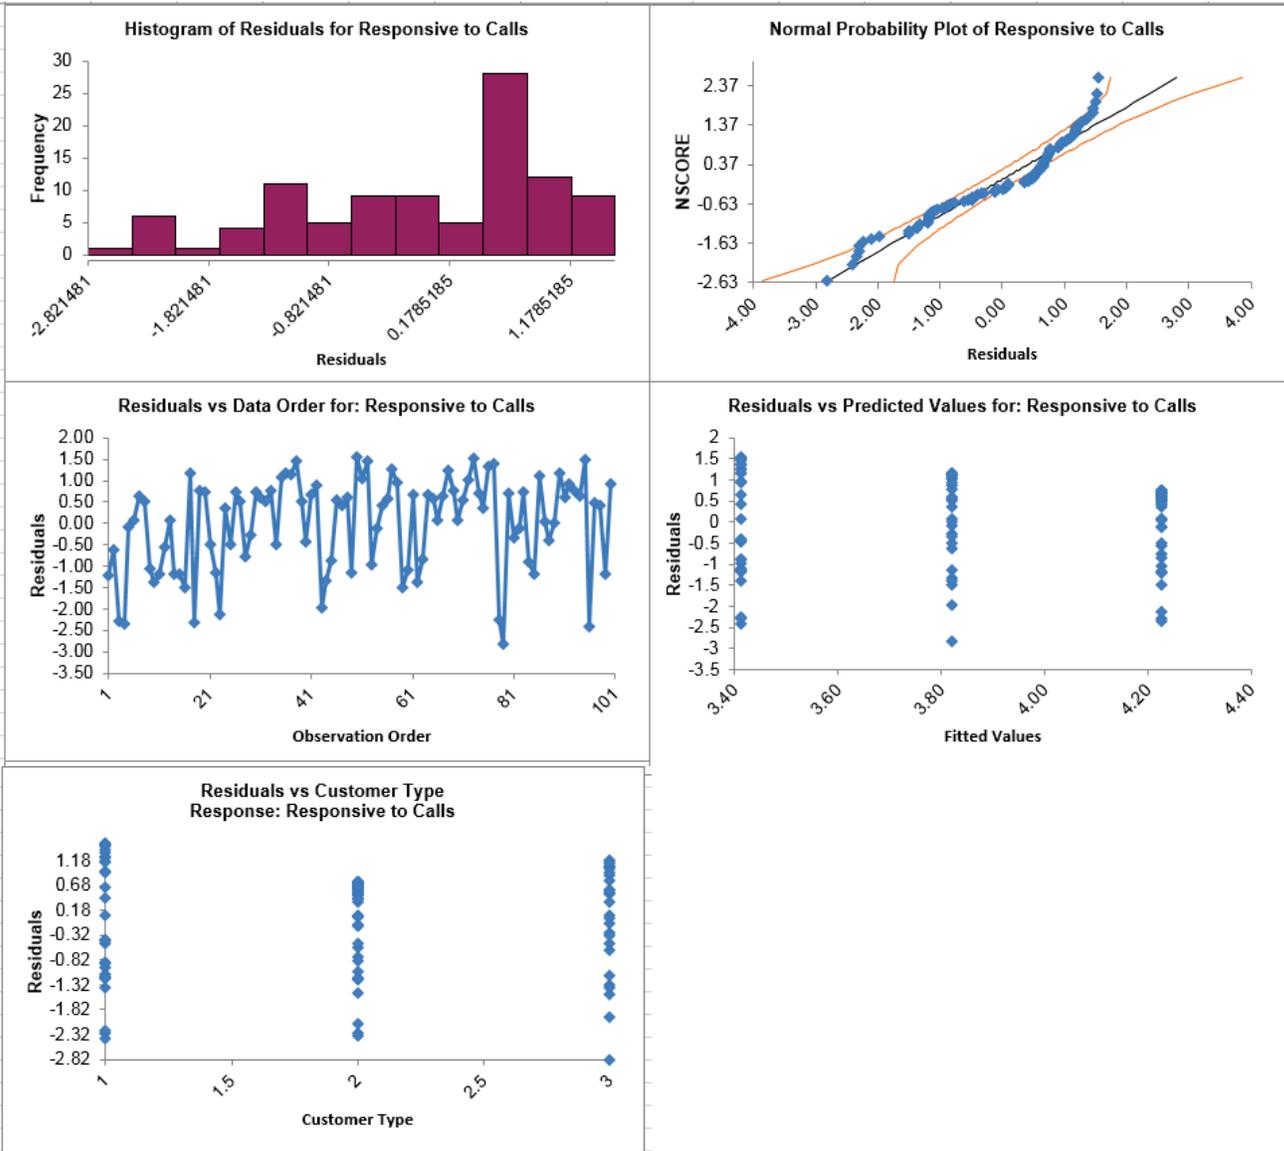
13. The significant Games-Howell Probability values above have not changed as compared to Welch Pairwise, but note that they are larger to compensate for the family-wise error rate:

Pairwise Mean Difference (row - column)	1	2	3
1	0	-0.814578	-0.409869
2		0	0.404709
3			0

Welch Pairwise Probabilities	1	2	3
1		0.0045	0.1981
2			0.1170
3			

Note also that the 1 – 2 Games-Howell probability is smaller than the Bonferroni corrected value = $.0045 * 3 = .0135$, so more powerful than Bonferroni. The difference in power between Games-Howell and Bonferroni becomes more prominent with a larger number of groups, so Bonferroni is not included as an option.

14. Click the **Welch Residuals** sheet tab to display the Residual Plots:



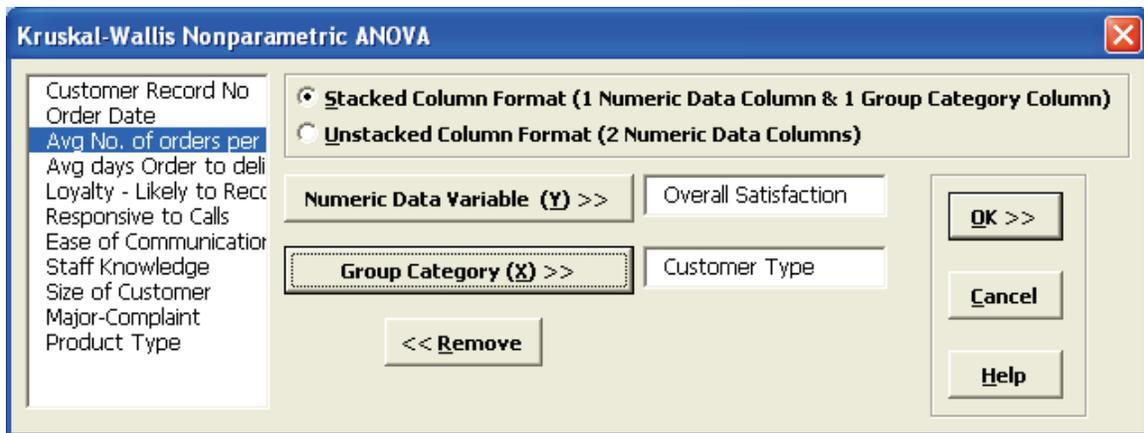
Residuals are the unexplained variation from the ANOVA model (Actual – Predicted or Fitted values). Note that the residuals are not normally distributed - as expected from the assumptions report - but like the regular ANOVA, Welch’s ANOVA is quite robust to the assumption of normality. Also, as expected from Levene’s test for equal variances, the variability for Type 2 is less than Type 1, but Welch’s ANOVA is robust to the assumption of equal variances, so we can trust that the P-Values are valid.

Part M – Nonparametric Multiple Comparison

Kruskal-Wallis

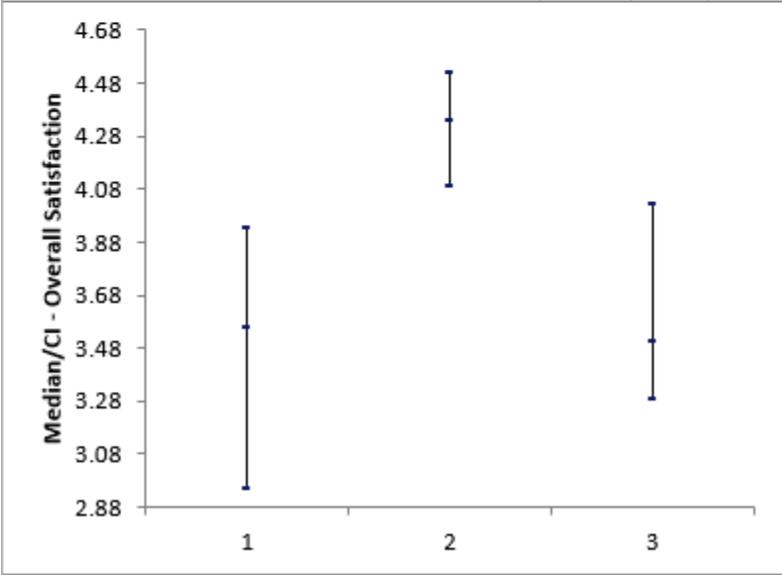
The Kruskal-Wallis test is an extension of the Mann-Whitney Rank test, allowing for more than 2 samples. It is a nonparametric equivalent to the parametric One-Way ANOVA. The Null Hypothesis is: $H_0: \text{Median}_1 = \text{Median}_2 = \dots = \text{Median}_K$. H_a : At least two Medians are different. The test does not assume sample normality but does assume that the samples have equal shapes. If the shapes are different, the null hypothesis is that the distributions are the same.

1. Open **Customer Data.xlsx**, click on **Sheet 1** tab (or press **F4** to activate last worksheet).
2. Click **SigmaXL > Statistical Tools > Nonparametric Tests > Kruskal-Wallis**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**.
3. Click **Next**. Ensure that **Stacked Column Format** is checked. Select *Overall Satisfaction*, click **Numeric Data Variable (Y) >>**; select *Customer Type*, click **Group Category (X) >>**.



4. Click **OK**. The results are shown below:

Kruskal-Wallis Nonparametric ANOVA: Overall Satisfacti			
Test Information			
H ₀ : Median 1 = Median 2 = ... = Median k			
H _a : At least one pair Median i ≠ Median j			
Customer Type	1	2	3
Count (N)	31	42	27
Median	3.56	4.34	3.51
UC Median (2-sided, 95%)	3.936	4.518	4.023
LC Median (2-sided, 95%)	2.954	4.095	3.289
Z	-3.339	4.529	-1.557
Kruskal-Wallis Statistic (H)	21.360		
DF	2		
P-Value (2-sided, adjusted for ties)	0.0000		



Click on cell B16 to view the P-Value with more decimal place precision (or change the cell format to scientific notation). The P-Value of 0.000023 (2.3 e-5) tells us that we reject H₀. At least one pairwise set of medians are not equal.

5. The Kruskal-Wallis Statistic is based on comparing mean ranks for each group versus the mean rank for all observations. The Z value for Customer Type 3 is -1.56, the smallest absolute Z-value. This size indicates that the mean rank for Type 3 differed least from the mean rank for all observations. The Z value for Customer Type 2 is 4.53, the largest absolute Z-value. This size indicates that the mean rank differed most from the mean rank for all observations.

6. A graphical view of the Overall Satisfaction Median and 95% Confidence Intervals are given to complement the Z scores. The fact that the CI's for Customer Type 2 do not overlap those of Type 1 or 3, clearly shows that Customer Type 2 has a significantly higher median satisfaction score. The overlap of CI's for Type 1 and 3 shows that the median satisfaction scores for 1 and 3 are not significantly different.

Kruskal-Wallis – Exact

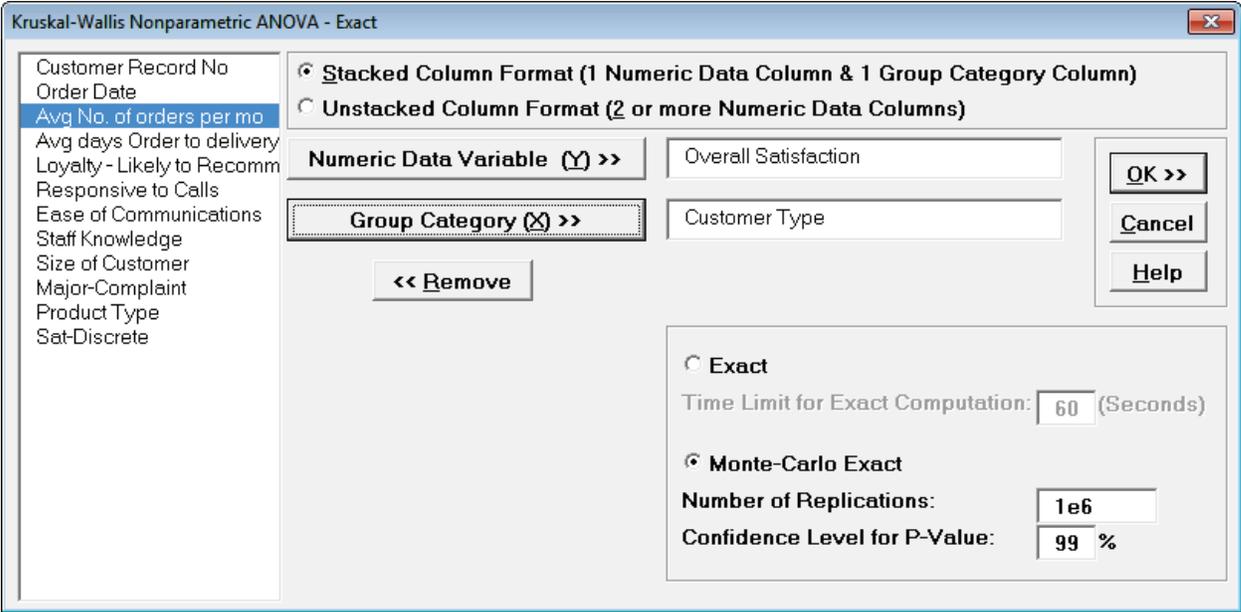
We will now redo the above example to estimate the exact P-Value using Monte Carlo. Typically, this would not be necessary unless the sample sizes were smaller (each sample $N \leq 5$ for Kruskal-Wallis), but this gives us continuity on the example. We will consider a small sample problem later.

Computing an exact P-Value for Kruskal-Wallis is very computationally intensive. The Network Model by Mehta and Patel cannot be used for this test (see Appendix [Exact and Monte Carlo P-Values for Nonparametric and Contingency Tests](#)). In this example, the total number of permutations are:

$$(31+42+27)! / (31! * 42! * 27!) = 7.42 \text{ E}44$$

(i.e., more than the number of stars in the observable universe). So, we will not attempt to compute the exact, but rather use Monte Carlo!

1. Open **Customer Data.xlsx**, click **Sheet 1** tab (or press **F4** to activate last worksheet).
2. Click **SigmaXL > Statistical Tools > Nonparametric Tests – Exact > Kuskal-Wallis – Exact**. If necessary, check **Use Entire Data Table**, click **Next**.
3. Ensure that **Stacked Column Format** is checked. Select *Overall Satisfaction*, click **Numeric Data Variable (Y) >>**; select *Customer Type*, click **Group Category (X) >>**. Select **Monte Carlo Exact** with the **Number of Replications** = 1e6 and **Confidence Level for P-Value** = 99%. One million replications are used because the expected P-Value is very small as estimated from the “large sample” Kruskal-Wallis above. This will take up to a minute to run, so if you have a slow computer, use 1e5 replications instead of 1e6.



Tip: The Monte Carlo 99% confidence interval for P-Value is **not** the same as a confidence interval on the test statistic due to data sampling error. The confidence level for the hypothesis test statistic is still 95%, so **all reported P-Values less than .05 will be highlighted in red** to indicate significance. The 99% Monte Carlo P-Value confidence interval is due to the uncertainty in Monte Carlo sampling, and it becomes smaller as the number of replications increases (irrespective of the data sample size). The Exact P-Value will lie within the stated Monte Carlo confidence interval 99% of the time.

4. Click **OK**.

Kruskal-Wallis Nonparametric ANOVA - Monte Carlo: Overall Satisfaction			
Test Information			
H ₀ : Median 1 = Median 2 = ... = Median k			
H _a : At least one pair Median i ≠ Median j			
Customer Type	1	2	3
Count (N)	31	42	27
Median	3.56	4.34	3.51
UC Median (2-sided, 95% approx.)	3.936	4.518	4.023
LC Median (2-sided, 95% approx.)	2.954	4.095	3.289
Z	-3.339	4.529	-1.557
Kruskal-Wallis Statistic (H)	21.360		
DF	2		
Monte Carlo P-Value (2-sided)	0.0000		
Monte Carlo P-Value 99% CI Upper	0.0000		
Monte Carlo P-Value 99% CI Lower	0.0000		

Click on cell B16 to view the P-Value with more decimal place precision (or change the cell format to scientific notation). The Monte Carlo P-Value here is 0.000009 (9 e-6) with a 99% confidence interval of .000002 (2 e-6) to 0.000016 (1.6 e-5). This will be slightly different every time it is run (the Monte Carlo seed value is derived from the system clock). So, we reject H₀: at least one pairwise set of medians are not equal.

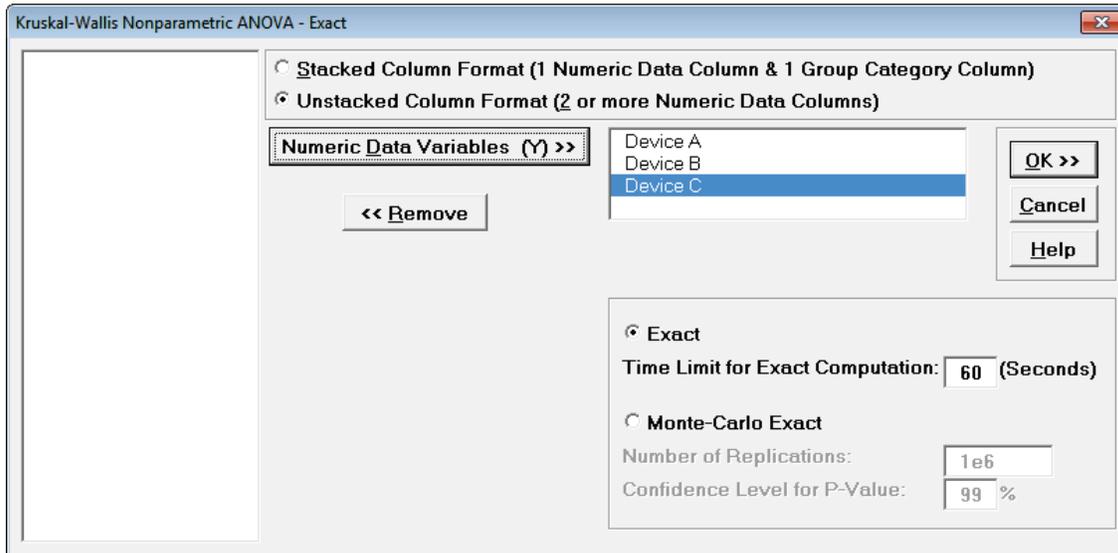
Note that the large sample (asymptotic) P-Value of 2.3 e-5 lies outside of the Monte Carlo exact confidence interval.

5. Now we will consider a small sample problem. Open **Snore Study.xlsx**. This data is from:

Gibbons, J.D. and Chakraborti, S. (2010). *Nonparametric Statistical Inference* (5th Edition). New York: Chapman & Hall, (Example 10.2.1 data, page 347; Example 10.4.2 analysis, pp. 360 – 362).

An experiment was conducted to determine which device is the most effective in stopping snoring or at least in reducing it. Fifteen men who are habitual snorers were divided randomly into three groups to test the devices. Each man’s sleep was monitored for one night by a machine that measures amount of snoring on a 100-point scale while using a device.

- Select **Snore Study Data** tab. Click **SigmaXL > Statistical Tools > Nonparametric Tests – Exact > Kruskal-Wallis - Exact**. If necessary, check **Use Entire Data Table**, click **Next**.
- With **Unstacked Column Format** checked, select *Device A*, *Device B* and *Device C*, click **Numeric Data Variables (Y) >>**. Select **Exact** with the default **Time Limit for Exact Computation = 60** seconds.



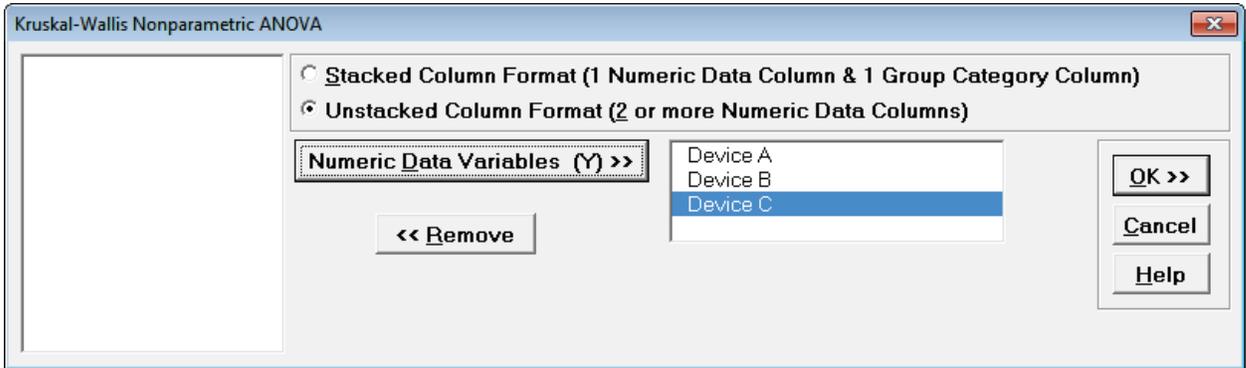
- Click **OK**. Results:

Kruskal-Wallis Nonparametric ANOVA - Exact			
Test Information			
H_0 : Median 1 = Median 2 = ... = Median k			
H_a : At least one pair Median $i \neq$ Median j			
Results:			
	Device A	Device B	Device C
Count (N)	5	5	5
Median	79	92	26
UC Median (2-sided, 95% approx.)	91	96	78
LC Median (2-sided, 95% approx.)	35	76	8
Z	0.244949	2.449	-2.694
Kruskal-Wallis Statistic (H)	8.880		
DF	2		
Exact P-Value (2-sided)	0.0042		

With the Exact P-Value = 0.0042 we reject H_0 , and conclude that there is a significant difference in median snore study scores. This exact P-Value matches that given in the reference textbook using SAS and StatXact.

By way of comparison we will now rerun the analysis using the “large sample” or “asymptotic” Kruskal-Wallis test.

9. Select **Snore Study Data** tab (or press **F4** to activate last worksheet). Click **SigmaXL > Statistical Tools > Nonparametric Tests > Kruskal-Wallis**. If necessary, check **Use Entire Data Table**, click **Next**.
10. With **Unstacked Column Format** checked, select *Device A*, *Device B* and *Device C*, click **Numeric Data Variables (Y) >>**.



11. Click **OK**. Results:

Kruskal-Wallis Nonparametric ANOVA			
Test Information			
H_0 : Median 1 = Median 2 = ... = Median k			
H_a : At least one pair Median $i \neq$ Median j			
Results:			
	Device A	Device B	Device C
Count (N)	5	5	5
Median	79	92	26
UC Median (2-sided, 95%)	91	96	78
LC Median (2-sided, 95%)	35	76	8
Z	0.244949	2.449	-2.694
Kruskal-Wallis Statistic (H)	8.880		
DF	2		
P-Value (2-sided, adjusted for ties)	0.0118		

With the P-Value = .0118 we reject H_0 (using alpha = .05), but note that if we were using alpha = 0.01, we would have incorrectly failed to reject the null hypothesis. This “large sample” P-Value matches that given in the reference textbook using Minitab.

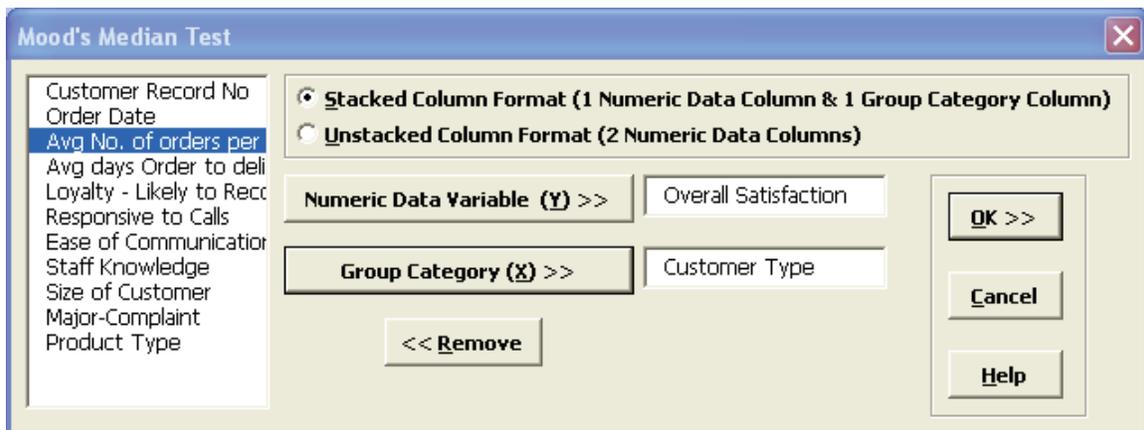
In conclusion, whenever you have a small sample size and are performing a nonparametric test, always use the Exact option.

Mood's Median Test

Mood's Median Test is an extension of the One Sample Sign Test, using Chi-Square as the test statistic. Like the Kruskal-Wallis test, Mood's median test can be used to test the equality of medians from multiple samples. It provides a nonparametric alternative to the one-way analysis of variance. The Null Hypothesis is: $H_0: \text{Median}_1 = \text{Median}_2 = \dots = \text{Median}_K$. H_a : At least two Medians are different.

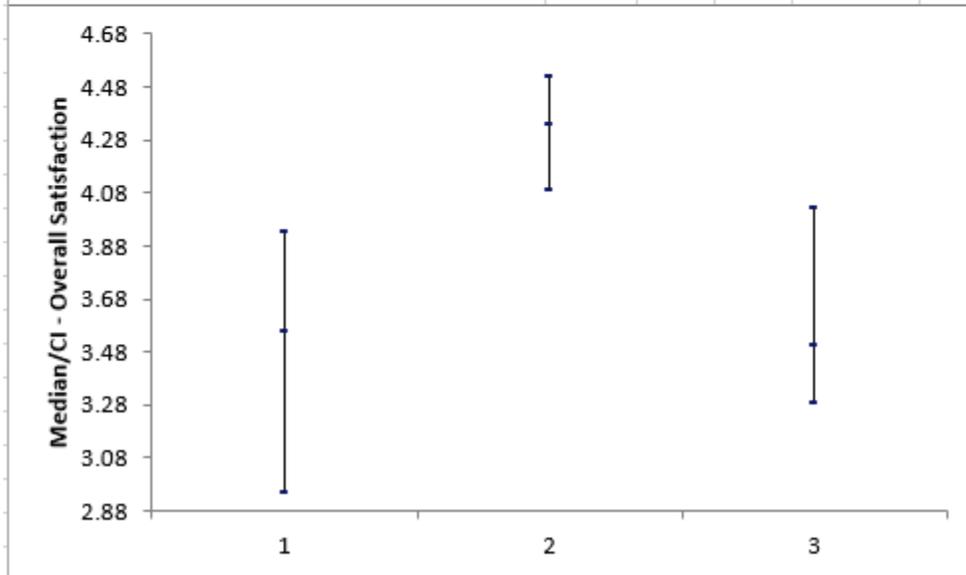
Mood's median test is more robust to outliers than the Kruskal-Wallis test, but is less powerful in the absence of outliers. You should first look at your data with Boxplots. If there are extreme outliers, then Mood's Median should be used rather than Kruskal-Wallis.

1. Open **Customer Data.xlsx**, click on **Sheet 1** tab (or press **F4** to activate last worksheet).
2. Click **SigmaXL > Statistical Tools > Nonparametric Tests > Mood's Median Test**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**.
3. Click **Next**. Ensure that **Stacked Column Format** is checked. Select *Overall Satisfaction*, click **Numeric Data Variable (Y) >>**; select *Customer Type*, click **Group Category (X) >>**.



4. Click **OK**. The results are shown below:

Mood's Median Test: Overall Satisfaction			
Test Information			
H ₀ : Median 1 = Median 2 = ... = Median k			
H _a : At least one pair Median i ≠ Median j			
Customer Type	1	2	3
Count (N ≤ Overall Median)	21	12	17
Count (N > Overall Median)	10	30	10
Median	3.560	4.340	3.510
UC Median (2-sided, 95%)	3.936	4.518	4.023
LC Median (2-sided, 95%)	2.954	4.095	3.289
Overall Median	3.945		
Chi-Square	13.432		
DF	2		
P-Value (2-sided)	0.0012		



The P-Value of 0.0012 tells us that we reject H₀. At least one pairwise set of medians are not equal.

- A graphical view of the Overall Satisfaction Median and 95% Confidence Intervals are given. This is the same graph provided in the Kruskal-Wallis test report.

Mood's Median Test - Exact

We will now redo the above example to compute exact P-Values. Typically, this would not be necessary unless the sample sizes were smaller (each sample $N \leq 10$ for Mood's Median), but this gives us continuity on the example. We will consider a small sample problem later. Unlike Kruskal-Wallis, computing an exact P-Value is efficient due to the reduction of the data to a Chi-Square table and the use of a Network Model (see Appendix [Exact and Monte Carlo P-Values for Nonparametric and Contingency Tests](#))

1. Open **Customer Data.xlsx**, click **Sheet 1** tab (or press **F4** to activate last worksheet).
2. Click **SigmaXL > Statistical Tools > Nonparametric Tests – Exact > Mood's Median Test – Exact**. If necessary, check **Use Entire Data Table**, click **Next**.
3. With **Stacked Column Format** checked, select *Overall Satisfaction*, click **Numeric Data Variable (Y) >>**; select *Customer Type*, click **Group Category (X) >>**. Select **Exact** with the default **Time Limit for Exact Computation = 60 seconds**. Set **Values Equal to Overall Median: to Counted as Below**.

The screenshot shows the 'Mood's Median Test - Exact' dialog box. On the left is a list of variables: Customer Record No, Order Date, Avg No. of orders per mo, Avg days Order to delivery, Loyalty - Likely to Recomm, Responsive to Calls, Ease of Communications, Staff Knowledge, Size of Customer, Major-Complaint, Product Type, and Sat-Discrete. The 'Stacked Column Format (1 Numeric Data Column & 1 Group Category Column)' option is selected. The 'Numeric Data Variable (Y) >>' field contains 'Overall Satisfaction' and the 'Group Category (X) >>' field contains 'Customer Type'. A '<< Remove' button is below the group category field. In the 'Values Equal to Overall Median:' section, 'Counted as Below' is selected. The 'Exact' method is selected, with a 'Time Limit for Exact Computation' of 60 (Seconds). The 'Monte-Carlo Exact' method is also visible, with 'Number of Replications' set to 10000 and 'Confidence Level for P-Value' set to 99 %.

Tip: If none of the observations are equal to the overall median, the options for “**Values Equal to Overall Median:**” do not affect the exact P-Value. That is the case with this Customer Data example. In cases where there are observations equal to the overall median, then the exact P-Value will be different for each option selected. For regular Mood's Median, SigmaXL uses **Counted as Below** (in agreement with Minitab) but will use **Counted as Above** for small sample cases if the average cell expected value is higher. StatXact uses **Counted as Below** for Mood's Median. Matlab uses **Not Counted**. Users should try **Counted as Below** and **Counted as Above** to ensure that the P-Values agree on reject or fail to reject H_0 . Use of this option will be demonstrated in the small sample example.

Tip: If the exact computation time limit is exceeded a dialog will prompt you to use Monte Carlo or to increase the computation time. When this occurs, Monte Carlo is recommended.

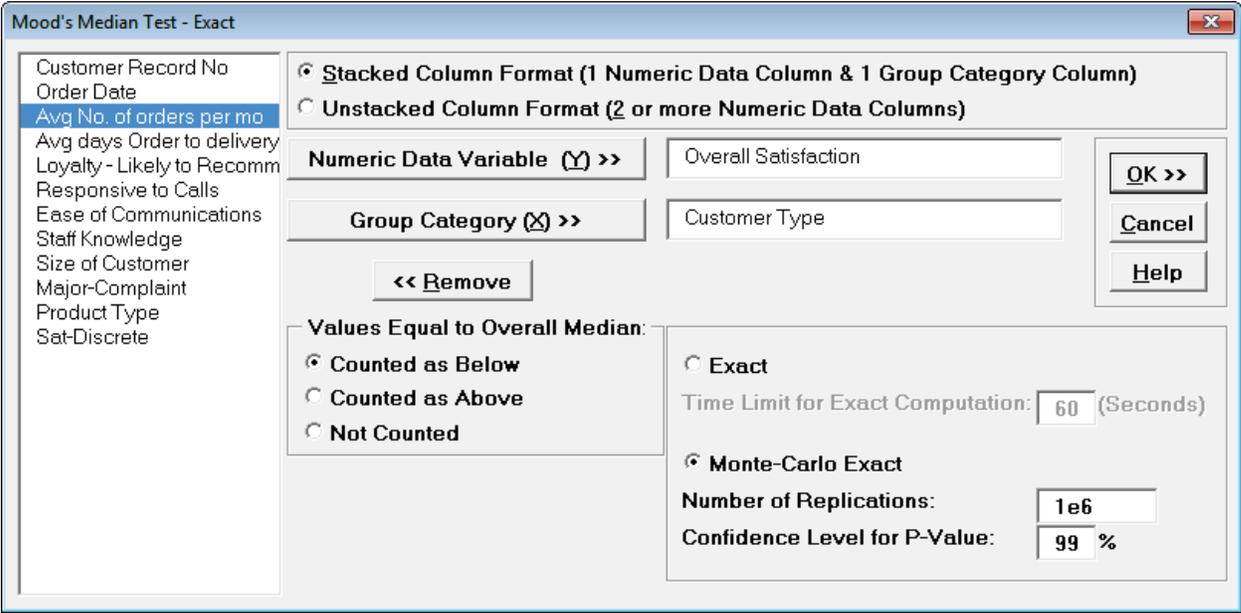
4. Click **OK**. Resulting output:

Mood's Median Test - Exact: Overall Satisfaction			
Test Information			
H ₀ : Median 1 = Median 2 = ... = Median k			
H _a : At least one pair Median i ≠ Median j			
Customer Type	1	2	3
Count (N ≤ Overall Median)	21	12	17
Count (N > Overall Median)	10	30	10
Median	3.560	4.340	3.510
UC Median (2-sided, 95% approx.)	3.936	4.518	4.023
LC Median (2-sided, 95% approx.)	2.954	4.095	3.289
Overall Median	3.945		
Chi-Square	13.432		
DF	2		
Exact P-Value (2-sided)	0.0012		

Click on cell **B17** to view the Exact P-Value with more decimal place precision. Given the P-Value of .00124 we reject H₀ and conclude that at least one pairwise set of medians are not equal. The Exact P-Value is very close to the above “large sample” or “asymptotic” result (.00121). This was expected because the sample size is reasonable (N > 10), so the “large sample” Mood’s Median P-Values are valid using a chi-square approximation. Note also that the counts in **B8:D9** form a two-way contingency table and the expected counts are all greater than 5.

The Exact P-Value was computed in seconds, but if the data set was larger, the required computation time could become excessive, and Monte Carlo would be required.

5. Press **F3** or click **Recall SigmaXL Dialog** to recall last dialog. Select **Monte Carlo Exact** with the **Number of Replications = 1e6** and **Confidence Level for P-Value = 99%**. One million replications are used because the expected P-Value is small as estimated from the “large sample” Mood’s Median above. This will take up to a minute to run, so if you have a slow computer, use 1e5 replications instead of 1e6.



Tip: The Monte Carlo 99% confidence interval for P-Value is **not** the same as a confidence interval on the test statistic due to data sampling error. The confidence level for the hypothesis test statistic is still 95%, so **all reported P-Values less than .05 will be highlighted in red** to indicate significance. The 99% Monte Carlo P-Value confidence interval is due to the uncertainty in Monte Carlo sampling, and it becomes smaller as the number of replications increases (irrespective of the data sample size). The Exact P-Value will lie within the stated Monte Carlo confidence interval 99% of the time.

6. Click **OK**.

Mood's Median - Monte Carlo: Overall Satisfaction			
Test Information			
H ₀ : Median 1 = Median 2 = ... = Median k			
H _a : At least one pair Median i ≠ Median j			
Customer Type	1	2	3
Count (N ≤ Overall Median)	21	12	17
Count (N > Overall Median)	10	30	10
Median	3.560	4.340	3.510
UC Median (2-sided, 95% approx.)	3.936	4.518	4.023
LC Median (2-sided, 95% approx.)	2.954	4.095	3.289
Overall Median	3.945		
Chi-Square	13.432		
DF	2		
Monte Carlo P-Value (2-sided)	0.0012		
Monte Carlo P-Value 99% CI Upper	0.0013		
Monte Carlo P-Value 99% CI Lower	0.0012		

Click on cell **B17** to view the Monte Carlo P-Value with more decimal place precision. The Monte Carlo P-Value here is .001249 with a 99% confidence interval of .00117 (**B19**) to .00133 (**B18**). This will be slightly different every time it is run (the Monte Carlo seed value is derived from the system clock). The true Exact P-Value = .00124 lies within this confidence interval.

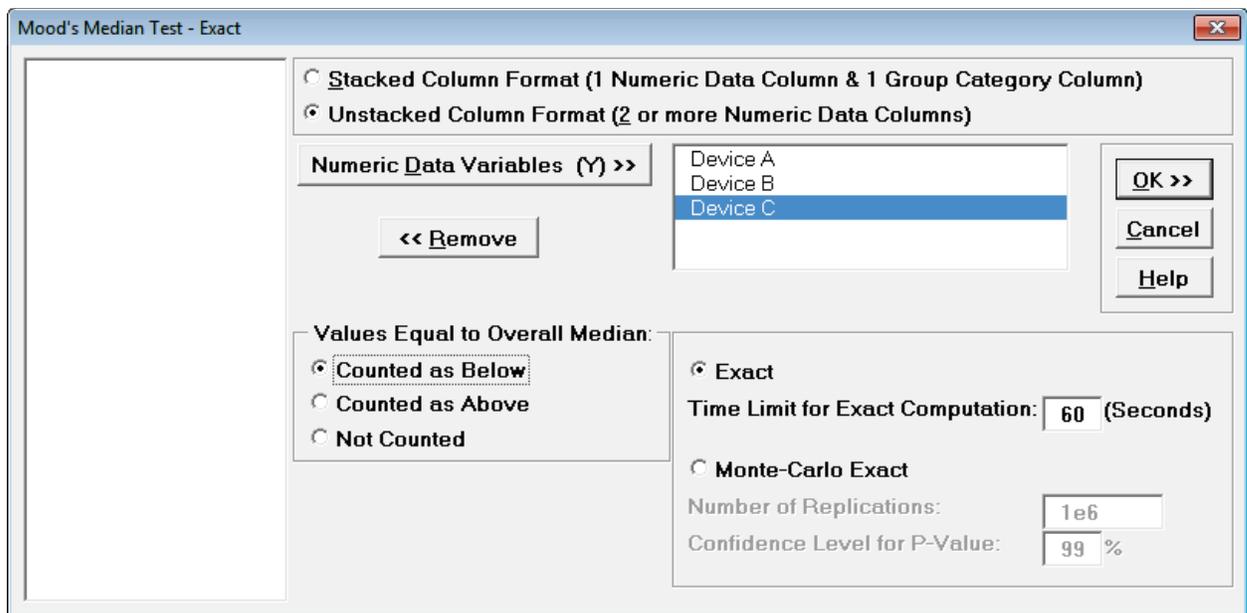
So, we reject H_0 : at least one pairwise set of medians are not equal.

7. Now we will consider the small sample problem used in Kruskal-Wallis. Open **Snore Study.xlsx**. This data is from:

Gibbons, J.D. and Chakraborti, S. (2010). *Nonparametric Statistical Inference* (5th Edition). New York: Chapman & Hall, (Example 10.2.1, pp. 347-348).

An experiment was conducted to determine which device is the most effective in stopping snoring or at least in reducing it. Fifteen men who are habitual snorers were divided randomly into three groups to test the devices. Each man's sleep was monitored for one night by a machine that measures amount of snoring on a 100-point scale while using a device.

8. Select **Snore Study Data** tab. Click **SigmaXL > Statistical Tools > Nonparametric Tests – Exact > Mood's Median Test - Exact**. If necessary, check **Use Entire Data Table**, click **Next**.
9. With **Unstacked Column Format** checked, select *Device A*, *Device B* and *Device C*, click **Numeric Data Variables (Y) >>**. Select **Exact** with the default **Time Limit for Exact Computation = 60** seconds.



10. Click **OK**. Results:

Mood's Median Test - Exact			
Test Information			
H ₀ : Median 1 = Median 2 = ... = Median k			
H _a : At least one pair Median i ≠ Median j			
Results:			
	Device A	Device B	Device C
Count (N ≤ Overall Median)	2	1	5
Count (N > Overall Median)	3	4	0
Median	79	92	26
UC Median (2-sided, 95% approx.)	91	96	78
LC Median (2-sided, 95% approx.)	35	76	8
Overall Median	78		
Chi-Square	6.964		
DF	2		
Exact P-Value (2-sided)	0.0676		

With the Exact P-Value = .0676 we fail to reject H₀, and cannot conclude that there is a significant difference in median snore study scores. This exact P-Value matches that given in the reference textbook using StatXact. Comparing Mood's Median to Kruskal-Wallis we see that the Exact P-Value is higher because Mood's Median is not as powerful as Kruskal-Wallis.

We will now rerun the analysis using the **Counted as Above** option.

- Press **F3** or click **Recall SigmaXL Dialog** to recall last dialog. Set **Values Equal to Overall Median:** to **Counted as Above**.

Mood's Median Test - Exact

Stacked Column Format (1 Numeric Data Column & 1 Group Category Column)
 Unstacked Column Format (2 or more Numeric Data Columns)

Numeric Data Variables (Y) >>

Device A
Device B
Device C

<< Remove

OK >>

Cancel

Help

Values Equal to Overall Median:

Counted as Below
 Counted as Above
 Not Counted

Exact
 Time Limit for Exact Computation: 60 (Seconds)

Monte-Carlo Exact
 Number of Replications: 1e6
 Confidence Level for P-Value: 99 %

12. Click **OK**. Results:

Mood's Median Test - Exact			
Test Information			
H_0 : Median 1 = Median 2 = ... = Median k			
H_a : At least one pair Median i \neq Median j			
Results:	Device A	Device B	Device C
Count (N < Overall Median)	2	1	4
Count (N \geq Overall Median)	3	4	1
Median	79	92	26
UC Median (2-sided, 95% approx.)	91	96	78
LC Median (2-sided, 95% approx.)	35	76	8
Overall Median	78		
Chi-Square	3.750		
DF	2		
Exact P-Value (2-sided)	0.3007		

With the Exact P-Value = 0.3007 we fail to reject H_0 , and cannot conclude that there is a significant difference in median snore study scores. Note that setting **Values Equal to Overall Median to Counted as Above** resulted in a different count for Device C, which results in the dramatic difference in Exact P-Value.

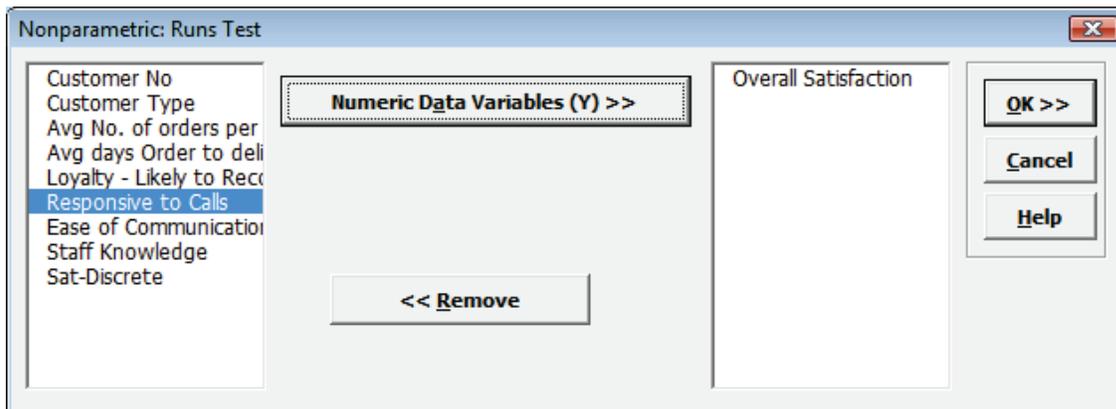
In conclusion, when using Mood's Median Exact, always try **Counted as Below** and **Counted as Above** to ensure that the P-Values agree with each other.

Part N – Nonparametric Runs Test

Nonparametric Runs Test for Randomness

The nonparametric runs test provides a test for randomness or independence. The null hypothesis is H_0 : The data is random (or independent). The alternative hypothesis is H_a : The data is not random (or independent). Note that this test is also provided as an option in Run Charts (**SigmaXL > Graphical Tools > Run Chart**). In addition to providing an overall test for randomness, 4 tests are performed to detect Clustering, Mixtures, Trends, and Oscillations. If any of these patterns are significant (typically using $\alpha = 0.01$), we would need to take corrective action before proceeding with further statistical analysis. (Note that SigmaXL will highlight any P-Values < .05 in red.)

1. Open **Customer Data.xlsx**, click on **Sheet 1** tab (or press **F4** to activate last worksheet).
2. Click **SigmaXL > Statistical Tools > Nonparametric Tests > Runs Test**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**.
3. Click **Next**. Select *Overall Satisfaction*, click **Numeric Data Variable (Y) >>**:



4. Click **OK**. The resulting report is shown:

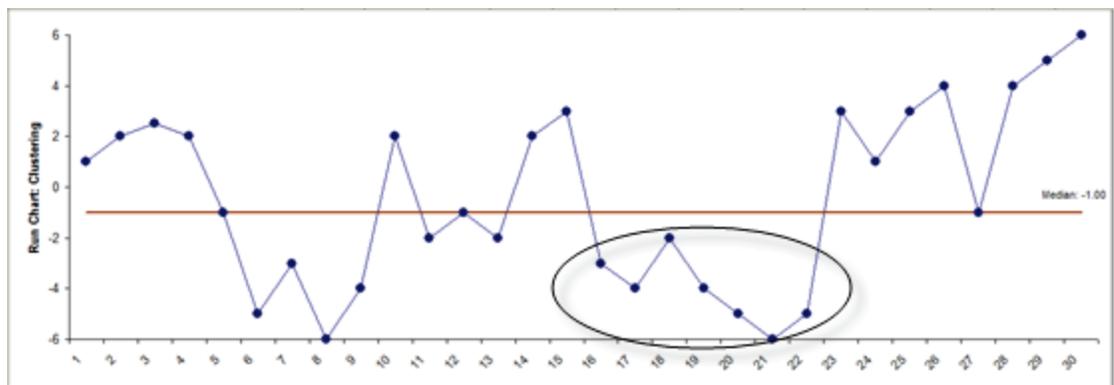
Nonparametric Runs Test	
Test Information	Overall Satisfaction
Number of Runs about Median:	44
Expected Number of Runs about Median:	51
Number of Points above Median:	50
Number of Points equal to or below Median:	50
P-Value for Clustering:	0.0797
P-Value for Mixtures:	0.9203
P-Value for Lack of Randomness (2-Sided):	0.1594
Number of Runs Up or Down:	64
Expected Number of Runs Up or Down:	66.333
P-Value for Trends:	0.2883
P-Value for Oscillation:	0.7117

5. With all of the P-Values being greater than 0.01, we fail to reject H0, and conclude that the data is random (or statistically independent). Recall from the run chart of this data that there were no obvious trends or patterns.

Nonparametric Runs Test for Randomness - Examples

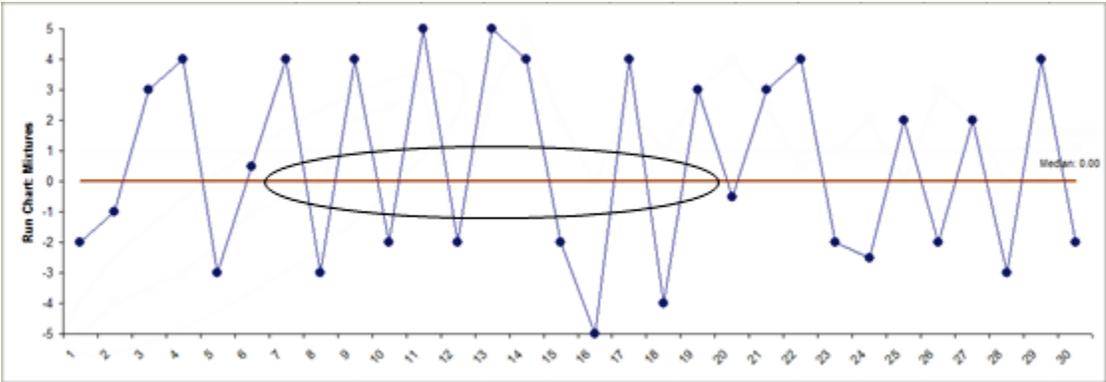
Examples of Clustering, Mixtures, Trends, and Oscillations are given below using the Run Chart to illustrate. The data for these examples are given in the file **Runs Test Example Data.xlsx**. (Use **SigmaXL > Graphical Tools > Run Charts** to create the run charts).

- a. **Clustering** appears as a group of points in one area of the chart. It may indicate special cause variation such as sampling or measurement problems.



Nonparametric Runs Test: Clustering	
Number of Runs about Median:	9
Expected Number of Runs about Median:	15.933
Number of Points above Median:	14
Number of Points equal to or below Median:	16
P-Value for Clustering:	0.0048
P-Value for Mixtures:	0.9952
P-Value for Lack of Randomness (2-Sided):	0.0096
Number of Runs Up or Down:	17
Expected Number of Runs Up or Down:	19.667
P-Value for Trends:	0.1168
P-Value for Oscillation:	0.8832

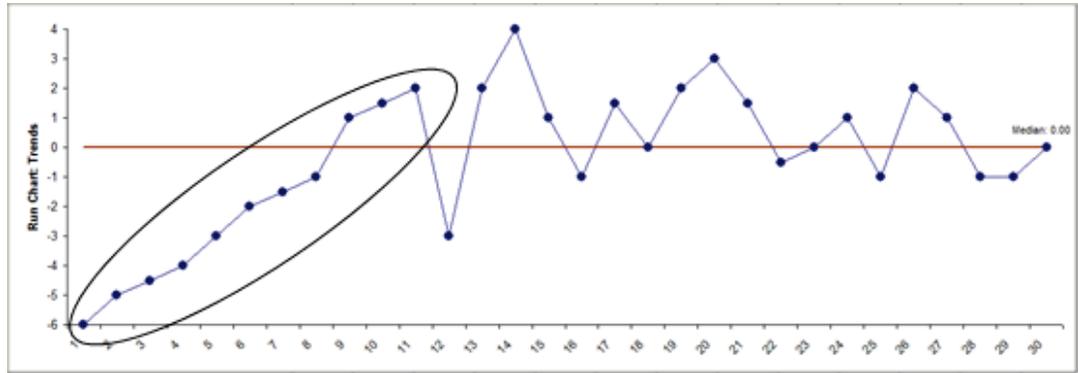
- b. **Mixtures** appear as an absence of data points near the center line. A mixture may indicate a bimodal distribution due to a regular change of shift, machinery, or raw materials.



Nonparametric Runs Test: Mixtures	
Number of Runs about Median:	23
Expected Number of Runs about Median:	16
Number of Points above Median:	15
Number of Points equal to or below Median:	15
P-Value for Clustering:	0.9954
P-Value for Mixtures:	0.0046
P-Value for Lack of Randomness (2-Sided):	0.0093
Number of Runs Up or Down:	22
Expected Number of Runs Up or Down:	19.667
P-Value for Trends:	0.8514
P-Value for Oscillation:	0.1486

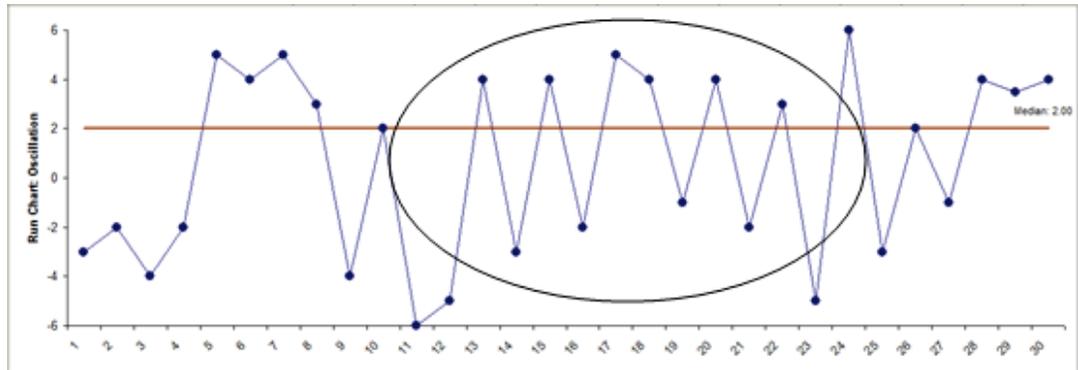
Note that the P-Value for Mixtures = 1 – P-Value for Clustering. They are mutually exclusive. The P-Value for Lack of Randomness = 2 * minimum of (P-Value Clustering, P-Value Mixtures).

- c. **Trends** appear as an upward or downward drift in the data and may be due to special causes such as tool wear.



Nonparametric Runs Test: Trends	
Number of Runs about Median:	13
Expected Number of Runs about Median:	15.733
Number of Points above Median:	13
Number of Points equal to or below Median:	17
P-Value for Clustering:	0.1504
P-Value for Mixtures:	0.8496
P-Value for Lack of Randomness (2-Sided):	0.3008
Number of Runs Up or Down:	13
Expected Number of Runs Up or Down:	19.667
P-Value for Trends:	0.0015
P-Value for Oscillation:	0.9985

d. **Oscillations** appear as rapid up/down fluctuations indicating process instability.



Nonparametric Runs Test: Oscillation	
Number of Runs about Median:	16
Expected Number of Runs about Median:	15.933
Number of Points above Median:	14
Number of Points equal to or below Median:	16
P-Value for Clustering:	0.5099
P-Value for Mixtures:	0.4901
P-Value for Lack of Randomness (2-Sided):	0.9801
Number of Runs Up or Down:	25
Expected Number of Runs Up or Down:	19.667
P-Value for Trends:	0.9914
P-Value for Oscillation:	0.0086

Note that the P-Value for Trends = 1 – P-Value for Oscillation. They are mutually exclusive.

Nonparametric Runs Test for Randomness - Exact

The Nonparametric Runs test does not assume that the sample data are normally distributed, but it does assume that the test statistic follows a Normal distribution when computing the “large sample” or “asymptotic” P-Value. With a small sample size ($N \leq 50$), this approximation may be invalid, so exact methods should be used. SigmaXL computes the exact P-Values utilizing permutations.

It is important to note that while exact P-Values are “correct,” they do not increase (or decrease) the power of a small sample test, so they are not a solution to the problem of failure to detect a trend due to inadequate sample size.

Clustering, Mixtures and Lack of Randomness (Runs Above/Below)

If Count (N) is greater than 1000, the exact P-Value is estimated using a continuity-corrected normal approximation. Since the Runs Test Exact P-Value is computed very quickly for sample sizes as large as 1000, Monte Carlo P-Values are not required.

Trends and Oscillation (Runs Up/Down)

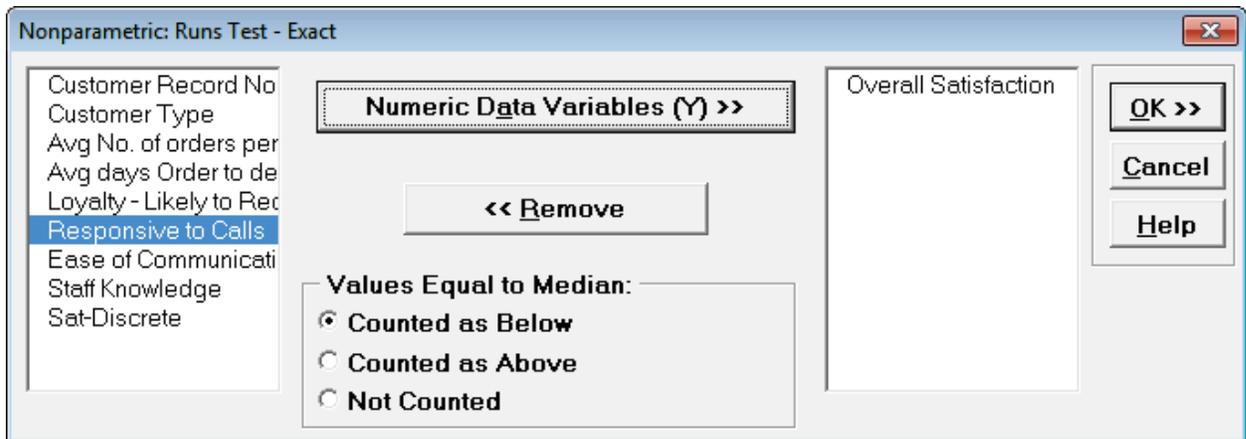
Exact P-Values are derived from published tables, given the sample size and the number of up/down runs. The exact tables apply to $N \leq 50$. If $N > 50$, a continuity-corrected normal approximation is used.

Exact P-Values are not available in Run Charts.

For further details and references refer to the Appendix [Exact and Monte Carlo P-Values for Nonparametric and Contingency Tests](#).

We will now redo the Customer Data Overall Satisfaction example to compute exact P-Values. Typically, this would not be necessary unless the sample sizes were smaller ($N \leq 50$), but this gives us continuity on the example. We will also consider the above small sample examples later.

1. Open **Customer Data.xlsx**, click **Sheet 1** tab.
2. Click **SigmaXL > Statistical Tools > Nonparametric Tests – Exact > Runs Test – Exact**. If necessary, check **Use Entire Data Table**, click **Next**.
3. Select *Overall Satisfaction*, click **Numeric Data Variables (Y) >>**. Set **Values Equal to Median: to Counted as Below**.



Note on Options for Clustering, Mixtures and Lack of Randomness (Runs Above/Below)

If none of the observations are equal to the sample median, the options for “**Values Equal to Median:**” do not affect the exact P-Value. That is the case with this Customer Data example. In cases where there are observations equal to the sample median, then the exact P-Value will be different for each option selected. For regular Runs Test, SigmaXL uses **Counted as Below** (in agreement with Minitab). StatXact uses **Counted as Above** for the Runs Test. Matlab uses **Not Counted**. Users should try **Counted as Below** and **Counted as Above** to ensure that the P-Values agree on reject or fail to reject H0.

Note on Options for Trends and Oscillation (Runs Up/Down)

The **Values Equal to Median** options also apply to the Runs Up/Down Test: zeros of the first order differences (i.e., two consecutive values that are the same). **Counted as Below** denotes “counted as negative.” **Counted as Above** denotes “counted as positive.” **Not Counted** denotes that zero first order differences are deleted. This consolidation of options was done to keep the dialog simple.

Missing values are ignored for all exact runs tests.

Use of these options will be demonstrated in the small sample example.

4. Click **OK**. Resulting output:

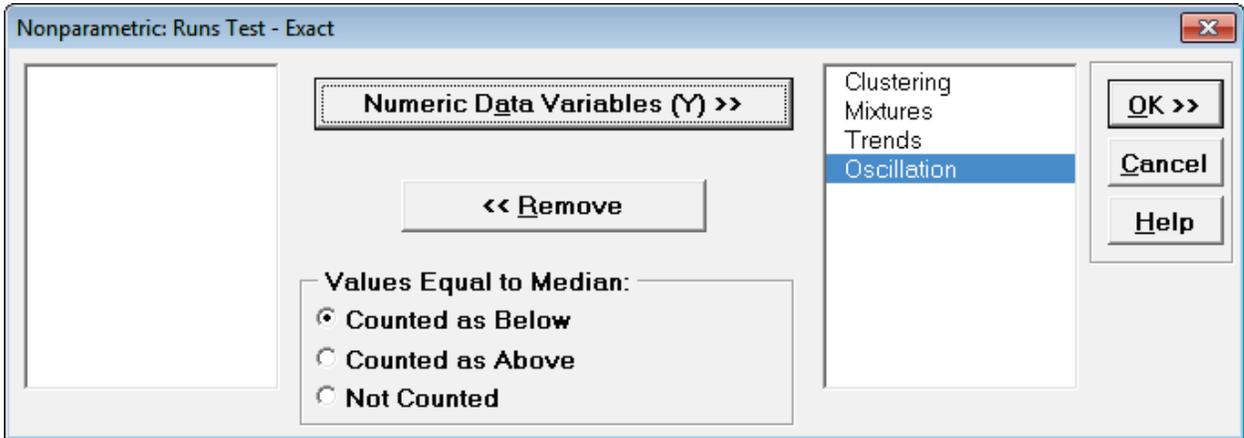
Nonparametric Runs Test - Exact	
Test Information	Overall Satisfaction
Number of Runs about Median:	44
Expected Number of Runs about Median:	51
Number of Points equal to or below Median:	50
Number of Points above Median:	50
Exact P-Value for Clustering:	0.0957
Exact P-Value for Mixtures:	0.9345
Exact P-Value for Lack of Randomness (2-Sided):	0.1914
Number of Runs Up or Down (zeros are negative):	64
Expected Number of Runs Up or Down:	66.333
P-Value for Trends:	0.3304
P-Value for Oscillation:	0.7512

With all of the P-Values being greater than 0.01 (alpha = .01 is preferred for the Runs Test to minimize false alarms), we fail to reject H_0 , and conclude that the data is random (or statistically independent).

The Above/Below and Up/Down Runs counts are identical to the above “large sample” or “asymptotic” results. The Clustering/Mixtures/Randomness (Above/Below) Exact P-Values are close but slightly different. This was expected because the sample size is reasonable ($N > 50$), so the “large sample” P-Values are valid using a normal approximation.

The P-Values for Trends and Oscillation (Up/Down) use a normal approximation because the sample size is greater than 50. They are, however, slightly different than the “large sample” above because a continuity correction is now applied to the normal approximation.

5. Now we will consider the small sample examples. Open **Runs Test Example Data.xlsx**, click **Runs Test Example Data** tab.
6. Click **SigmaXL > Statistical Tools > Nonparametric Tests – Exact > Runs Test – Exact**. If necessary, check **Use Entire Data Table**, click **Next**.
7. Select *Clustering* and shift click to *Oscillation*, click **Numeric Data Variables (Y) >>**. Set **Values Equal to Median:** to **Counted as Below**.



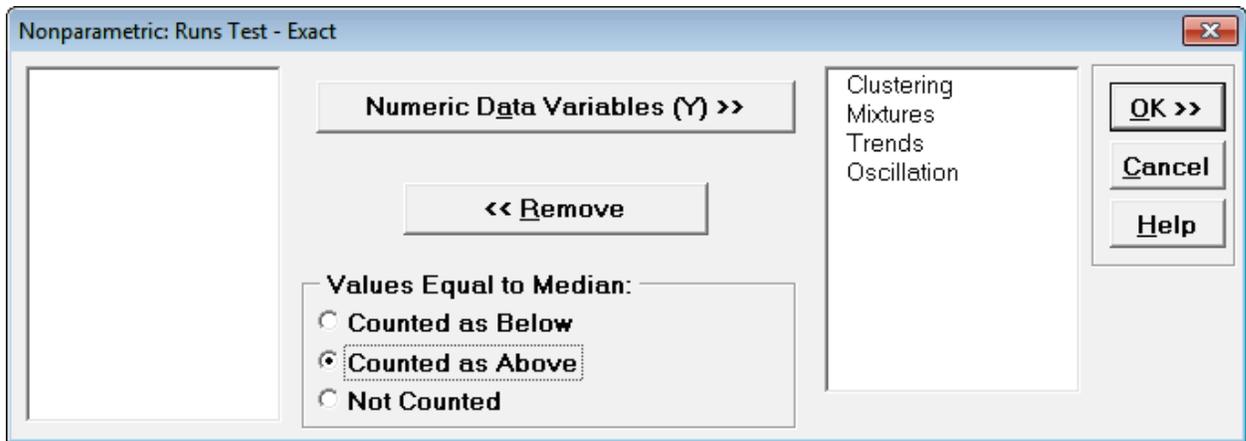
8. Click **OK**. Resulting output:

Nonparametric Runs Test - Exact				
Test Information	Clustering	Mixtures	Trends	Oscillation
Number of Runs about Median:	9	23	13	16
Expected Number of Runs about Median:	15.933	16	15.733	15.933
Number of Points above Median:	14	15	13	14
Number of Points equal to or below Median:	16	15	17	16
Exact P-Value for Clustering:	0.0073	0.9977	0.1980	0.5854
Exact P-Value for Mixtures:	0.9976	0.0070	0.8887	0.5664
Exact P-Value for Lack of Randomness (2-Sided):	0.0135	0.0139	0.3451	1.0000
Number of Runs Up or Down (zeros are negative):	17	22	13	25
Expected Number of Runs Up or Down:	19.667	19.667	19.667	19.667
Exact P-Value for Trends:	0.1662	0.8988	0.0028	0.9968
Exact P-Value for Oscillation:	0.9213	0.2072	0.9994	0.0130

The Exact P-Values are close to the above “large sample” or “asymptotic” results, but note that some of the values are now greater than .01 so they would fail-to reject H0. Note that this does not imply that the large sample runs test is more powerful, but rather we cannot conclude some of the previously identified patterns at the 99% confidence level.

We will now rerun the analysis using the **Counted as Above** option.

9. Press **F3** or click **Recall SigmaXL Dialog** to recall last dialog. Set **Values Equal to Median:** to **Counted as Above**.



10. Click **OK**. Results:

Nonparametric Runs Test - Exact				
Test Information	Clustering	Mixtures	Trends	Oscillation
Number of Runs about Median:	9	23	12	20
Expected Number of Runs about Median:	15.733	16	15.933	15.933
Number of Points equal to or above Median:	17	15	16	16
Number of Points below Median:	13	15	14	14
Exact P-Value for Clustering:	0.0084	0.9977	0.1007	0.9574
Exact P-Value for Mixtures:	0.9973	0.0070	0.9524	0.0919
Exact P-Value for Lack of Randomness (2-Sided):	0.0127	0.0139	0.1926	0.1394
Number of Runs Up or Down (zeros are positive):	17	22	13	25
Expected Number of Runs Up or Down:	19.667	19.667	19.667	19.667
Exact P-Value for Trends:	0.1662	0.8988	0.0028	0.9968
Exact P-Value for Oscillation:	0.9213	0.2072	0.9994	0.0130

Note the change in **Trends** and **Oscillation** Number of Runs and resulting dramatic change to the Exact P-Values.

In conclusion, when using Runs Test Exact, always try **Counted As Below** and **Counted as Above** to ensure that the P-Values agree with each other. Also, whenever you have a small sample size and are performing a nonparametric test, always use Exact.

Part 0 – Attribute/Discrete Data Tests

1 Proportion Test and Confidence Interval Template

Click **SigmaXL > Templates and Calculators > Basic Statistical Templates > 1 Proportion Test and Confidence Interval** or **SigmaXL > Statistical Tools > Basic Statistical Templates > 1 Proportion Test and Confidence Interval**.

See [Basic Statistical Templates – 1 Proportion Test and Confidence Interval](#) for an example of the 1 Proportion Test and Confidence Interval template.

2 Proportions Test and Confidence Interval Template

Click **SigmaXL > Templates and Calculators > Basic Statistical Templates > 2 Proportions Test and Confidence Interval** or **SigmaXL > Statistical Tools > Basic Statistical Templates > 2 Proportions Test and Confidence Interval**.

See below for an example of the 2 Proportions Test and Confidence Interval template.

2 Proportions Equivalence Test Template

Click **SigmaXL > Templates and Calculators > Basic Statistical Templates > 2 Proportions Equivalence Test** or **SigmaXL > Statistical Tools > Basic Statistical Templates > 2 Proportions Equivalence Test**.

See [Basic Statistical Templates – 2 Proportions Equivalence Test](#) for an example of the 2 Proportions Equivalence Test template.

One-Way Chi-Square Goodness-of-Fit Template

Click **SigmaXL > Templates & Calculators > Basic Statistical Templates > One-Way Chi-Square Goodness-of-Fit Test** (or **SigmaXL > Statistical Tools > Chi-Square Tests > One Way Chi-Square Goodness-of-Fit Template**) to access the One-Way Chi-Square Goodness-of-Fit Test calculator.

See [Basic Statistical Templates – One-Way Chi-Square Goodness-of-Fit Test](#) for an example of the One-Way Chi-Square Goodness-of-Fit Test template.

One-Way Chi-Square Goodness-of-Fit Template - Exact

Click **SigmaXL > Templates & Calculators > Basic Statistical Templates > One-Way Chi-Square Goodness-of-Fit Test Exact** (or **SigmaXL > Statistical Tools > Chi-Square Tests – Exact > One Way Chi-Square Goodness-of-Fit Exact Template**) to access the One-Way Chi-Square Goodness-of-Fit Test – Exact calculator.

See [Basic Statistical Templates – One-Way Chi-Square Goodness-of-Fit Test - Exact](#) for an example of the One-Way Chi-Square Goodness-of-Fit Test - Exact template.

1 Poisson Rate Test and Confidence Interval Template

Click **SigmaXL > Templates and Calculators > Basic Statistical Templates > 1 Poisson Rate Test and Confidence Interval** or **SigmaXL > Statistical Tools > Basic Statistical Templates > 1 Poisson Rate Test and Confidence Interval**.

See [Basic Statistical Templates – 1 Poisson Rate Test and Confidence Interval](#) for an example of the 1 Poisson Rate Test and Confidence Interval template.

2 Poisson Rates Test and Confidence Interval Template

Click **SigmaXL > Templates and Calculators > Basic Statistical Templates > 2 Poisson Rates Test and Confidence Interval** or **SigmaXL > Statistical Tools > Basic Statistical Templates > 2 Poisson Rates Test and Confidence Interval**.

See [Basic Statistical Templates – 2 Poisson Rates Test and Confidence Interval](#) for an example of the 2 Poisson Rates Test and Confidence Interval template.

2 Poisson Rates Equivalence Test Template

Click **SigmaXL > Templates and Calculators > Basic Statistical Templates > 2 Poisson Rates Equivalence Test** or **SigmaXL > Statistical Tools > Basic Statistical Templates > 2 Poisson Rates Equivalence Test**.

See [Basic Statistical Templates – 2 Poisson Rates Equivalence Test](#) for an example of the 2 Poisson Equivalence Test template.

2 Proportions Test and Confidence Interval Template Example

1. We begin with a scenario where Day Shift is running at 80% Yield and Night Shift has 70% Yield. This data is based on a random sample of 100 units for each Shift, each unit is either good or bad. Question: Is Day Shift running differently than Night Shift? Statistically, we call the Null Hypothesis, H_0 : Proportion $P_1 =$ Proportion P_2 ; the alternative hypothesis, H_a is: $P_1 \neq P_2$. If the calculated P-Value $< .05$, then we reject the null hypothesis and conclude that Day Shift and Night Shift are different.
2. Click **SigmaXL > Templates & Calculators > Basic Statistical Templates > 2 Proportions Test and Confidence Interval** (or **SigmaXL > Statistical Tools > Basic Statistical Templates > 2 Proportions Test and Confidence Interval**).
3. Enter $x_1 = 80$, $n_1 = 100$, $x_2 = 70$, $n_2 = 100$ as shown:

2 Proportions Test and Confidence Interval			
Sample Data (user inputs):		Sample 1	Sample 2
Number of Events	x	80	70
Sample Size	n	100	100
Null Hypothesis (hypothesized difference)	$H_0: P_1 - P_2 =$	0	
Alternative Hypothesis	$H_a: P_1 - P_2$	Not Equal To	
Confidence Level (enter .95 for 95%)	$100 \times (1 - \alpha)\%$	95.0%	
Hypothesis Test Method		Fisher's Exact	
Confidence Interval Method		Newcombe-Wilson Score	
Results:			
Sample proportion (x/n)		0.8000	0.7000
Sample proportion difference		0.1000	
alpha		0.0500	
Minimum expected value (should be ≥ 5 for normal approximation)		25.0000	
Fisher's Exact probability P-Value (2-sided)		0.1412	
Upper Confidence Limit (2-sided)		0.2167	
Lower Confidence Limit (2-sided)		-0.0202	

Since the Fisher's exact P-Value of 0.141 is greater than .05, we fail to reject H_0 . We do not have enough evidence to show that there is a significant difference between Day Shift and Night shift. This does not mean that we have proven that they are the same. In practice however, we either assume that they are the same or we collect more data.

Note: Fisher's exact P-Values should be used for any real-world problem. The approximate P-Values based on the normal distribution are provided for instructional purposes, e.g., comparing to hand calculations.

4. Now enter $x_1 = 160$, $n_1 = 200$, $x_2 = 140$, $n_2 = 200$. Note that the Fisher's exact P-Value is now .028, so we reject H_0 .

Chi-Square Test – Two-Way Table Data

1. Open the file **Attribute Data.xlsx**, ensure that **Example 1** Sheet is active. This data is in Two – Way Table format, or pivot table format. Note that cells B2:D4 have been pre-selected.
2. Click **SigmaXL > Statistical Tools > Chi-Square Tests > Chi-Square Test & Association – Two-Way Table Data**. Note the selection of data includes the Row and Column labels (if we had Row and Column Totals these would NOT be selected). Do not check **Advanced Tests and Measures of Association**.

Chi-Square Test & Association - Two Way Table Data

Please select your data

\$B\$2:\$D\$4

Note:
 1. Two-Way (pivot) table must be based on a single Y or count variable.
 2. Selection should include only one row and only one column label.
 3. Selection should exclude Grand Totals.
 See Help for example.

Advanced Tests and Measures of Association

Nominal Categories

Ordinal Categories

Confidence Level: 95.0

Help Cancel Next >>

See Appendix [Chi-Square Tests and \(Contingency\) Table Associations](#) for the Chi-Square Table Statistics formula details.

3. Click **Next**. The resulting output is:

Chi-Square 2 Way Table Statistics		
Observed Counts	Day Shift	Night Shift
Pass	80	70
Fail	20	30
Expected Counts	Day Shift	Night Shift
Pass	75	75
Fail	25	25
Std. Residuals	Day Shift	Night Shift
Pass	0.577350	-0.577350
Fail	-1	1
Chi-Square	2.667	
DF	1	
P-Value	0.1025	

The P-Value matches that of the 2 Proportions test. Since the P-Value of 0.1 is greater than .05, we fail to reject H0.

- Now click **Example 2** Sheet tab. The Yields have not changed but we have doubled the sample size. Repeat the above analysis. The resulting output is:

Chi-Square 2 Way Table Statistics		
Observed Counts	Day Shift	Night Shift
Pass	160	140
Fail	40	60
Expected Counts	Day Shift	Night Shift
Pass	150	150
Fail	50	50
Std. Residuals	Day Shift	Night Shift
Pass	0.816497	-0.816497
Fail	-1.414	1.414
Chi-Square	5.333	
DF	1	
P-Value	0.0209	

Since the P-Value is $< .05$, we now reject the Null Hypothesis, and conclude that Day Shift and Night Shift are significantly different. The Std. (Standardized) Residuals tell us that Day Shift failures are less than expected (assuming equal proportions), and Night Shift failures are more than expected.

Note, by doubling the sample size, we improved the power or sensitivity of the test.

- Click the **Example 3** Sheet tab. In this scenario we have 3 suppliers and an additional marginal level. A random sample of 100 units per supplier is tested. The null hypothesis here is: No relationship between Suppliers and Pass/Fail/Marginal rates, but in this case, we can state it as No difference across suppliers. Redoing the above analysis (for selection B2:E5) yields the following:

Chi-Square 2 Way Table Statistics			
Observed Counts	Supplier A	Supplier B	Supplier C
Pass	80	70	75
Fail	10	15	18
Marginal	10	15	7
Expected Counts	Supplier A	Supplier B	Supplier C
Pass	75	75	75
Fail	14.333	14.333	14.333
Marginal	10.667	10.667	10.667
Std. Residuals	Supplier A	Supplier B	Supplier C
Pass	0.577350	-0.577350	0
Fail	-1.145	0.176090	0.968496
Marginal	-0.204124	1.327	-1.123
Chi-Square	6.008		
DF	4		
P-Value	0.1985		

The P-Value tells us that we do not have enough evidence to show that there is a difference across the 3 suppliers.

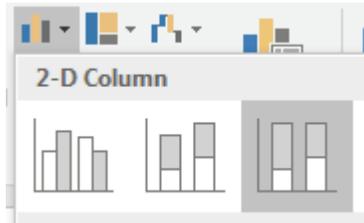
- 6. Click the **Example 4** Sheet tab. Here we have doubled the sample size to 200 per supplier. Note that the percentages are identical to example 3. Redoing the above analysis yields the following:

Chi-Square 2 Way Table Statistics			
Observed Counts	Supplier A	Supplier B	Supplier C
Pass	160	140	150
Fail	20	30	36
Marginal	20	30	14
Expected Counts	Supplier A	Supplier B	Supplier C
Pass	150	150	150
Fail	28.667	28.667	28.667
Marginal	21.333	21.333	21.333
Std. Residuals	Supplier A	Supplier B	Supplier C
Pass	0.816497	-0.816497	0
Fail	-1.619	0.249029	1.370
Marginal	-0.288675	1.876	-1.588
Chi-Square	12.016		
DF	4		
P-Value	0.0172		

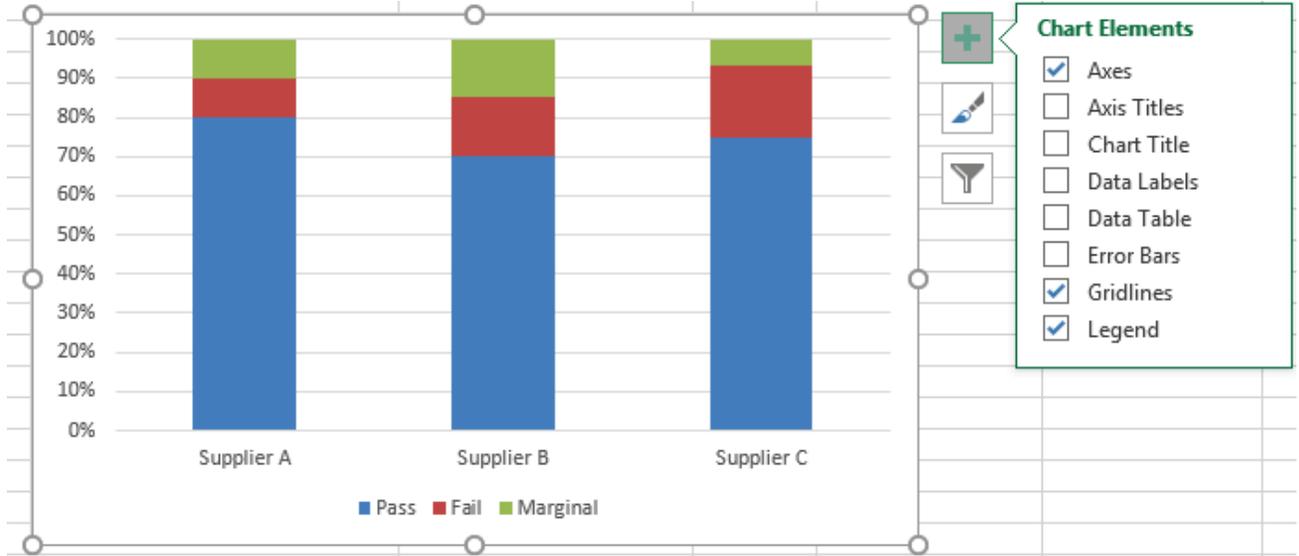
With the P-Value < .05 we now conclude that there is a significant difference across suppliers. Examining the Std. (Standardized) Residuals tells us that Supplier A has fewer failures than

expected (if there was no difference across suppliers), Supplier B has more marginal parts than expected and Supplier C has fewer marginal parts than expected.

7. The table row and column cell percentages can be visualized using Excel's 100% Stacked Column Chart. Select cells **A3:D6** of the **Chi-Square** sheet. Click Excel's **Insert > Insert Column or Bar Chart** and select **100% Stacked Column** as shown.

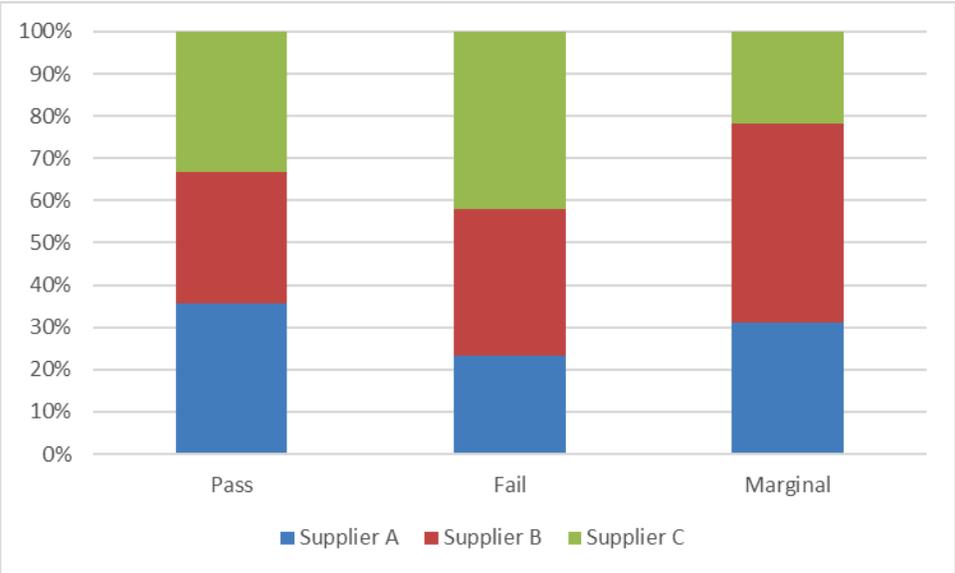


8. Click to create the 100% stacked column chart. Uncheck the Chart Title as shown.



9. The rows and columns can easily be switched by clicking **Design > Switch Row/Column**





10. These charts make it easy to visualize the cell row and column percentages.

Chi-Square Test – Two-Way Table Data: Advanced Tests and Measures of Association – Nominal Categories

Checking the **Nominal Categories** option provides additional chi-square statistics and measures of association, including:

- Adjusted Residuals
 - Equivalent to normal z score
 - Red font highlight denotes significant cell residual value
 - Bold red highlight denotes significant cell residual value with Bonferroni adjustment
 - Note: red highlight is only active if Chi-Square P-Value is significant
- Cell's Contribution to Chi-Square
- Additional Chi-Square Tests
 - Likelihood Ratio
 - McNemar-Bowker Symmetry (Square Table)
 - For a 2x2 table, McNemar's test is equivalent to a paired two-proportions test, for example applicable to studying before versus after change in proportion on the same subject. The returned P-Value is exact, based on the binomial distribution.
 - Bowker extended McNemar's test for square tables larger than 2x2. The null hypothesis is that the table is symmetrical (i.e., symmetry of disagreement). This uses Chi-Square as the test statistic.
- Measures of Association for Nominal Categories
 - Pearson's Phi
 - Pearson's Phi is equivalent to Pearson's correlation coefficient for a 2x2 table. It is the most popular measure of association for 2x2 tables.
 - We recommend the following rules-of-thumb, adapted from Cohen (1988):
 - < 0.1 = Very Weak
 - 0.1 to < 0.3: Weak ("Small" Effect)
 - 0.3 to < 0.5: Moderate ("Medium" Effect)
 - > 0.5: Strong ("Large" Effect)
 - Although Phi is equivalent to Pearson's correlation for a 2x2 table, we recommend these rules-of-thumb for use in typical contingency tables, rather than those commonly used for correlation (i.e., > 0.9 = Strong).

- Cramer's V
 - Cramer's V is an extension of Phi for larger tables. It is the most popular measure of association for tables of any size.
 - It varies from 0 to 1, with 0 = no association and 1 = perfect association.
 - Use Cohen's rules-of-thumb given above for Phi.
- Contingency Coefficient
 - An alternative to Phi, varies from 0 to < 1.
 - Use Cohen's rules-of-thumb given above for Phi.
- Cohen's Kappa (Agreement - Square Table)
 - Kappa is used to measure agreement between two assessors evaluating the same parts or items.
 - For an extended Attribute Measurement Systems Analysis use **SigmaXL > Measurement Systems Analysis > Attribute MSA**.
 - For Attribute MSA used in Six Sigma quality, the recommendation is Kappa > 0.9 is strong agreement and < 0.7 is weak agreement, but for general use, the less stringent guidelines by Fleiss are recommended:

Kappa: ≥ 0.75 or so signifies excellent agreement, for most purposes, and ≤ 0.40 or so signifies poor agreement.
 - See Appendix [Kappa](#) for further details.
- Goodman-Kruskal Lambda & Tau and Theil's Uncertainty
 - Measures of Proportional Reduction in Predictive Error. The basic concept is a measure that indicates how much knowing the value of the independent variable improves our ability to estimate the value of the dependent variable.
 - They are Directional Measures. If the Y dependent variable is in the **Rows Category**, then use the *Rows Dependent* measure. If the Y dependent variable is in the **Columns Category**, then use the *Cols Dependent* measure. If there is no clear X-Y dependent-independent relationship, then use the *Symmetric* measures (not available for Tau).
 - Use Cohen's rules-of-thumb for these measures.

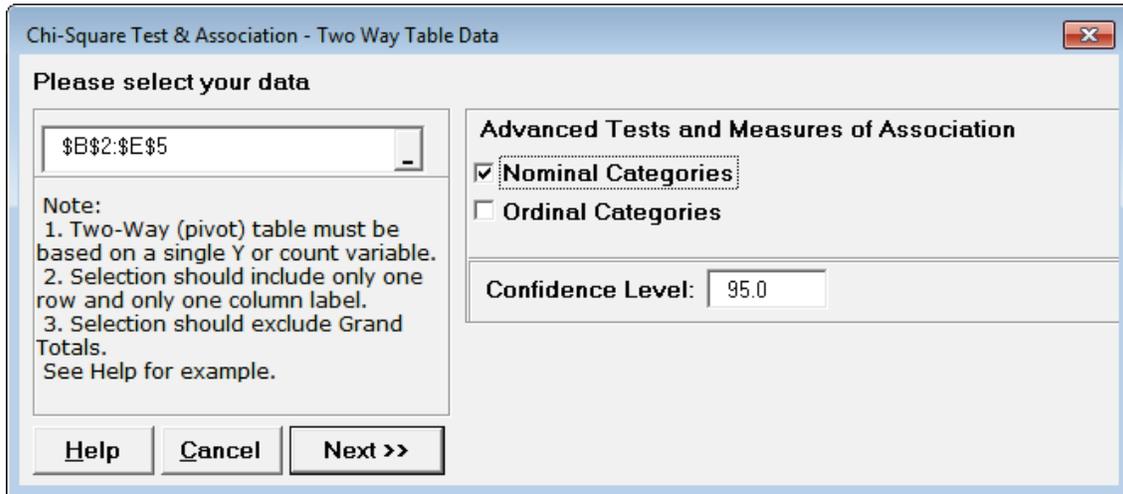
See Appendix [Chi-Square Tests and \(Contingency\) Table Associations](#) for further formula details and references. The following are external links with helpful presentations on measures of association for contingency tables:

- http://course1.winona.edu/bdeppa/STAT%20701%20Online/stat_701%20home.htm
Click **4a) Measures of Association**. This is a narrated presentation, but a regular PowerPoint

is also available on the web site. It is part of a Biostatistics course by Dr. Brant Deppa, Department of Mathematics and Statistics, Winona State University.

- <http://uregina.ca/~gingrich/ch11a.pdf> by Dr. Paul Gingrich, University of Regina.

1. Press **F3** or click **Recall SigmaXL Dialog** to recall last dialog. Check **Nominal Categories** as shown:



2. Click **Next**. The resulting output is:

Chi-Square 2 Way Table Statistics			
Observed Counts	Supplier A	Supplier B	Supplier C
Pass	160	140	150
Fail	20	30	36
Marginal	20	30	14
Expected Counts	Supplier A	Supplier B	Supplier C
Pass	150	150	150
Fail	28.667	28.667	28.667
Marginal	21.333	21.333	21.333
Std. Residuals	Supplier A	Supplier B	Supplier C
Pass	0.816497	-0.816497	0
Fail	-1.619	0.249029	1.370
Marginal	-0.288675	1.876	-1.588
Adjusted Residuals	Supplier A	Supplier B	Supplier C
Pass	2	-2	0
Fail	-2.14191972	0.329526111	1.812393613
Marginal	-0.374066	2.431429007	-2.05736301
Cell's Contribution to Chi-Square	Supplier A	Supplier B	Supplier C
Pass	0.66666667	0.66666667	0
Fail	2.62015504	0.062015504	1.875968992
Marginal	0.08333333	3.520833333	2.520833333

The adjusted residuals are equivalent to normal z values, so for a specified 95% confidence level, any value greater than 1.96 (or less than -1.96) is highlighted in red. This results in a slight difference in interpretation from that of the standardized residuals, but the 3 largest magnitude residuals are consistent.

Chi-Square	12.016			
DF	4			
P-Value	0.0172			
Additional Chi-Square Tests				
Test	Chi-Square	DF	P-Value	
Likelihood Ratio	12.1425	4	0.0163	
McNemar-Bowker Symmetry (Square Table)	189.9572	3	0.0000	
Measures of Association for Nominal Categories				
Measure	Value	Std. Error	95% Lower Bound	95% Upper Bound
Pearson's Phi	0.1415			
Cramer's V	0.1001			
Contingency Coefficient	0.1401			
Cohen's Kappa (Agreement - Square Table)	0.0100	0.0211	-0.0314	0.0514
Goodman-Kruskal Lambda (Cols Dependent)	0.0650	0.0249	0.0162	0.1138
Goodman-Kruskal Lambda (Rows Dependent)	0.0000	0.0000	0.0000	0.0000
Goodman-Kruskal Lambda (Symmetric)	0.0473	0.0181	0.0118	0.0827
Goodman-Kruskal Tau (Cols Dependent)	0.0100	0.0057	0.0000	0.0211
Goodman-Kruskal Tau (Rows Dependent)	0.0095	0.0058	0.0000	0.0209
Theil's Uncertainty (Cols Dependent)	0.0092	0.0052	-0.0010	0.0194
Theil's Uncertainty (Rows Dependent)	0.0138	0.0078	-0.0015	0.0291
Theil's Uncertainty (Symmetric)	0.0110	0.0062	-0.0012	0.0233

As noted above, the Chi-Square P-Value tells us that there is a significant difference across suppliers, in other words, there is association between Supplier and Pass/Fail/Marginal, but it does not tell us the degree or strength of that association. Cramer's V is used for tables larger than 2x2 and from the rules-of-thumb, the 0.1 value is considered weak (or small effect).

Computing Odds Ratio and Confidence Interval for 2x2 Table

Odds Ratio and Confidence Intervals are not directly available for 2x2 Tables but can be obtained using Logistic Regression. The table data must be rearranged to stacked column format and the response changed to “Event/Trial” format. **Example 2 of Attribute Data.xlsx**:

	Day Shift	Night Shift
Pass	160	140
Fail	40	60

would be changed to:

Shift (X)	Pass (Y)	Trials (N)
1	160	200
2	140	200

Note, coding Shift as continuous numeric is easier to interpret and is valid because it has a range of 1.

Analyze using **SigmaXL > Statistical Tools > Regression > Binary Logistic Regression**. Select **Response Count (Y) / Sample Size**. *Pass(Y)* is selected as **Numeric Response Count (Y)**, *Trials (N)* is **Numeric Sample Size (Trials)** and *Shift (X)* is selected as **Continuous Predictors (X)**:

The screenshot shows the 'Binary Logistic Regression' dialog box. It has several sections for configuration:

- Response Type:** Radio buttons for 'Binary Response (Y)' and 'Response Count (Y) / Sample Size (Trials)'. The second option is selected.
- Numeric Response Count (Y) >>:** A text box containing 'Pass (Y)'.
- Numeric Sample Size (Trials) >>:** A text box containing 'Trials (N)'.
- Continuous Predictors (X) >>:** A list box containing 'Shift (X)', which is highlighted in blue.
- Categorical Predictors (X) >>:** An empty list box.
- Reference Event >>:** A text box containing '1'.

Click **OK**. The Odds Ratio and Confidence Limits are given as:

Odds Ratio	Lower 95% Odds Ratio	Upper 95% Odds Ratio
0.583333	0.368282	0.923960

The shift change 1 to 2 is 0.583 times as likely to produce passed (good) product with a 95% confidence interval of 0.368 to 0.924.

Chi-Square Test – Two-Way Table Data: Advanced Tests and Measures of Association – Ordinal Categories

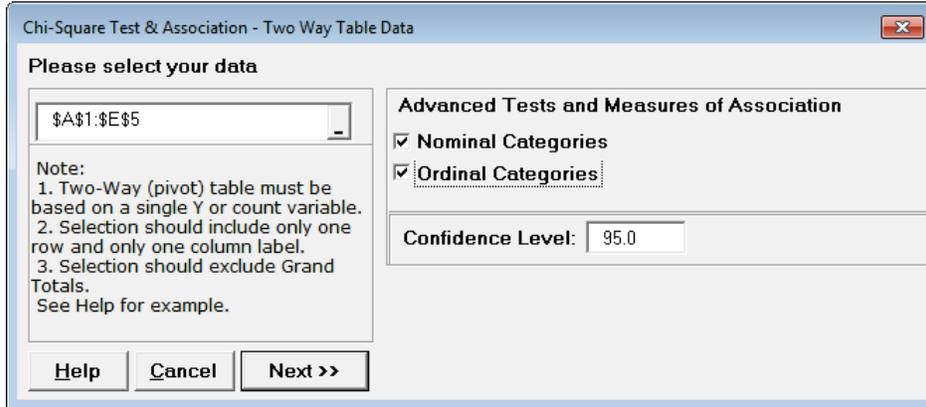
Checking the **Ordinal Categories** option provides statistics and measures of association appropriate when both row and column category variables are ordinal:

- Adjusted Residuals and Cell's Contribution to Chi-Square
- Tests of Association for Ordinal Categories
 - Concordant – Discordant
 - The P-Value for this hypothesis test is from Kendall's Tau-B, but is the same for all of the Concordant – Discordant ordinal measures: Tau-C, Gamma and Somers' D. See Agresti (2010). This may differ from other software using an approximation formula.
 - Spearman Rank Correlation
- Measures of Association for Ordinal Categories with Confidence Intervals
 - Spearman Rank Correlation
 - Equivalent to Pearson's correlation on ranks
 - Kendall's Tau-B (Square Table), Kendall-Stuart Tau-C (Rectangular Table), Goodman-Kruskal Gamma
 - Use Tau-B for square tables (no. rows = no. columns) and Tau-C for rectangular tables (no. rows <> no. columns).
 - Tau-B, Tau-C and Gamma are asymmetric measures, so will give the same result regardless of variable assignment to Rows and Columns.
 - These are all Concordant – Discordant measures.
 - Somers' D (Cols & Rows Dependent, Symmetric)
 - Also a Concordant – Discordant measure but directional. If the Y dependent variable is in the Rows Category, then use the Rows Dependent measure. If the Y dependent variable is in the Columns Category, then use the Cols Dependent measure. If there is no clear X-Y dependent-independent relationship, then use the Symmetric measure.
- SigmaXL provides rules-of-thumb for Kendall's Correlation in Ordinal Attribute MSA (strong association is > 0.8) and Pearson or Spearman Correlation (strong association is > 0.9), however these are in the context of measurement systems analysis, design of experiments or a controlled process study. For typical contingency table applications, we recommend the rules-of-thumb, adapted from Cohen (1988):

- 0.5+: Strong (Large Effect)
- 0.3 to < 0.5: Moderate (Medium Effect)
- 0.1 to < 0.3: Weak (Small Effect)
- < 0.1: Very Weak

See Appendix [Chi-Square Tests and \(Contingency\) Table Associations](#) for further formula details and references. See also Agresti (2010, Chapter 7 *Non-Model-Based Analysis of Ordinal Association*) and Deppa's presentation on Measures of Association at: http://course1.winona.edu/bdeppa/STAT%20701%20Online/stat_701%20home.htm.

1. Open the file **Attribute Data.xlsx**, click **Example 5 – Salary Sat** Sheet tab. This data is in two-way table format and has ordinal categories: *Salary* in the Rows and *Satisfaction Level* in the Columns. Note that cells A1:E5 have been pre-selected.
2. Click **SigmaXL > Statistical Tools > Chi-Square Tests > Chi-Square Test & Association – Two-Way Table Data**. Note the selection of data includes the Row and Column labels (if we had Row and Column Totals these would NOT be selected). Check **Nominal Categories** and **Ordinal Categories** as shown:



Tip: Even if the categories are ordinal, it is sometimes useful to select nominal categories as well for comparison purposes.

3. Click **Next**. The resulting output is:

Chi-Square 2 Way Table Statistics				
Observed Counts	Very Dissatisfied	Somewhat Dissatisfied	Somewhat Satisfied	Very Satisfied
< 20K	5	6	19	12
20 to < 30K	5	7	20	14
30 to < 40K	3	12	28	24
>= 40K	1	2	18	26
Expected Counts	Very Dissatisfied	Somewhat Dissatisfied	Somewhat Satisfied	Very Satisfied
< 20K	2.911	5.614	17.673	15.802
20 to < 30K	3.188	6.149	19.356	17.307
30 to < 40K	4.644	8.955	28.193	25.208
>= 40K	3.257	6.282	19.777	17.683
Std. Residuals	Very Dissatisfied	Somewhat Dissatisfied	Somewhat Satisfied	Very Satisfied
< 20K	1.224	0.162972	0.315591	-0.956432
20 to < 30K	1.014757961	0.343393	0.146278	-0.794905
30 to < 40K	-0.762713	1.017372859	-0.036361522	-0.240586
>= 40K	-1.251	-1.708	-0.399632	1.978
Adjusted Residuals	Very Dissatisfied	Somewhat Dissatisfied	Somewhat Satisfied	Very Satisfied
< 20K	1.426134732	0.196736411	0.465932976	-1.360693153
20 to < 30K	1.196940639	0.419818923	0.218713418	-1.145298952
30 to < 40K	-0.967089216	1.33704372	-0.05844315	-0.37262257
>= 40K	-1.480071878	-2.095444235	-0.599449996	2.858762784
Cell's Contribution to Chi-Square	Very Dissatisfied	Somewhat Dissatisfied	Somewhat Satisfied	Very Satisfied
< 20K	1.499326463	0.026559799	0.099597859	0.914762153
20 to < 30K	1.029733719	0.117919038	0.02139728	0.631873485
30 to < 40K	0.581730668	1.035047535	0.00132216	0.057881515
>= 40K	1.564416624	2.918900046	0.15970582	3.911611766

Note: 4 out of 16 cells have expected counts less than 5.

Chi-Square	14.572			
DF	9			
P-Value	0.1034			
Additional Chi-Square Tests				
Test	Chi-Square	DF	P-Value	
Likelihood Ratio	15.4156	9	0.0801	
McNemar-Bowker Symmetry (Square Table)	32.8921	6	0.0000	
Measures of Association for Nominal Categories				
Measure	Value	Std. Error	95% Lower Bound	95% Upper Bound
Pearson's Phi	0.2686			
Cramer's V	0.1551			
Contingency Coefficient	0.2594			
Cohen's Kappa (Agreement - Square Table)	0.0752	0.0428	-0.0087	0.1592
Goodman-Kruskal Lambda (Cols Dependent)	0.0684	0.0547	0.0000	0.1756
Goodman-Kruskal Lambda (Rows Dependent)	0.0296	0.0556	0.0000	0.1385
Goodman-Kruskal Lambda (Symmetric)	0.0476	0.0483	0.0000	0.1424
Goodman-Kruskal Tau (Cols Dependent)	0.0230	0.0136	0.0000	0.0497
Goodman-Kruskal Tau (Rows Dependent)	0.0233	0.0115	0.0008	0.0458
Theil's Uncertainty (Cols Dependent)	0.0322	0.0153	0.0022	0.0621
Theil's Uncertainty (Rows Dependent)	0.0279	0.0134	0.0017	0.0541
Theil's Uncertainty (Symmetric)	0.0299	0.0143	0.0019	0.0578

Tests of Association for Ordinal Categories		
Test	Value	P-Value
Concordant - Discordant	2707	0.0009
Spearman Rank Correlation	0.2205	0.0016

Measures of Association for Ordinal Categories				
Measure	Value	Std. Error	95% Lower Bound	95% Upper Bound
Spearman Rank Correlation	0.2205	0.0669	0.0834	0.3495
Kendall's Tau-B (Square Table)	0.1899	0.0570	0.0782	0.3017
Kendall-Stuart Tau-C (Rectangular Table)	0.1769	0.0535	0.0721	0.2817
Goodman-Kruskal Gamma	0.2700	0.0795	0.1141	0.4258
Somers' D (Cols Dependent)	0.1791	0.0540	0.0732	0.2850
Somers' D (Rows Dependent)	0.2014	0.0603	0.0833	0.3196
Somers' D (Symmetric)	0.1896	0.0569	0.0781	0.3012

Note that the Chi-Square P-Value is 0.1, indicating that there is no significant association between Salary and Satisfaction when they are treated as nominal categories (although the significant result for McNemar-Bowker does show that there is lack of symmetry in the off diagonals).

Since the Chi-Square P-Value is not significant, the **Adjusted Residuals** are not highlighted, even though some values are greater than 1.96 (and less than -1.96). This follows the concept used in ANOVA called "Fisher Protected" where one considers the significance of post-hoc tests only when the overall test is significant.

Note: 4 out of 16 cells have expected counts less than 5. If more than 20% of the cells have expected counts less than 5 (or if any of the cells have an expected count less than 1), the Chi-Square approximation may be invalid, and Fisher's Exact should be used. This will be discussed

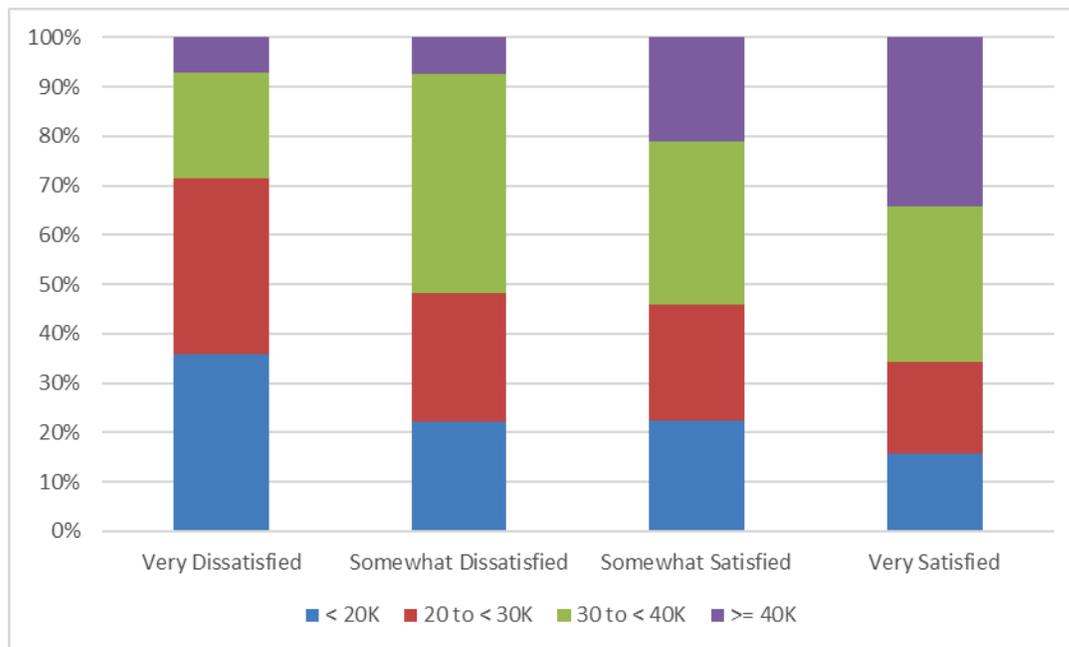
later, but for this example the Fisher's Monte-Carlo Exact P-Value = 0.095 so does not change the interpretation of the results for the above Chi-Square analysis.

When Salary and Satisfaction are treated as ordinal categories, the more powerful **Concordant – Discordant** and **Spearman Rank Correlation** P-Values clearly show that there is a significant association. The **Measures of Association for Ordinal Categories** show that this is positive, i.e., an increase in Salary is associated with an increase in Satisfaction. However, using the rules-of-thumb given above, we see that the association is weak.

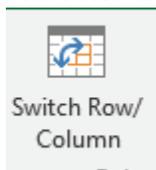
- The table row and column cell percentages can be visualized using Excel's 100% Stacked Column Chart. Select cells **A3:E7** of the **Chi-Square** sheet. Click Excel's **Insert > Insert Column or Bar Chart** and select **100% Stacked Column** as shown.

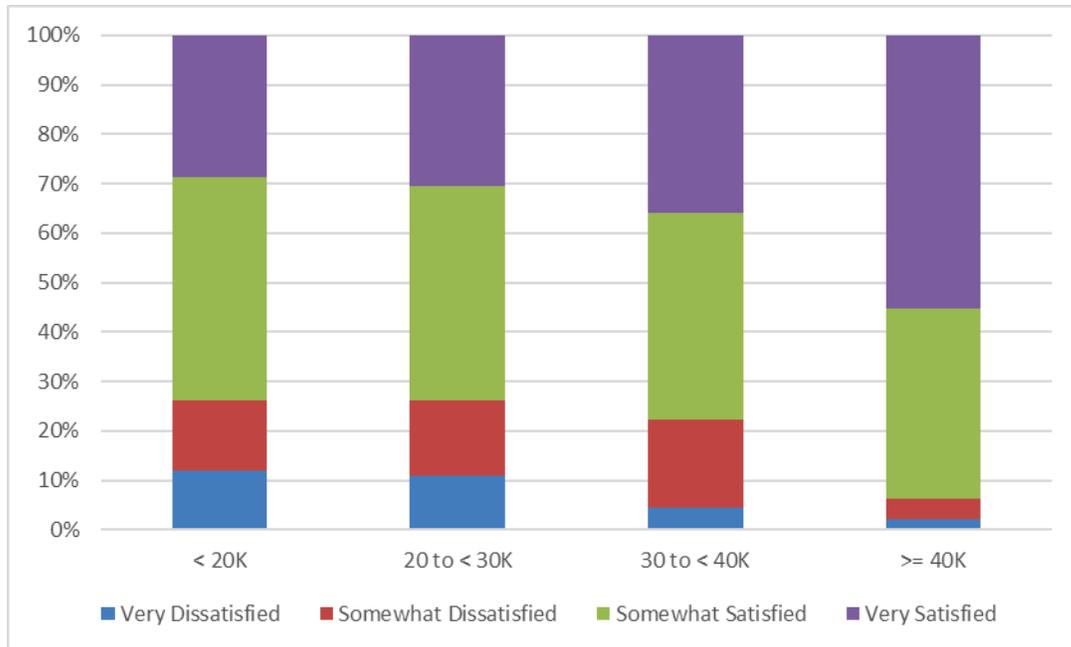


- Click to create the 100% stacked column chart (uncheck the Chart Title):



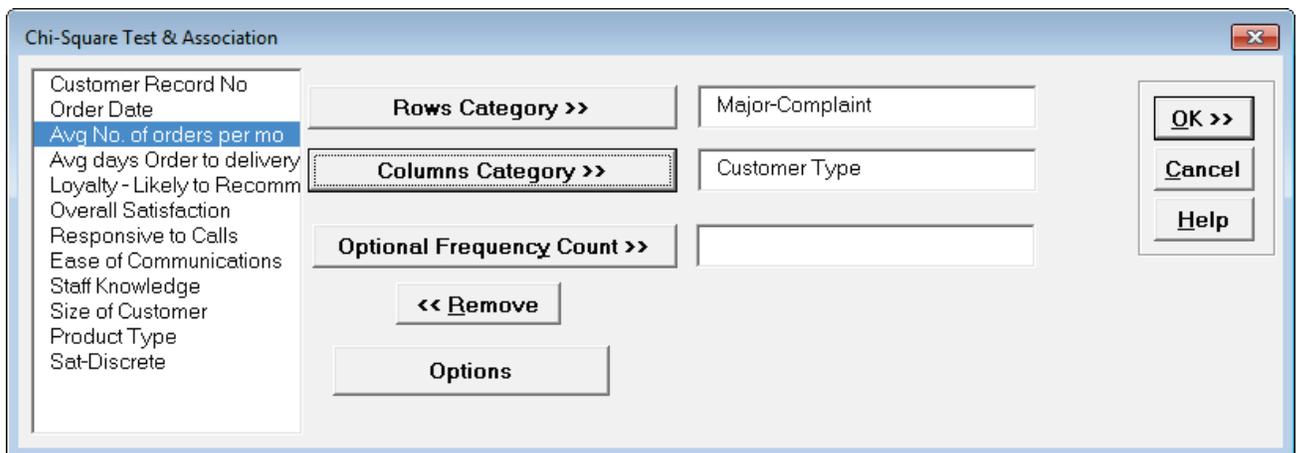
- The rows and columns can easily be switched by clicking **Design > Switch Row/Column**





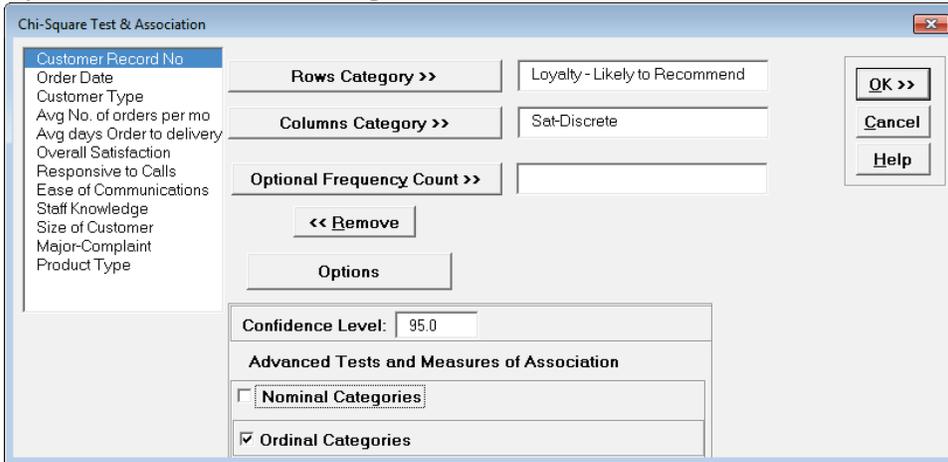
Chi-Square Test (Stacked Column Format Data)

1. Open **Customer Data.xlsx**. Click **Sheet 1** tab. The discrete data of interest is Complaints and Customer Type, i.e., does the type of complaint differ across customer type? Formally the Null Hypothesis is that there is no relationship (or independence) between Customer Type and Complaints.
2. Click **SigmaXL > Statistical Tools > Chi-Square Tests > Chi-Square Test & Association**. Ensure that the entire data table is selected. If not, check Use **Entire Data Table**. Click **Next**.
3. Select *Major-Complaint*, click **Rows Category >>**; select *Customer Type*, click **Columns Category >>**. We will not be using **Options** in this example.



4. Click **OK**. Results:

- Press **F3** or click **Recall SigmaXL Dialog** to recall last dialog. Select *Loyalty – Likely to Recommend*, click **Rows Category >>**; select *Sat-Discrete*, click **Columns Category >>**. Click **Options**, check **Ordinal Categories**.



Sat-Discrete is derived from Overall Satisfaction where a score ≥ 3.5 is considered a 1, and scores < 3.5 are considered a 0.

- Click **OK**. Results:

Chi-Square Test		
Loyalty - Likely to Recommend - Sat-Discrete		
Observed Counts	0	1
1	2	0
2	8	2
3	17	16
4	6	37
5	0	12
Expected Counts	0	1
1	0.660000	1.340
2	3.300	6.700
3	10.890	22.110
4	14.190	28.810
5	3.960	8.040
Std. Residuals	0	1
1	1.649	-1.158
2	2.587	-1.816
3	1.852	-1.299
4	-2.174	1.526
5	-1.990	1.397
Adjusted Residuals	0	1
1	2.03555	-2.036
2	3.33183	-3.332
3	2.76346	-2.763
4	-3.51818	3.518
5	-2.59161	2.592
Cell's Contribution to Chi-Square	0	1
1	2.72061	1.34
2	6.69394	3.297
3	3.42811	1.688
4	4.727	2.328
5	3.96	1.95
Note: 4 out of 10 cells have expected counts less than 5.		

Chi-Square	32.134			
DF	4			
P-Value	0.0000			
Tests of Association for Ordinal Categories				
Test	Value	P-Value		
Concordant - Discordant	1417	0.0000		
Spearman Rank Correlation	0.5558	0.0000		
Measures of Association for Ordinal Categories				
Measure	Value	Std. Error	95% Lower Bound	95% Upper Bound
Spearman Rank Correlation	0.5558	0.0943	0.3910	0.6861
Kendall's Tau-B (Square Table)	0.5163	0.0637	0.3915	0.6410
Kendall-Stuart Tau-C (Rectangular Table)	0.5668	0.0814	0.4073	0.7263
Goodman-Kruskal Gamma	0.8330	0.0678	0.7001	0.9660
Somers' D (Cols Dependent)	0.4159	0.0553	0.3076	0.5243
Somers' D (Rows Dependent)	0.6409	0.0795	0.4850	0.7967
Somers' D (Symmetric)	0.5044	0.0622	0.3826	0.6263

As expected, there is a strong positive association between Loyalty and Discrete Satisfaction. Note: Kendall-Stuart Tau-C should be used here rather than Tau-B because it is a rectangular table. Since Satisfaction leads to Loyalty, Loyalty is the Dependent variable, so Rows Dependent Somers' D should be used rather than Cols Dependent or Symmetric.

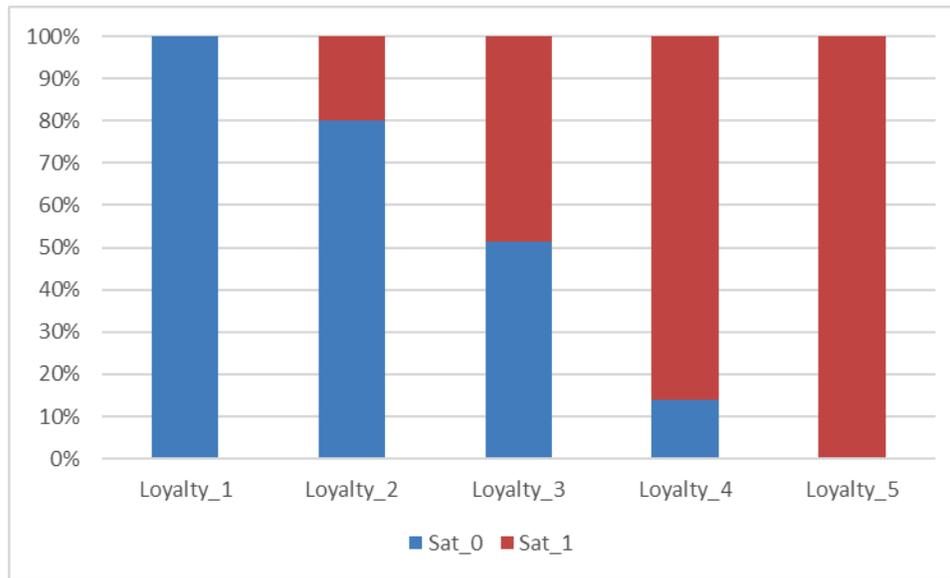
- In order to visualize the row and column percentages with Excel's 100% Stacked Column Chart, we will need to modify the numeric row and column labels as shown, converting to text as shown:

Observed Counts	Sat_0	Sat_1
Loyalty_1	2	0
Loyalty_2	8	2
Loyalty_3	17	16
Loyalty_4	6	37
Loyalty_5	0	12

- Select cells **A3:C8** of the **Chi-Square** sheet. Click Excel's **Insert > Insert Column or Bar Chart** and select **100% Stacked Column** as shown.



9. Click to create the 100% stacked column chart (uncheck the Chart Title):



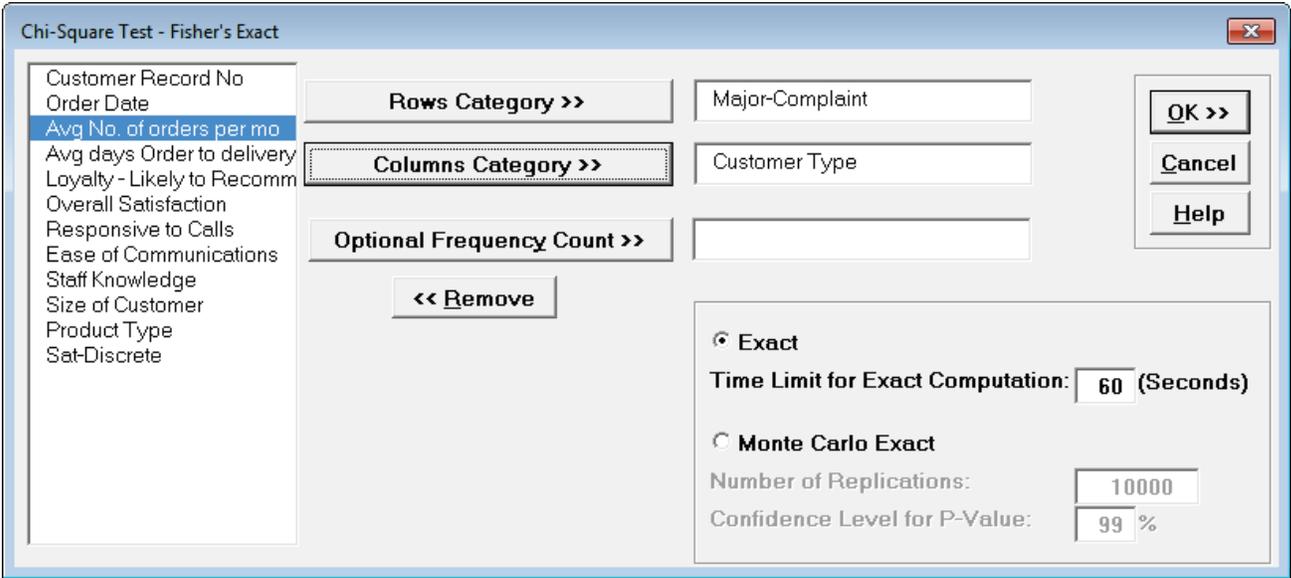
This clearly shows the strong relationship between the Loyalty scores and Discrete Satisfaction.

Chi-Square Test – Fisher’s Exact

The **Chi-Square Test – Fisher’s Exact** utilizes permutations and fast network algorithms to solve the Exact Fisher P-Value for contingency (two-way row*column) tables. This is an extension of the Fisher Exact option provided in the Two Proportion Test template. For data that requires more computation time than specified, Monte Carlo P-Values provide an approximate (but unbiased) P-Value that typically matches exact to two decimal places using 10,000 replications. One million replications give a P-Value that is typically accurate to three decimal places. A confidence interval (99% default) is given for the Monte Carlo P-Values. For further details refer to the Appendix [Exact and Monte Carlo P-Values for Nonparametric and Contingency Tests](#).

It is important to note that while exact P-Values are “correct,” they do not increase (or decrease) the power of a small sample test, so they are not a solution to the problem of failure to detect a change due to inadequate sample size!

1. We will now re-analyze the above Major Complaint by Customer Type data. Open **Customer Data.xlsx**. Click **Sheet 1** tab. The discrete data of interest is Complaints and Customer Type, i.e., does the type of complaint differ across customer type? Formally the Null Hypothesis is that there is no relationship (or independence) between Customer Type and Complaints.
2. Click **SigmaXL > Statistical Tools > Chi-Square Tests – Exact > Chi-Square Test – Fisher’s Exact**. Ensure that the entire data table is selected. If not, check Use **Entire Data Table**. Click **Next**.
3. Select *Major-Complaint*, click **Rows Category >>**; select *Customer Type*, click **Column Category >>**. Select **Exact** with the default **Time Limit for Exact Computation = 60** seconds.



Tip: If the exact computation time limit is exceeded a dialog will prompt you to use Monte Carlo or to increase the computation time. When this occurs, Monte Carlo is recommended.

4. Click **OK**. Results:

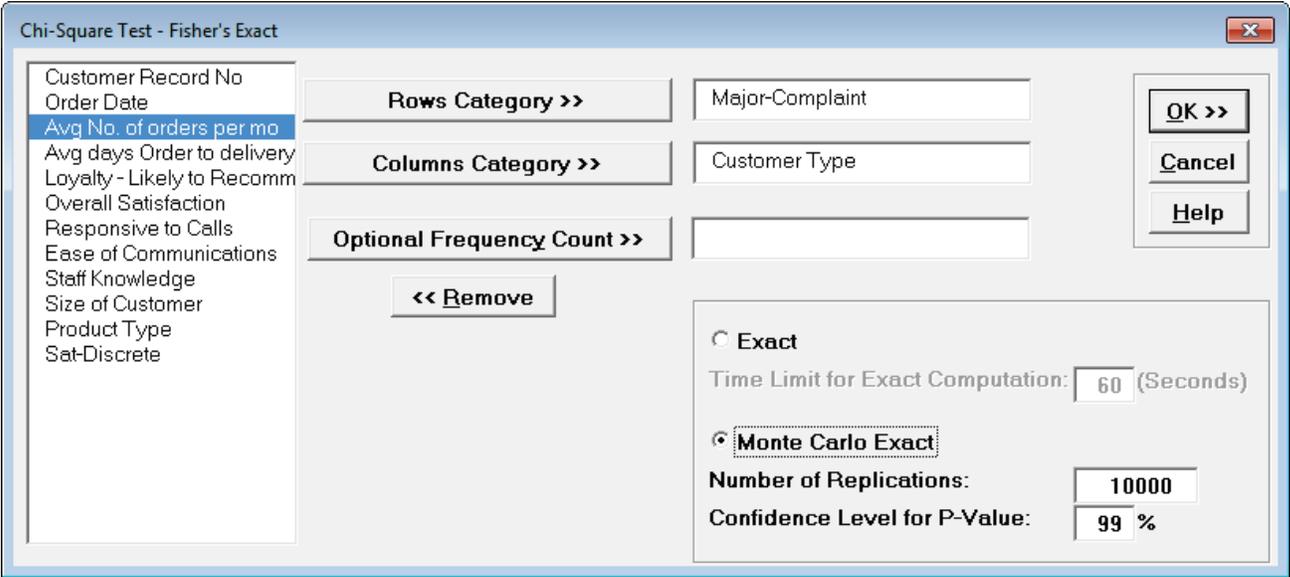
Chi-Square - Fisher's Exact			
Major-Complaint - Customer Type			
Observed Counts	1	2	3
Difficult-to-order	5	9	5
Not-available	2	0	2
Order-takes-too-long	1	3	6
Return-calls	19	28	13
Wrong-color	4	2	1
Expected Counts	1	2	3
Difficult-to-order	5.890	7.980	5.130
Not-available	1.240	1.680	1.08
Order-takes-too-long	3.100	4.200	2.700
Return-calls	18.600	25.200	16.200
Wrong-color	2.170	2.940	1.890
Std. Residuals	1	2	3
Difficult-to-order	-0.366718	0.361076	-0.057396402
Not-available	0.682500	-1.296	0.885270
Order-takes-too-long	-1.193	-0.585540	2.008
Return-calls	0.092747779	0.557773	-0.795046
Wrong-color	1.242	-0.548219	-0.647380
Chi-Square	12.211		
DF	8		
Chi-Square P-Value	0.1420		
Fisher's Exact P-Value	0.1469		
Note: 9 out of 15 cells have expected counts less than 5.			

With Fisher's Exact P-Value = 0.1469 we fail to reject H0, so we cannot conclude that there is a difference in customer complaints across customer types. This is close to the approximate Chi-Square P-Value of 0.142, so either P-Value results in the same conclusion.

However, there are cases where the "large sample" Chi-Square P-Value leads to one conclusion but Fisher's Exact P-Value gives another. This will be demonstrated later using **Two Way Table Data – Fisher's Exact**.

The Exact P-Value was computed very quickly, but if the data set was larger, the required computation time could become excessive, and Monte Carlo would be required. We will rerun this analysis with Monte Carlo and discuss the output report.

5. Press **F3** or click **Recall SigmaXL Dialog** to recall last dialog. Select **Monte Carlo Exact** with the default **Number of Replications** = 10000 and **Confidence Level for P-Value** = 99%.



Tip: As discussed above 10,000 replications will result in a Monte Carlo P-Value that is correct to two decimal places. One million (1e6) replications will result in three decimal places of accuracy and typically require less than 60 seconds to solve for any data set.

Tip: The Monte Carlo 99% confidence interval for P-Value is **not** the same as a confidence interval on the test statistic due to data sampling error. The confidence level for the hypothesis test statistic is still 95%, so **all reported P-Values less than .05 will be highlighted in red** to indicate significance. The 99% Monte Carlo P-Value confidence interval is due to the uncertainty in Monte Carlo sampling, and it becomes smaller as the number of replications increases (irrespective of the data sample size). The Exact P-Value will lie within the stated Monte Carlo confidence interval 99% of the time.

6. Click **OK**. Results:

Chi-Square Test - Fisher's Monte-Carlo			
Major-Complaint - Customer Type			
Observed Counts	1	2	3
Difficult-to-order	5	9	5
Not-available	2	0	2
Order-takes-too-long	1	3	6
Return-calls	19	28	13
Wrong-color	4	2	1
Expected Counts	1	2	3
Difficult-to-order	5.890	7.980	5.130
Not-available	1.240	1.680	1.08
Order-takes-too-long	3.100	4.200	2.700
Return-calls	18.600	25.200	16.200
Wrong-color	2.170	2.940	1.890
Std. Residuals	1	2	3
Difficult-to-order	-0.366718	0.361076	-0.057396402
Not-available	0.682500	-1.296	0.885270
Order-takes-too-long	-1.193	-0.585540	2.008
Return-calls	0.092747779	0.557773	-0.795046
Wrong-color	1.242	-0.548219	-0.647380
Chi-Square	12.211		
DF	8		
Chi-Square P-Value	0.1420		
Fisher's Monte-Carlo P-Value	0.1486		
Fisher's Monte-Carlo P-Value 99% CI Upper	0.1569		
Fisher's Monte-Carlo P-Value 99% CI Lower	0.1403		
Note: 9 out of 15 cells have expected counts less than 5.			

Fisher's Monte Carlo P-Value here is 0.1486 with a 99% confidence interval of 0.1403 to 0.1569. This will be slightly different every time it is run (the Monte Carlo seed value is derived from the system clock). The true Exact P-Value = 0.1469 lies within this confidence interval. If the confidence interval is too wide (e.g., upper "fail-to-reject" H0 and lower "reject" H0), simply rerun the Monte Carlo option using a larger number of replications (use 1e5 or 1e6).

Chi-Square Test – Two Way Table Data – Fisher’s Exact

Now we will consider a sparse data set where the Chi-Square approximation fails and Fisher’s Exact is required to give a correct conclusion for the hypothesis test. This is adapted from a subset of dental health data (oral lesions) obtained from house to house surveys that were conducted in three geographic regions of rural India [1, 2]. The Fisher’s Exact P-Value obtained with SigmaXL may be validated using these references. The data labels have been modified to a generic “A”, “B”, “C”, etc. for the oral lesions location (rows) and “Region1”, “Region2” and “Region3” for the geographic regions (columns).

1. Open the file **Oral_Lesions.xlsx**. This data is in Two –Way Contingency Table (or pivot table) format. Note that cells **A1:D10** have been pre-selected.

Site of Oral Lesion (Damage)	Region1	Region2	Region3
A	0	1	0
B	8	1	8
C	0	1	0
D	0	1	0
E	0	1	0
F	0	1	0
G	0	1	0
H	1	0	1
I	1	0	1

2. Click **SigmaXL > Statistical Tools > Chi-Square Tests – Exact > Chi-Square Test – Two-Way Table Data – Fisher’s Exact**. Note the selection of data includes the Row and Column labels.

Chi-Square Two-Way Table Data - Fisher's Exact X

Please select your data

Note:

1. Two-Way (pivot) table must be based on a single Y or count variable.
2. Selection should include only one row and only one column label.
3. Selection should exclude Grand Totals.

See Help for example.

Exact

Time Limit for Exact Computation: (Seconds)

Monte-Carlo Exact

Number of Replications:

Confidence Level for P-Value: %

3. Click **Next**. Results:

Chi-Square 2 Way Table Statistics - Fisher's Exact

Observed Counts	Region1	Region2	Region3
A	0	1	0
B	8	1	8
C	0	1	0
D	0	1	0
E	0	1	0
F	0	1	0
G	0	1	0
H	1	0	1
I	1	0	1

Expected Counts	Region1	Region2	Region3
A	0.370370	0.259259	0.370370
B	6.296	4.407	6.296
C	0.370370	0.259259	0.370370
D	0.370370	0.259259	0.370370
E	0.370370	0.259259	0.370370
F	0.370370	0.259259	0.370370
G	0.370370	0.259259	0.370370
H	0.740741	0.518519	0.740741
I	0.740741	0.518519	0.740741

Std. Residuals	Region1	Region2	Region3
A	-0.608581	1.455	-0.608581
B	0.678971	-1.623	0.678971
C	-0.608581	1.455	-0.608581
D	-0.608581	1.455	-0.608581
E	-0.608581	1.455	-0.608581
F	-0.608581	1.455	-0.608581
G	-0.608581	1.455	-0.608581
H	0.301232	-0.720082	0.301232
I	0.301232	-0.720082	0.301232

Chi-Square	22.099
DF	16
Chi-Square P-Value	0.1400
Fisher's Exact P-Value	0.0101

Note: 25 out of 27 cells have expected counts less than 5.

Note: 25 out of 27 cells have expected counts less than 5. Since more than 20% of the cells have expected counts less than 5 and several cells have an expected count less than 1, the Chi-Square P-Value = 0.14 is invalid and leads to an incorrect conclusion: Fail to Reject H0, Site of Oral Lesion and Geographic Region are independent. However, with Fisher's Exact P-Value =

0.0101, we have strong evidence to reject H_0 , and conclude that there is indeed a relationship between Site of Oral Lesion and Region.

While this example is originally from a dental health study, it could have been Defect Type versus Supplier and highlights that the use of Fisher's Exact for contingency tables with sparse data can make the difference between a good business decision and a bad business decision!

The Exact P-Value was computed very quickly, but if the data set was larger, the required computation time could become excessive, and Monte Carlo would be required. We will rerun this analysis with Monte Carlo and discuss the output report.

- Press **F3** or click **Recall SigmaXL Dialog** to recall last dialog. Select **Monte Carlo Exact** with the default **Number of Replications** = 10000 and **Confidence Level for P-Value** = 99%.

Chi-Square Two-Way Table Data - Fisher's Exact

Please select your data

\$A\$1:\$D\$10

Note:

1. Two-Way (pivot) table must be based on a single Y or count variable.
2. Selection should include only one row and only one column label.
3. Selection should exclude Grand Totals. See Help for example.

Exact
 Time Limit for Exact Computation: 60 (Seconds)

Monte-Carlo Exact
 Number of Replications: 10000
 Confidence Level for P-Value: 99 %

Tip: As discussed above 10,000 replications will result in a Monte Carlo P-Value that is correct to two decimal places. One million (1e6) replications will result in three decimal places of accuracy and typically require less than 60 seconds to solve for any data set.

Tip: The Monte Carlo 99% confidence interval for P-Value is **not** the same as a confidence interval on the test statistic due to data sampling error. The confidence level for the hypothesis test statistic is still 95%, so **all reported P-Values less than .05 will be highlighted in red** to indicate significance. The 99% Monte Carlo P-Value confidence interval is due to the uncertainty in Monte Carlo sampling, and it becomes smaller as the number of replications increases (irrespective of the data sample size). The Exact P-Value will lie within the stated Monte Carlo confidence interval 99% of the time.

- Click **Next**. Results:

Chi-Square 2 Way Table Statistics - Fisher's Monte Carlo			
Observed Counts	Region1	Region2	Region3
A	0	1	0
B	8	1	8
C	0	1	0
D	0	1	0
E	0	1	0
F	0	1	0
G	0	1	0
H	1	0	1
I	1	0	1
Expected Counts	Region1	Region2	Region3
A	0.370370	0.259259	0.370370
B	6.296	4.407	6.296
C	0.370370	0.259259	0.370370
D	0.370370	0.259259	0.370370
E	0.370370	0.259259	0.370370
F	0.370370	0.259259	0.370370
G	0.370370	0.259259	0.370370
H	0.740741	0.518519	0.740741
I	0.740741	0.518519	0.740741
Std. Residuals	Region1	Region2	Region3
A	-0.608581	1.455	-0.608581
B	0.678971	-1.623	0.678971
C	-0.608581	1.455	-0.608581
D	-0.608581	1.455	-0.608581
E	-0.608581	1.455	-0.608581
F	-0.608581	1.455	-0.608581
G	-0.608581	1.455	-0.608581
H	0.301232	-0.720082	0.301232
I	0.301232	-0.720082	0.301232
Chi-Square	22.099		
DF	16		
Chi-Square P-Value	0.1400		
Fisher's Monte Carlo P-Value	0.0094		
Fisher's Monte Carlo P-Value 99% CI Upper	0.0116		
Fisher's Monte Carlo P-Value 99% CI Lower	0.0072		

Fisher's Monte Carlo P-Value here is 0.0094 with a 99% confidence interval of 0.0072 to 0.0116. This will be slightly different every time it is run (the Monte Carlo seed value is derived from the

system clock). The true Exact P-Value = 0.0101 lies within this confidence interval. If the confidence interval is too wide (e.g., upper is a “fail-to-reject” H0 and lower is a “reject” H0), simply rerun the Monte Carlo option using a larger number of replications (use 1e5 or 1e6).

References for Fisher’s Exact

[1] Mehta, C. R.; Patel, N. R. (1997) "Exact inference in categorical data," unpublished preprint, <http://www.cytel.com/Papers/sxpaper.pdf>. See Table 7 (p. 33) for validation of Fisher’s Exact P-Value.

[2] Mehta, C.R. ; Patel, N.R. (1998). "Exact Inference for Categorical Data." In P. Armitage and T. Colton, eds., *Encyclopedia of Biostatistics*, Chichester: John Wiley, pp. 1411–1422.

Power & Sample Size for One Proportion Test

To determine Power & Sample Size for a 1 Proportion Test, you can use the Power & Sample Size Calculator or Power & Sample Size with Worksheet.

1. Click **SigmaXL > Statistical Tools > Power & Sample Size Calculators > 1 Proportion Test Calculator**.
2. Select **Solve For Power (1 – Beta)**. Enter **Sample Size** and **Alternative Proportion** as shown:

Note that we are calculating the power or likelihood of detection given that the hypothesized proportion is 0.5, but the alternative proportion is 0.6, sample size = 300, significance level = .05, and Ha: Not Equal To (two-sided test).

3. Click **OK**. The resulting report is displayed:

Power and Sample Size: 1 Proportion Test				
H0: P0 = 0.5				
Ha: P0 ≠ 0.5				
Solve For: Power (1 - Beta)				
Sample Size (N)	Alternative Proportion (P1)	Hypothesized Proportion (P0)	Significance Level (Alpha)	Power (1 - Beta)
300	0.6	0.5	0.05	0.937627018

A power value of 0.94 is acceptable, but note that the sample size $n = 300$, and the difference in proportion value is 0.1 or 10%! The sample size requirements for discrete data are much higher than those for continuous data.

4. To determine Power & Sample Size using a Worksheet, click **SigmaXL > Statistical Tools > Power & Sample Size with Worksheets > 1 Proportion Test**.
5. A graph showing the relationship between Power, Sample Size and Proportion Value can then be created using **SigmaXL > Statistical Tools > Power & Sample Size Chart**. See Part E for an example using the 1 Sample t-Test.

Power & Sample Size for Two Proportions Test

To determine Power & Sample Size for a 2 Proportions Test, you can use the Power & Sample Size Calculator or Power & Sample Size with Worksheet.

1. Click **SigmaXL > Statistical Tools > Power & Sample Size Calculators > 2 Proportions Test Calculator**.
2. Select **Solve For Power (1 – Beta)**. Enter **Sample Size** and **Proportion 1** values as shown:

Note that we are calculating the power or likelihood of detection given that $P1 = 0.5$ and $P2 = 0.6$, sample size for each group = 300, significance level = .05, and H_a : Not Equal To (two-sided test).

3. Click **OK**. The resulting report is displayed:

Power and Sample Size: 2 Proportions Test				
H0: P1 = P2				
Ha: P1 ≠ P2				
Solve For: Power (1 - Beta)				
Sample Size (N)	Proportion 1 (P1)	Proportion 2 (P2)	Significance Level (Alpha)	Power (1 - Beta)
300	0.6	0.5	0.05	0.693021232

A power value of 0.69 is unacceptable. Note that this value is much less than the power for the one proportion test (0.94).

4. To compensate, we will double the sample size per group. Press **F3** or click **Recall SigmaXL Dialog** to recall last dialog. Change the sample size per group from 300 to 600. Note that the power value is now 0.94.
5. To determine Power & Sample Size using a Worksheet, click **SigmaXL > Statistical Tools > Power & Sample Size with Worksheets > 2 Proportions Test**.
6. A graph showing the relationship between Power, Sample Size and Proportion Values can then be created using **SigmaXL > Statistical Tools > Power & Sample Size Chart**. See Part E for an example using the 1 Sample t-Test.

Part P – Analysis of Means (ANOM) Charts

Introduction – What is ANOM

From <http://asq.org/glossary/a.html>:

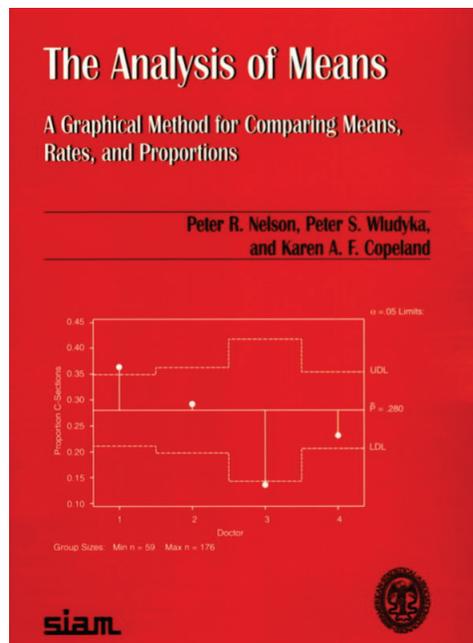
A statistical procedure for troubleshooting industrial processes and analyzing the results of experimental designs with factors at fixed levels. It provides a graphical display of data. Ellis R. Ott developed the procedure in 1967 because he observed that nonstatisticians had difficulty understanding analysis of variance. Analysis of means is easier for quality practitioners to use because it is (like) an extension of the control chart.

From the preface of the book, “The Analysis of Means: A Graphical Method for Comparing Means, Rates, and Proportions” by Peter R. Nelson, Peter S. Wludyka, and Karen A. F. Copeland:

The goal of statistical data analysis is to use data to gain and communicate knowledge about processes and phenomena. Comparing means is often part of an analysis, for data arising in both experimental and observational studies.

The analysis of means (ANOM) is an alternative procedure (to ANOVA) for comparing means.

ANOM has the advantages of being much more intuitive and providing an easily understood graphical result, which clearly indicates any means that are different (from the overall mean) and allows for easy assessment of practical as well as statistical significance.



Formula for One-Way Balanced* Normal (from ANOM book 2.4)

$$\begin{aligned} \text{UDL} &= \bar{y}_{..} + h(\alpha; I, N - I) \sqrt{MS_e} \sqrt{\frac{I-1}{N}} \\ \text{LDL} &= \bar{y}_{..} - h(\alpha; I, N - I) \sqrt{MS_e} \sqrt{\frac{I-1}{N}} \end{aligned}$$

where:

$\bar{y}_{..}$ = overall mean

h = critical value from multivariate t distribution – SigmaXL uses table exact critical values (Table B.1)

N = sample size

I = number of levels

$\sqrt{MS_e}$ SQRT (Mean Square Error) = pooled standard deviation.

* Unbalanced uses critical values from studentized maximum modulus (SMM) distribution. SigmaXL uses table exact critical values (Table B.2). An adjustment is also made for varying sample size that results in varying decision limit values. See Appendix [Analysis of Means \(ANOM\) Charts](#).

Formula for One-Way Balanced Binomial Proportions (from ANOM book 2.6)

$$\begin{aligned} \text{UDL} &= \bar{p} + h(\alpha; I, \infty) \sqrt{\bar{p}(1 - \bar{p})} \sqrt{\frac{I-1}{N}} \\ \text{LDL} &= \bar{p} - h(\alpha; I, \infty) \sqrt{\bar{p}(1 - \bar{p})} \sqrt{\frac{I-1}{N}} \end{aligned}$$

This is based on the normal approximation to the Binomial, so requires sample sizes large enough that np and $n(p-1) > 5$. SigmaXL automatically checks this and gives a warning message if the condition is not met.

Formula for One-Way Balanced Poisson Rates (from ANOM book 2.9)

$$\begin{aligned} \text{UDL} &= \bar{u} + h(\alpha; I, \infty) \sqrt{\bar{u}} \sqrt{\frac{I-1}{N}} \\ \text{LDL} &= \bar{u} - h(\alpha; I, \infty) \sqrt{\bar{u}} \sqrt{\frac{I-1}{N}} \end{aligned}$$

This is based on the normal approximation to the Poisson, so requires sample sizes large enough that $nu > 5$. SigmaXL automatically checks this and gives a warning message if the condition is not met.

Formula for Nonparametric Transformed Ranks (from ANOM book 9.3)

The values from the combined sample are ranked from smallest to largest (R_{ij}) and transformed to a z-score using the inverse normal:

$$E_{ij} = \Phi^{-1}[0.5 + R_{ij}/(2N + 1)]$$

and the above formula for balanced (or unbalanced) Normal is applied to compute decision limits.

Formula for Balanced Variances (from ANOM book 4.4)

$$\begin{aligned} \text{UDL} &= U(\alpha; I, n - 1) \text{IMS}_e \\ \text{CL} &= \text{MS}_e \\ \text{LDL} &= L(\alpha; I, n - 1) \text{IMS}_e \end{aligned}$$

where U, L are critical values from Table B.4 ($n \leq 35$). For unbalanced and large sample size see Appendix [Analysis of Means \(ANOM\) Charts](#).

Formula for Levene Robust Variances (from ANOM book 4.16)

The Levene Absolute Deviation from Median (ADM) is a simple transformation:

$$y_{ij} = |x_{ij} - \tilde{x}_i|$$

and the above formula for balanced (or unbalanced) Normal is applied to compute decision limits.

Two-Way Normal

Main Effects for Two-Way ANOM are similar to One-Way but the mean square error (MSE) is derived from the ANOVA.

Slice Charts are a modified ANOM chart developed by Dr. Peter Wludyka that enables one to easily interpret the effects in the presence of an interaction (Wludyka 2013, 2015). The basic idea is to compare the levels of one factor for each level of the other factor with MSE derived from the Two-Way ANOVA. A yellow highlight automatically recommends Main Effects (if interaction is not significant, i.e., P-Value ≥ 0.1) or Slice Chart (if interaction is significant, P-Value < 0.1). The interaction P-Value is determined from ANOVA

There is an option to specify correction to alpha for multiple chart family-wise error rate:

Bonferroni $\alpha' = \alpha/m$; m = number of charts

Note that this is strictly to adjust for multiple charts. As seen in the formulas above, the individual charts also have correction built-in for the number of levels.

A Two-Way ANOM analysis will typically have balanced data from a designed experiment, with an equal number of observations for each combination level of X1 and X2. SigmaXL will allow unbalanced data, but we recommend that you use balanced data or slightly unbalanced data. ANOM charts use actual data means, not fitted (predicted least squares) means as used in Two-Way ANOVA Main Effects and Interaction Plots.

Two-Way Binomial Proportions and Poisson Rates

In collaboration, Peter Wludyka and John Noguera of SigmaXL extended the Slice Charts to Binomial and Poisson (Wludyka and Noguera 2016).

As with Normal, the basic idea is to compare the levels of one factor for each level of the other factor. MSE is derived from the whole model

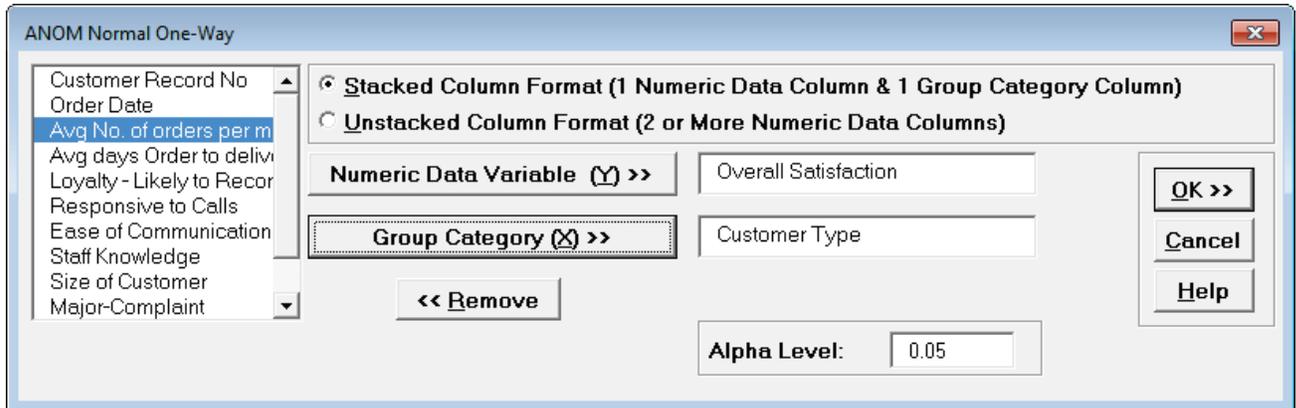
A yellow highlight automatically recommends Main Effects (if interaction is not significant) or Slice Chart (if interaction is significant).

The interaction P-Value is automatically determined from Logistic regression for Binomial Proportions and Poisson regression for Poisson Rates.

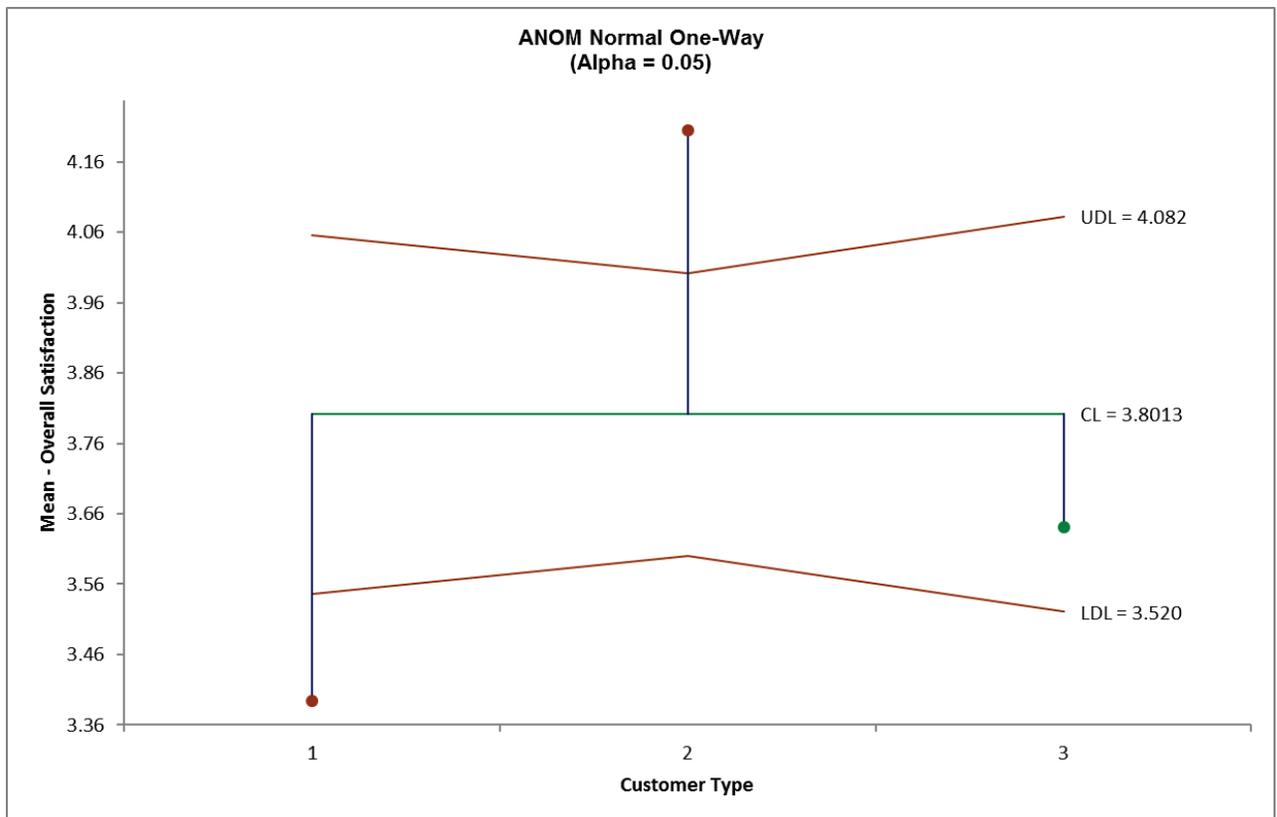
See Appendix [Analysis of Means \(ANOM\) Charts](#) for further formula details and references.

ANOM Normal One-Way

1. Open **Customer Data.xlsx**, click on **Sheet 1** tab.
2. Click **SigmaXL > Graphical Tools > Analysis of Means (ANOM) > ANOM Normal One-Way**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**.
3. Click **Next**. Ensure that **Stacked Column Format** is checked. Select *Overall Satisfaction*, click **Numeric Data Variable (Y) >>**; select *Customer Type*, click **Group Category (X) >>**. Set **Alpha Level = 0.05**.



4. Click **OK**. The ANOM Normal One-Way chart is shown below:



5. Here we see that Customer Type 1 mean satisfaction score is significantly below the overall mean and Customer Type 2 is significantly higher. This is consistent with the results that were observed in the ANOVA analysis, but is easier to interpret.
6. The varying decision limits are due to the varying sample sizes for each Customer Type, with smaller sample size giving wider limits in a manner similar to a control chart. If the data are balanced the decision limit lines will be constant.

ANOM Normal Two-Way (with Main Effects and Slice Charts)

1. We will look at two balanced examples from the ANOM book (used with author permission), one with no interaction and another with a strong interaction, and then a slightly unbalanced version of the latter (Wludyka, 2013).
2. Open **ANOM Examples.xlsx**, click on **Hemoglobin – Normal** tab. This data is from a factorial design to study hemoglobin levels in males (Example 5.1). The factors are Therapy (Chemo or Radiation) and Drug (Type 1, 2, 3).
3. Click **SigmaXL > Graphical Tools > Analysis of Means (ANOM) > ANOM Normal Two-Way**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**.
4. Click **Next**. Select *Hemoglobin Level*, click **Numeric Data Variable (Y) >>**; select *Therapy*, click **Group Category Factor (X1) >>**; select *Drug*, click **Group Category Factor (X2) >>**. **Alpha Level = 0.05**, Adjust chart alpha for family-wise error rate is unchecked:

ANOM Normal Two-Way

Numeric Data Variable (Y) >> Hemoglobin Level

Group Category Factor (X1) >> Therapy

Group Category Factor (X2) >> Drug

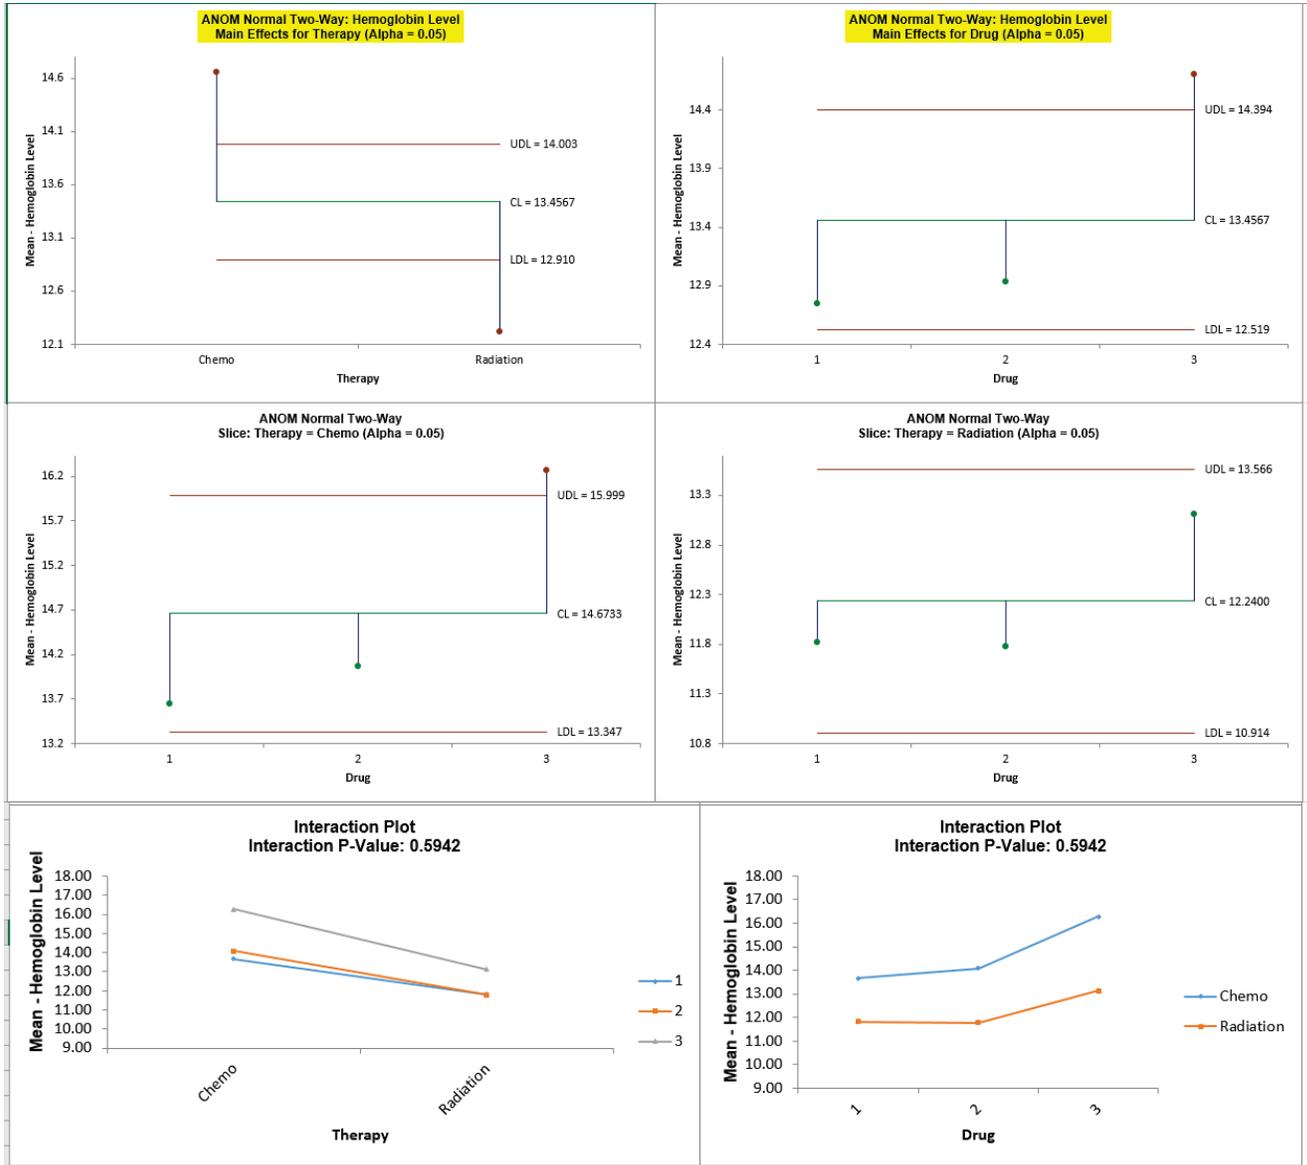
<< Remove

Alpha Level: 0.05

Adjust chart alpha for family-wise error rate

OK >> Cancel Help

5. Click **OK**. The ANOM Normal Two-Way Main Effects, Slice charts and Interaction plots are shown below:



6. Since the Interaction is not significant (P-Value is ≥ 0.1), the Main Effects charts are highlighted. The decision limits match the manual calculations given in the ANOM book (pages 87, 88). Since the average hemoglobin level for radiation falls below the lower decision limit, and the average for chemotherapy falls above the upper decision limit, therapy has an ($\alpha = 0.05$) effect. There is also an effect on hemoglobin level due to the drug since the average hemoglobin level for drug 3 is above the upper decision line.

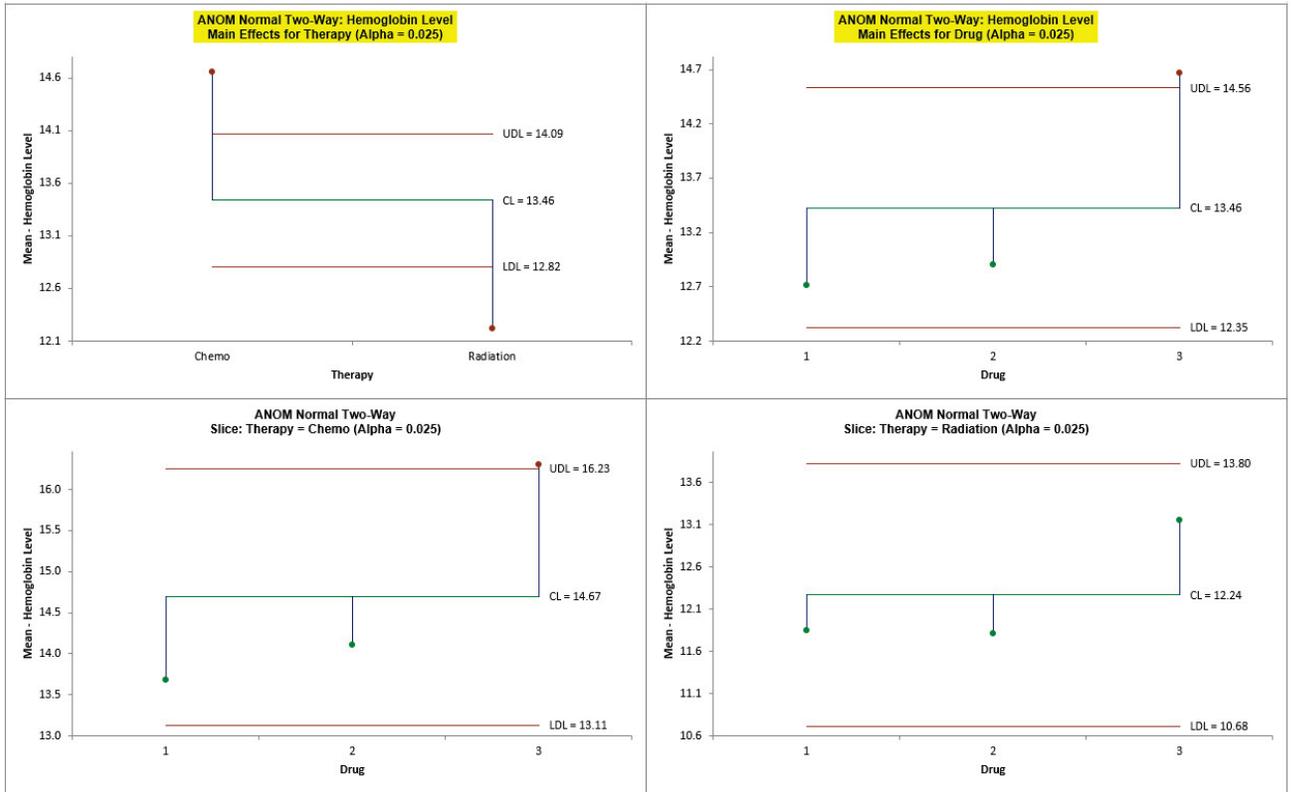
Tip: Two-Way ANOM Normal is complimentary to Two-Way ANOVA. Analyzing the above data using Two-Way ANOVA shows that Therapy is significant with a P-Value = 0.0001, Drug is also significant with a P-Value = 0.0105 and the Interaction is not significant, with a P-Value = 0.5942.

7. Click **Recall SigmaXL Dialog** menu or press **F3** to recall last dialog. Check **Adjust chart alpha for family-wise error rate**.

The screenshot shows the 'ANOM Normal Two-Way' dialog box. On the left is an empty list box. To its right are three rows of controls: 'Numeric Data Variable (Y) >>' with 'Hemoglobin Level' in the text box; 'Group Category Factor (X1) >>' with 'Therapy' in the text box; and 'Group Category Factor (X2) >>' with 'Drug' in the text box. Below these is a '<< Remove' button. To the right of the list box are three buttons: 'OK >>', 'Cancel', and 'Help'. At the bottom right, there is an 'Alpha Level:' label with a text box containing '0.05' and a checked checkbox labeled 'Adjust chart alpha for family-wise error rate'.

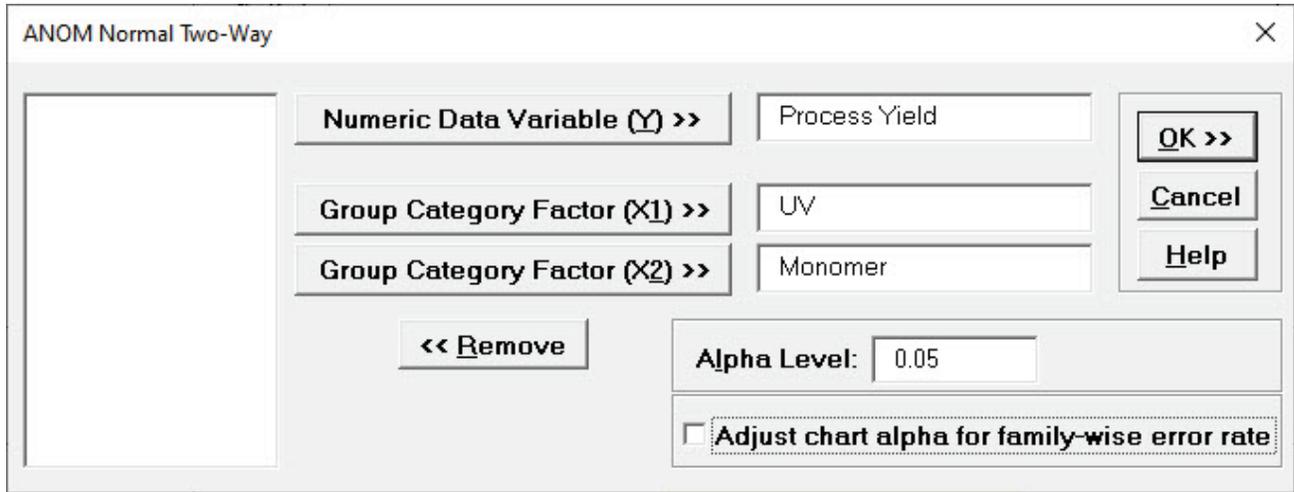
This will apply a Bonferroni correction to the specified alpha values to control the overall family-wise error rate for the charts. So, the Main Effects alpha will be $= 0.05/2 = 0.025$ and the Slice Charts alpha will be $= 0.05/2 = 0.025$ as well.

8. Click **OK**. The ANOM Normal Two-Way Main Effects and Slice charts are shown below:

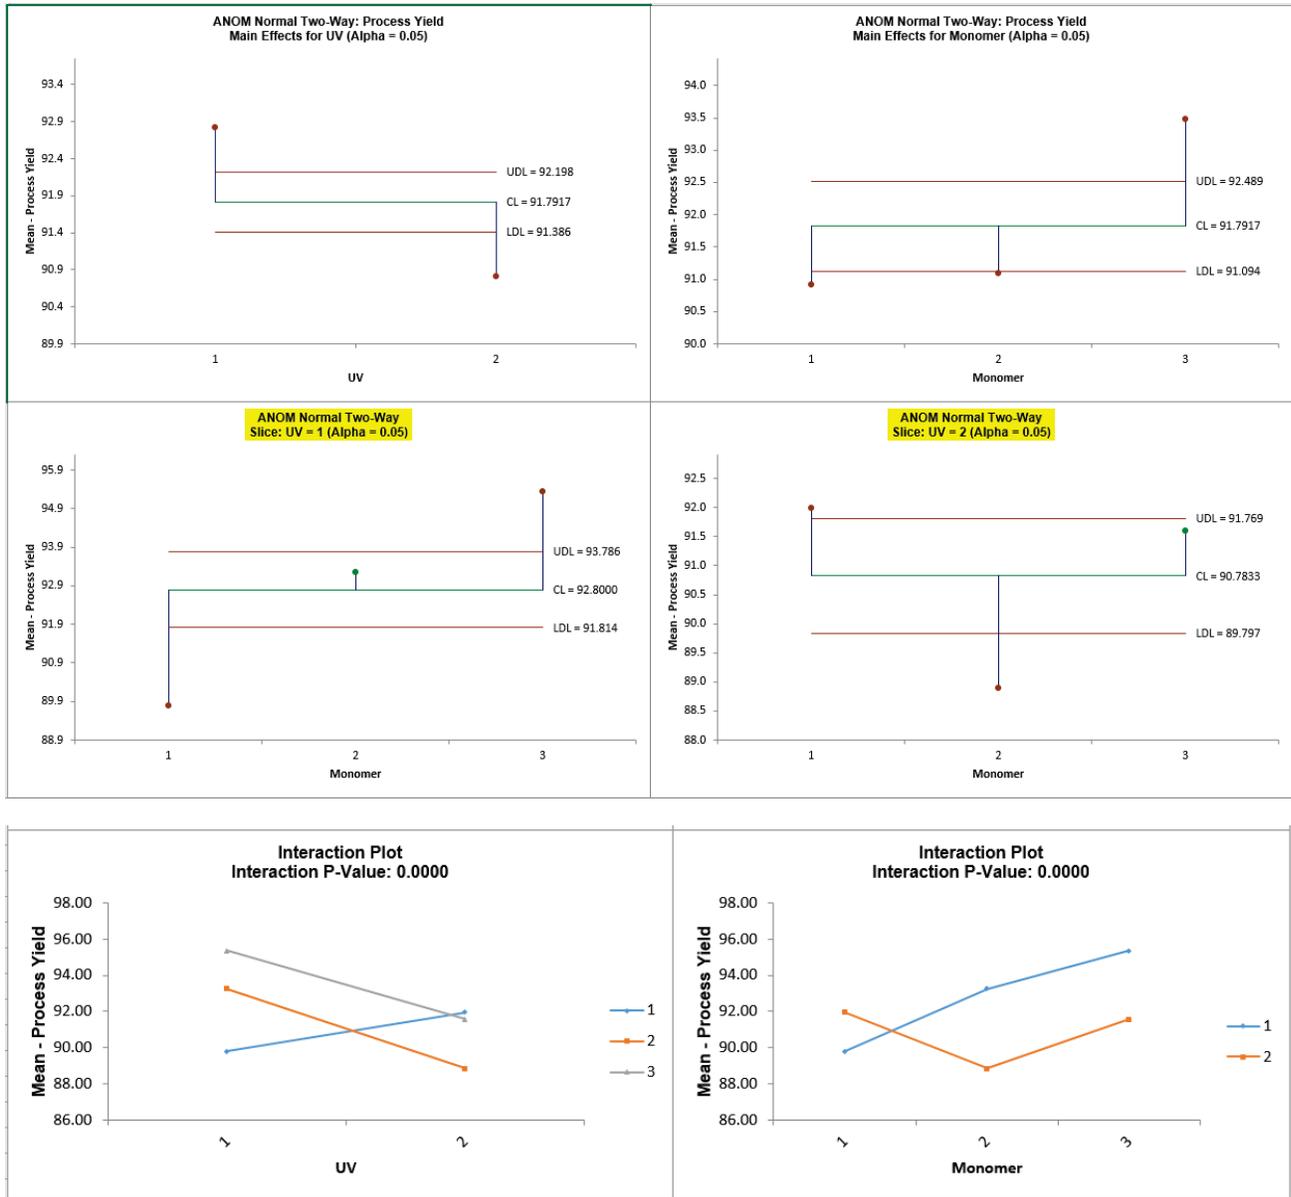


9. With the Bonferroni corrected alpha of 0.025, the significant results do not change, but they are closer to the decision lines. The trade-off for this chart family-wise adjustment is loss of power.

10. In **ANOM Examples.xlsx**, click on **Process Yield – Normal** tab. This data is from a factorial design to study the effect of three monomers and two levels of UV exposure (1 = UV; 2 = No UV) on the percent yield for contact lens production (Example 5.3).
11. Click **SigmaXL > Graphical Tools > Analysis of Means (ANOM) > ANOM Normal Two-Way**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**.
12. Click **Next**. Select *Process Yield*, click **Numeric Data Variable (Y) >>**; select *UV*, click **Group Category Factor (X1) >>**; select *Monomer*, click **Group Category Factor (X2) >>**. **Alpha Level = 0.05**, **Adjust chart alpha for family-wise error rate** is unchecked:



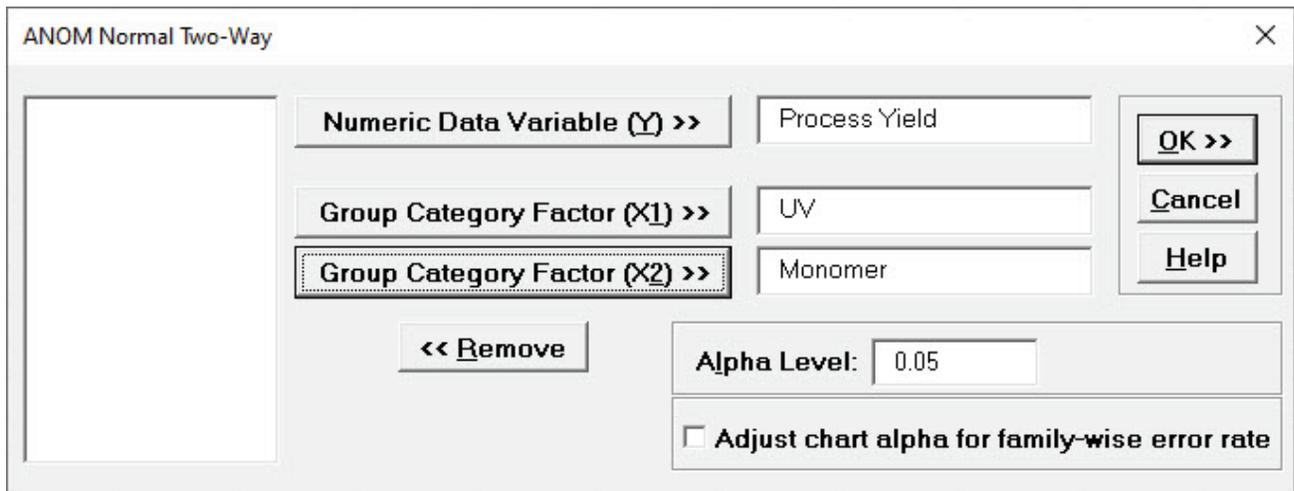
13. Click **OK**. The ANOM Normal Two-Way Main Effects, Slice charts and Interaction plots are shown below:



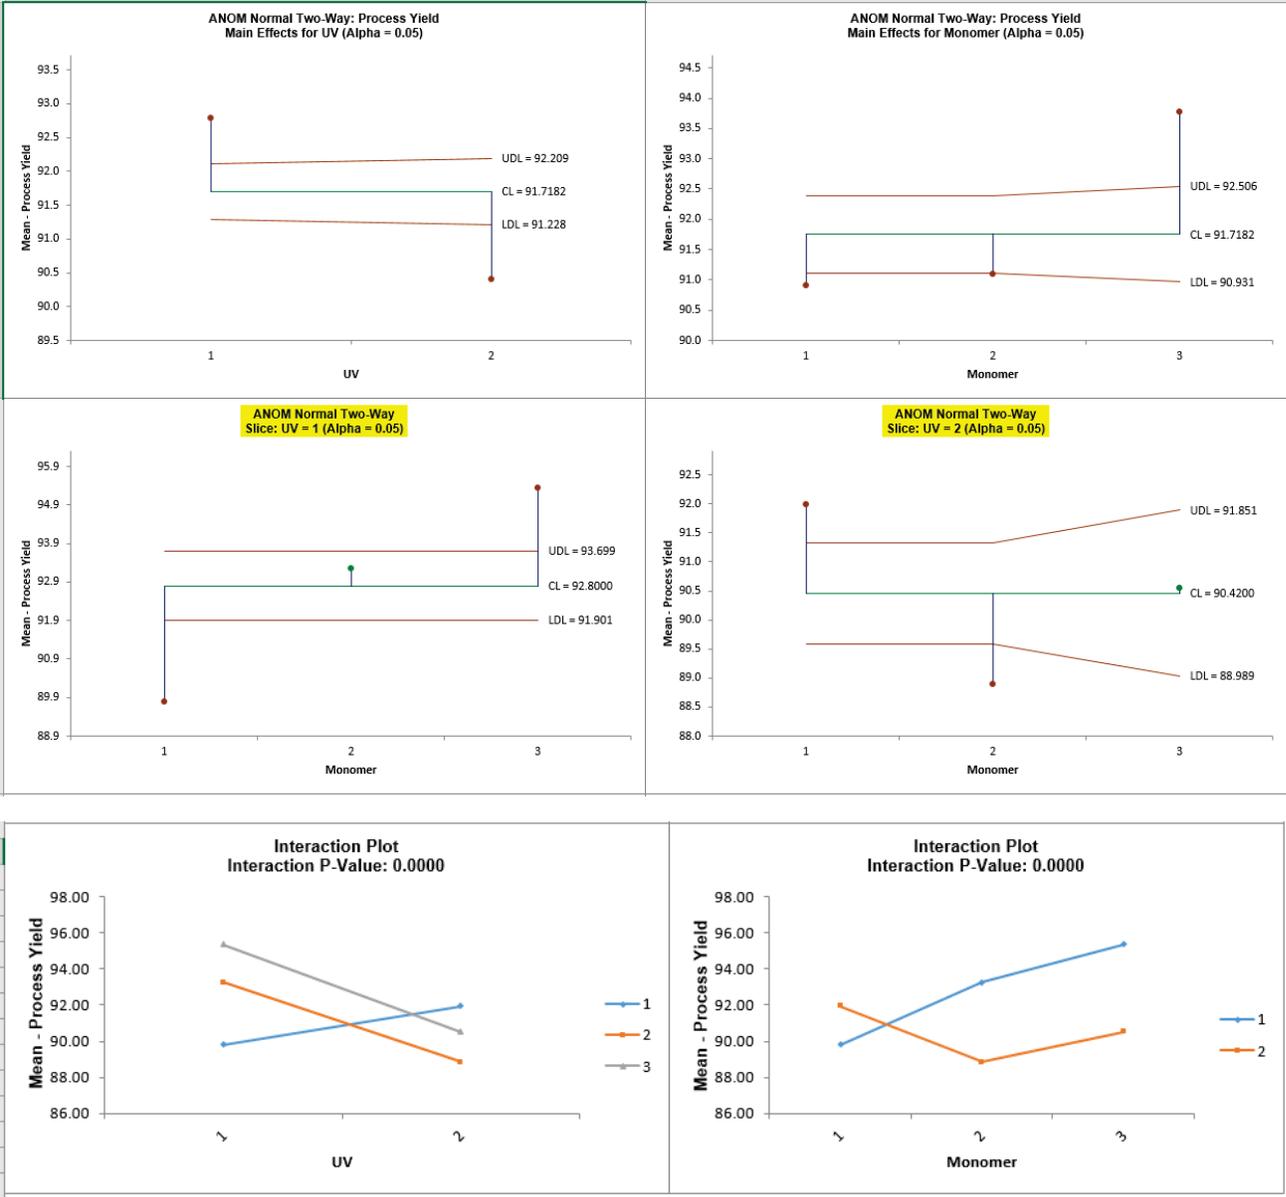
14. Since the Interaction P-Value is < 0.1 , the Slice charts are highlighted. The decision limits for the Slice charts match the calculations given in Wludyka (2013). See Appendix [References for Analysis of Means \(ANOM\) Charts](#). The Slice plot clearly shows the effect of the strong interaction. When UV = 1 (UV Exposure), the three monomers differ. In particular, the average for monomer 1 is below the overall average and the average for monomer 3 is above the overall average. However, for UV = 2 (no UV), the monomers also differ but in an entirely different manner than in the UV = 1 case, revealing the nature of the interaction.

Tip: As previously noted, Two-Way ANOM Normal is complimentary to Two-Way ANOVA. Analyzing the above data using Two-Way ANOVA shows that UV is significant with a P-Value = 0.0001, Monomer is significant with a P-Value = 0.0000 (4.4 e-5) and the Interaction is also significant, with P-Value = 0.0000 (2.2e-6).

15. In **ANOM Examples.xlsx**, click on **Process Yield – Normal Unbal** tab. This data is from a factorial design to study the effect of three monomers and two levels of UV exposure (1 = UV; 2 = No UV) on the percent yield for contact lens production, but two of the observations are missing for A = 2 (no UV) and B = 3 (monomer 3).
16. Click **SigmaXL > Graphical Tools > Analysis of Means (ANOM) > ANOM Normal Two-Way**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**.
17. Click **Next**. Select *Process Yield*, click **Numeric Data Variable (Y) >>**; select *UV*, click **Group Category Factor (X1) >>**; select *Monomer*, click **Group Category Factor (X2) >>**. **Alpha Level = 0.05**, **Adjust chart alpha for family-wise error rate** is unchecked:



18. Click **OK**. The ANOM Normal Two-Way Main Effects, Slice charts and Interaction plots are shown below:



19. The significant results of the experiment have not changed, but the decision limits are now varying due to the unbalanced data and use critical values for unbalanced data. The decision limits for the Slice charts match the calculations given in Wludyka (2013).

ANOM Binomial Proportions One-Way

1. Open **ANOM Examples.xlsx**, click on the **Test Scores – Binomial Prop** tab. This is standardized math test score data from 10 elementary schools (Example 3.4 from the ANOM book – used with author permission). We are testing to see if there is a difference between schools at an $\alpha = 0.01$ level.

School	Proficient	Enrollment	Proportion
1	45	52	0.8654
2	15	55	0.2727
3	89	105	0.8476
4	75	81	0.9259
5	40	46	0.8696
6	67	84	0.7976
7	13	18	0.7222
8	28	33	0.8485
9	17	48	0.3542
10	21	41	0.5122

2. Click **SigmaXL > Graphical Tools > Analysis of Means (ANOM) > ANOM Binomial Proportions One-Way**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**.
3. Click **Next**. Select *Proficient*, click **Numeric Data Variable (Y) >>**; select *Enrollment*, click **Subgroup Column or Size >>**; select *School*, click **Optional Group Category (X) >>**. Set **Alpha Level = 0.01**:

ANOM Binomial Proportions One-Way

Proportion

Numeric Data Variable (Y) >> Proficient

Subgroup Column or Size >> Enrollment

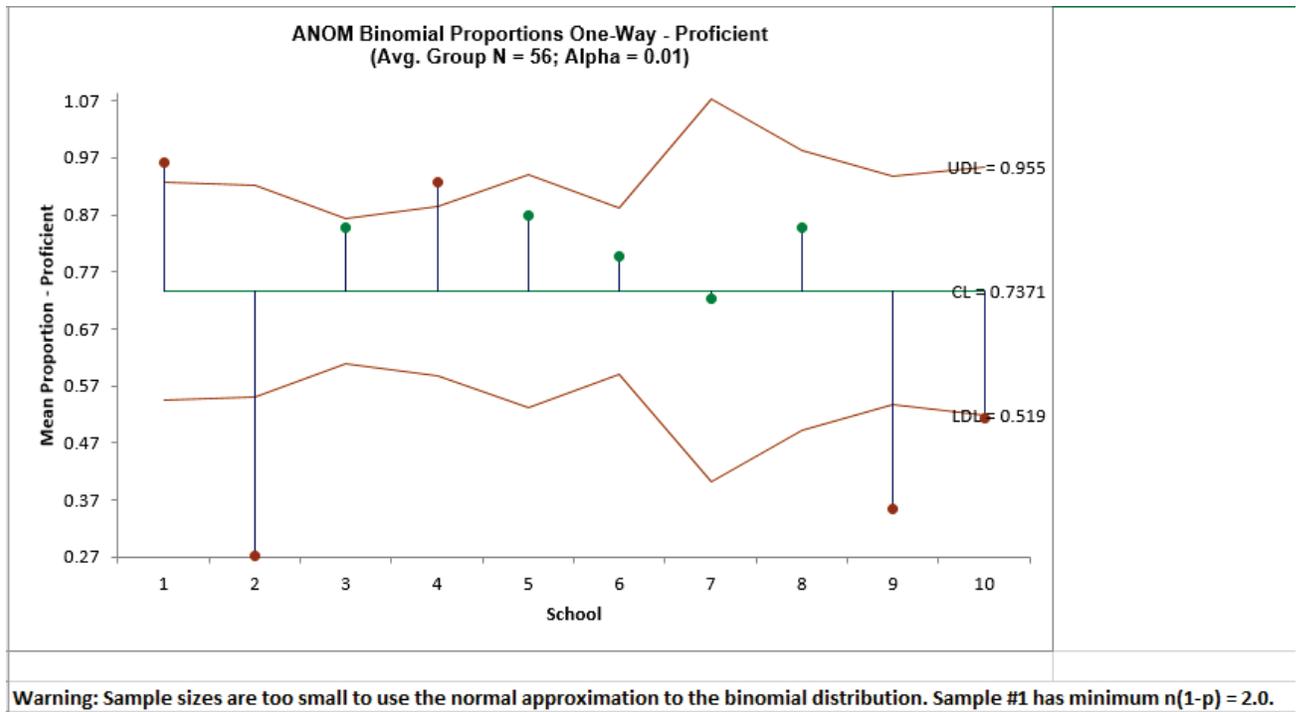
Optional Group Category (X) >> School

<< Remove

Alpha Level: 0.01

OK >> Cancel Help

4. Click **OK**. The ANOM Binomial Proportions One-Way chart is shown below:



5. The resulting ANOM decision chart shows that three schools are performing at significantly low levels and two schools are performing at significantly high levels.
6. SigmaXL automatically checks to see if the sample sizes are large enough for the normal approximation to the Binomial to be valid, i.e., np and $n(p-1)$ are > 5 . Here we see: **Warning: Sample sizes are too small to use the normal approximation to the binomial distribution. Sample #1 has minimum $n(1-p) = 2.0$.** Note that the warning does not show all occurrences, just the sample(s) with smallest failed np or $n(1-p)$.
7. This does not mean that the chart results are invalid, an obvious “out” or “in” will not likely be affected, but the results should be used with caution, and if possible, more data collected.

ANOM Binomial Proportions Two-Way

1. Open **ANOM Examples.xlsx**, click on the **Length of Stay - Binomial Prop** tab. This is hospital length of stay (LOS > 6 days) after bypass surgery. The data is from a factorial design with Factors: Drug (A, B) and Treatment (1, 2, 3). (Example 5.15 from the ANOM book – used with author permission).

Drug	Treatment	LOS > 6 days	Sample Size
A	1	10	50
A	2	6	50
A	3	20	50
B	1	15	50
B	2	16	50
B	3	12	50

2. Click **SigmaXL > Graphical Tools > Analysis of Means (ANOM) > ANOM Binomial Proportions Two-Way**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**.
3. Click **Next**. Select *LOS > 6 days*, click **Numeric Data Variable (Y) >>**; select *Sample Size*, click **Subgroup Column or Size >>**; select *Drug*, click **Group Category Factor (X1) >>**; select *Treatment*, click **Group Category Factor (X2) >>**. **Alpha Level = 0.05**, **Adjust chart alpha for family-wise error rate** is unchecked:

ANOM Binomial Proportions Two-Way

Drug
Treatment
LOS > 6 days
Sample Size

Numeric Data Variable (Y) >> LOS > 6 days

Subgroup Column or Size >> Sample Size

Group Category Factor (X1) >> Drug

Group Category Factor (X2) >> Treatment

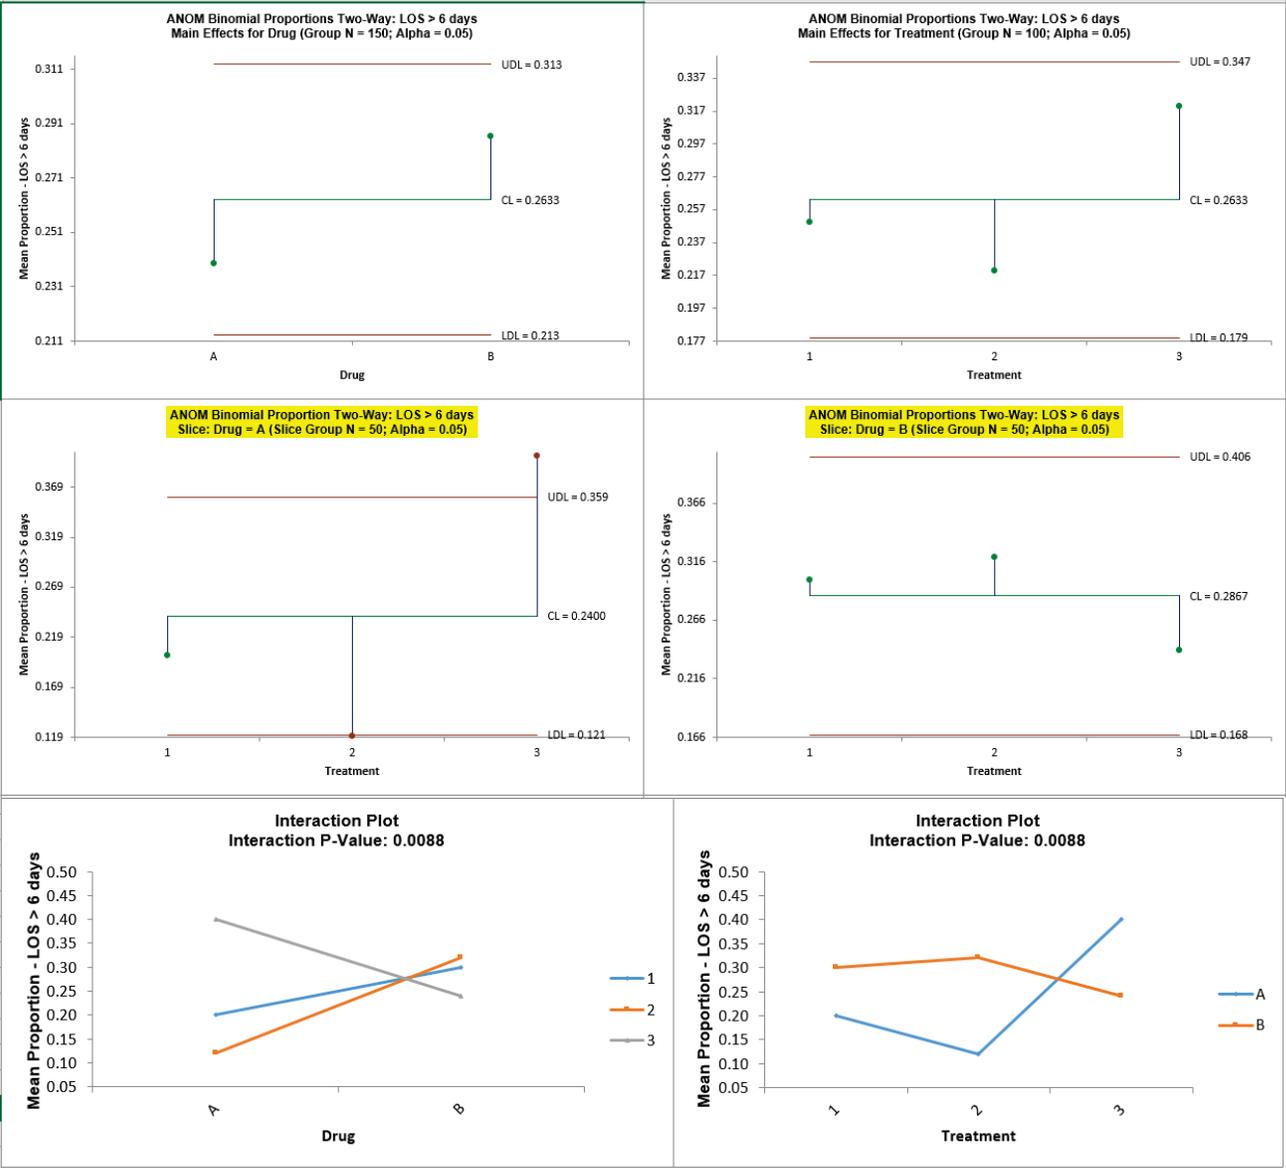
<< Remove

Alpha Level: 0.05

Adjust chart alpha for family-wise error rate

OK >>
Cancel
Help

4. Click **OK**. The ANOM Binomial Proportions Two-Way Main Effects, Slice charts and Interaction plots are shown below:



5. Since the Interaction P-Value is < 0.1 (automatically determined from Logistic regression), the Slice charts are highlighted. The Slice plot clearly shows the effect of the strong interaction. When Drug = A, Treatment 2 results in a significantly lower mean proportion LOS and Treatment 3 results in a significantly higher LOS. When Drug = B, the Treatment does not have a significant effect on LOS.

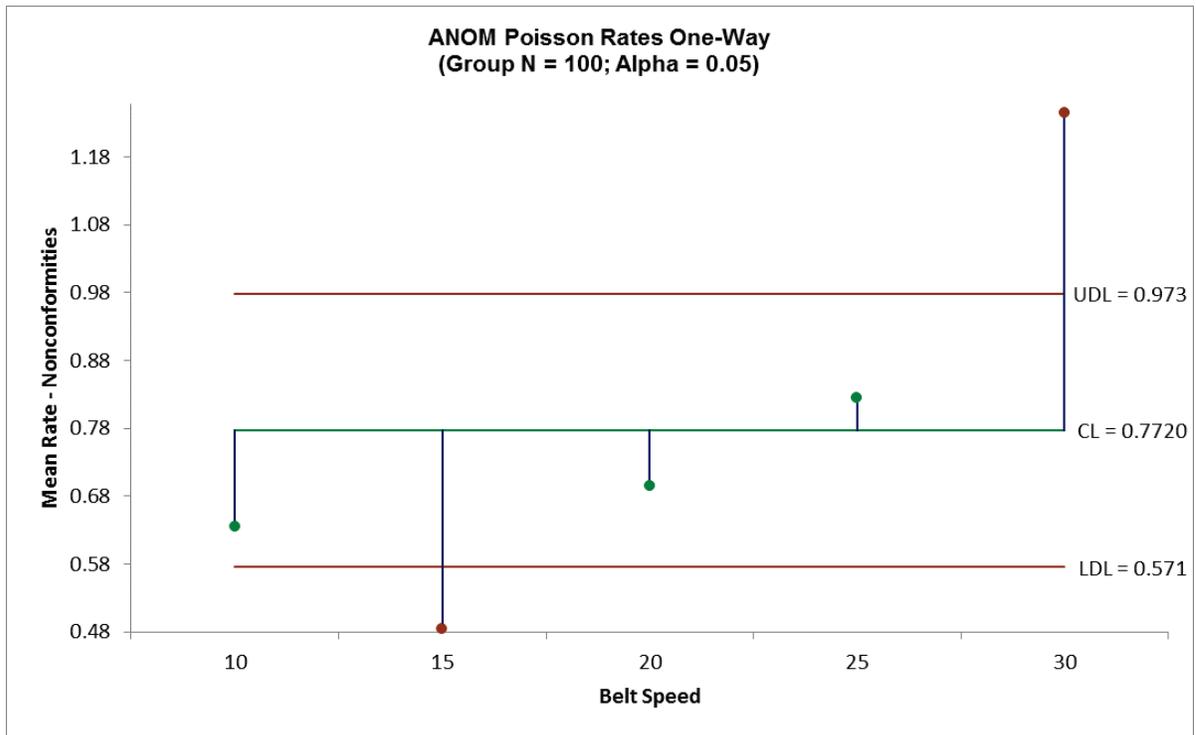
ANOM Poisson Rates One-Way

1. Open **ANOM Examples.xlsx**, click on the **Injection Molding-Poisson Rate** tab. This is injection molding data for contact lenses. It is from a study on the extent to which belt speed (cure time) influences the incidence of nonconformities (such as bubbles and tears) in the lenses. Since there can be more than one nonconformity per unit, Poisson ANOM is appropriate, similar to the U Control Chart (Example 2.16 from the ANOM book – used with author permission).

Belt Speed	Nonconformities	Sample Size
10	63	100
15	48	100
20	69	100
25	82	100
30	124	100

2. Click **SigmaXL > Graphical Tools > Analysis of Means (ANOM) > ANOM Poisson Rates One-Way**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**.
3. Click **Next**. Select *Nonconformities*, click **Numeric Data Variable (Y) >>**; enter 100 (or select *Sample Size*) for **Subgroup Column or Size >>**; select *Belt Speed*, click **Optional Group Category (X) >>**. Set **Alpha Level** = 0.05:

4. Click **OK**. The ANOM Poisson Rates One-Way chart is shown below:



5. From the ANOM chart, one sees that a belt speed of 15 results in a significantly low number of nonconformities and a belt speed of 30 results in a significantly high number of nonconformities.

ANOM Poisson Rates Two-Way

1. Open **ANOM Examples.xlsx**, click on the **ER Visits – Poisson Rate** tab. This is Emergency Room Visits Data. It is a factorial design study to investigate the effect of nurse intervention (1 = home nurse assigned, 2 = home nurse not assigned) and payor groups (1, 2, 3) on the ER utilization rates for patients with COPD (a serious lung disease) with an alpha level = 0.01 (Example 5.16 from the ANOM book – used with author permission).

Intervention	Payor	ER Visits	Sample Size
1	1	10	50
1	2	30	50
1	3	14	50
2	1	30	50
2	2	60	50
2	3	40	50

2. Click **SigmaXL > Graphical Tools > Analysis of Means (ANOM) > ANOM Poisson Rates Two-Way**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**.
3. Click **Next**. Select *ER Visits*, click **Numeric Data Variable (Y) >>**; select *Sample Size*, click **Subgroup Column or Size >>**; select *Intervention*, click **Group Category Factor (X1) >>**; select *Payor*, click **Group Category Factor (X2) >>**. **Alpha Level = 0.01**, **Adjust chart alpha for family-wise error rate** is unchecked:

ANOM Poisson Rates Two-Way

Intervention
Payor
ER Visits
Sample Size

Numeric Data Variable (Y) >>
Subgroup Column or Size >>
Group Category Factor (X1) >>
Group Category Factor (X2) >>

<< Remove

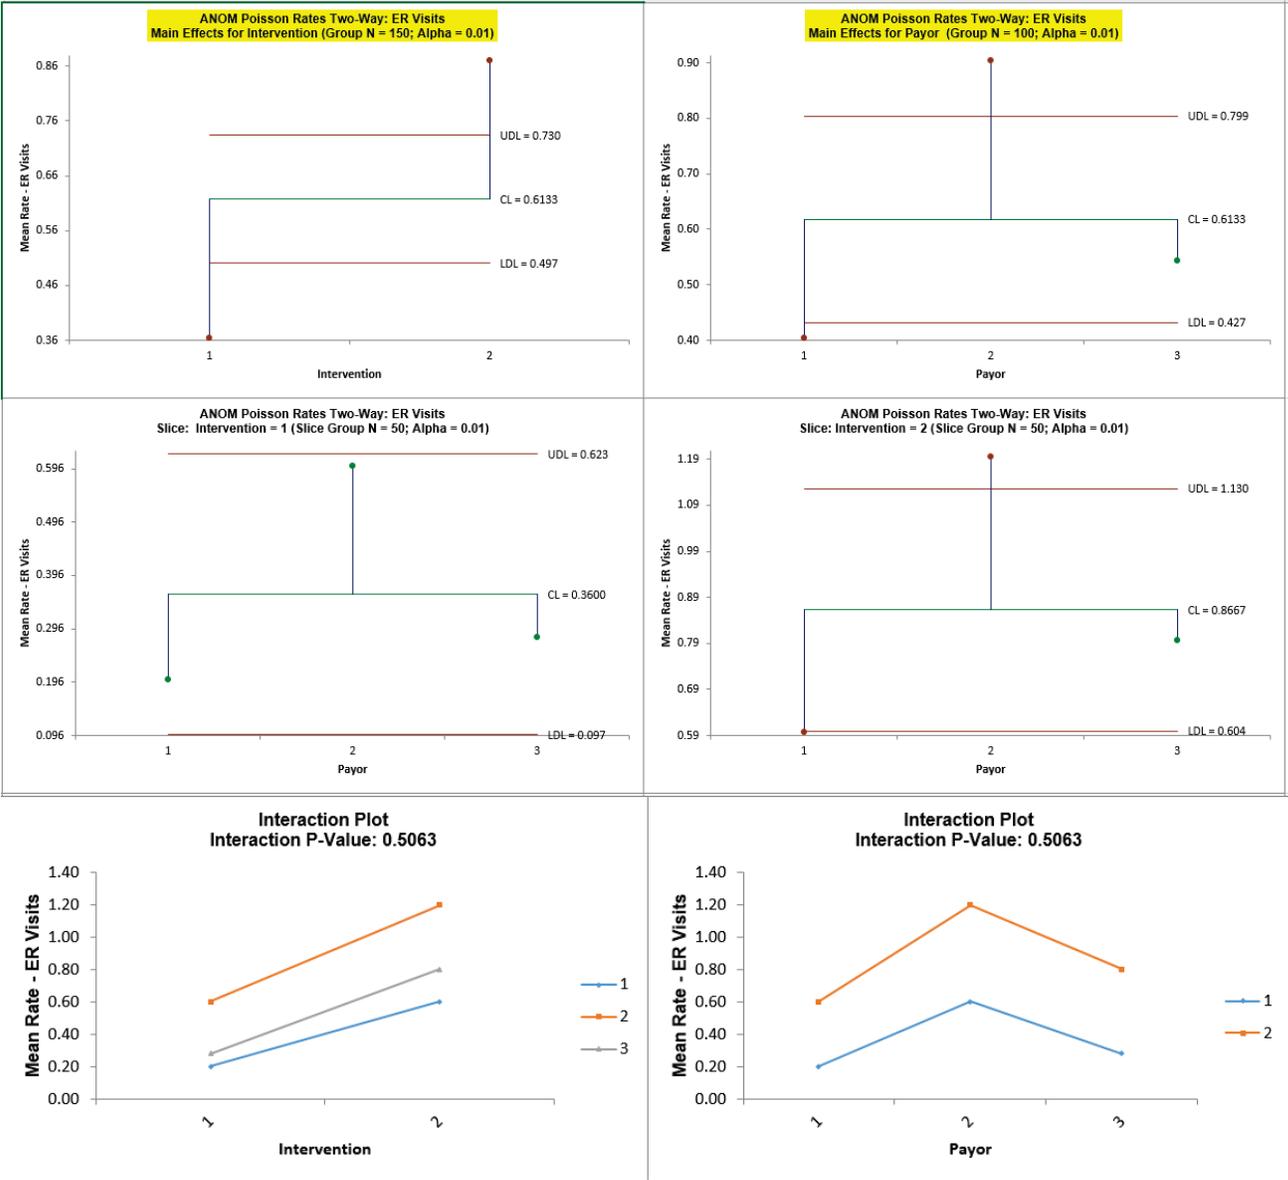
ER Visits
Sample Size
Intervention
Payor

OK >>
Cancel
Help

Alpha Level: 0.01

Adjust chart alpha for family-wise error rate

4. Click **OK**. The ANOM Poisson Rates Two-Way Main Effects, Slice charts and Interaction plots are shown below:

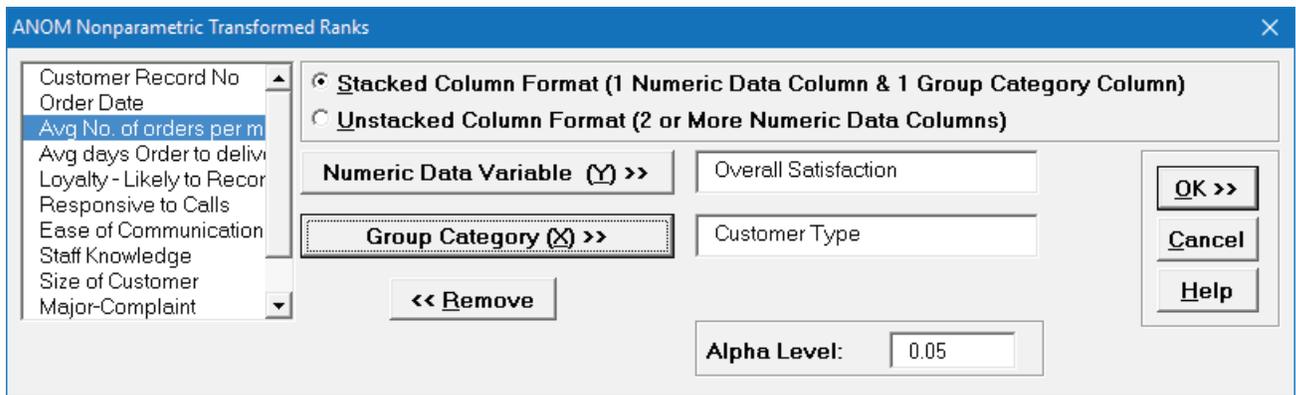


5. Since the Interaction is not significant (P-Value ≥ 0.1 , automatically determined from Poisson regression), the Main Effects charts are highlighted. Both factors Intervention and Payor are significant at alpha = 0.01. “Home nurse assigned” results in a significantly lower rate of ER visits versus “no home nurse assigned.” Payor Type 1 has a significantly lower rate of ER visits and Type 2 has a significantly higher rate.

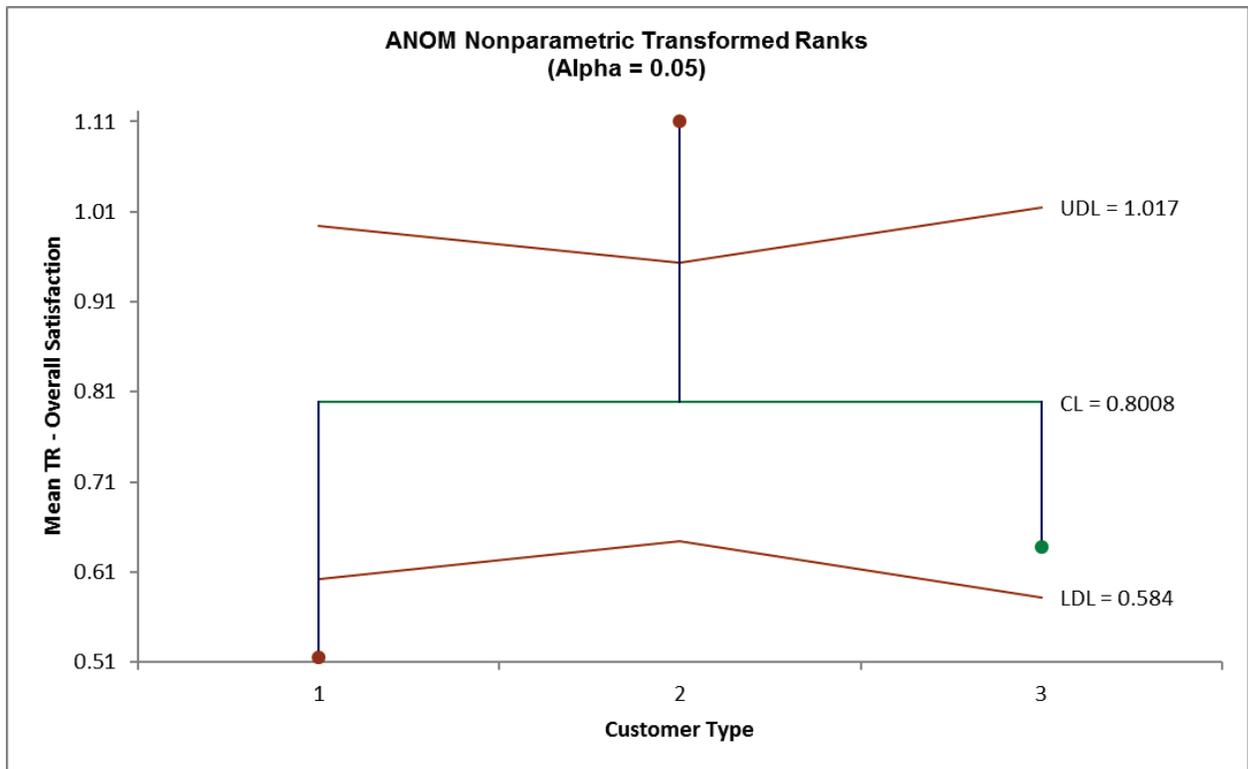
ANOM Nonparametric Transformed Ranks

ANOM Nonparametric Transformed Ranks is complementary to a Kruskal-Wallis test. Note that it is not included as an option in the Kruskal-Wallis dialog because there is a difference in the statistic used (Kruskal-Wallis uses ANOVA on ranks, whereas the ANOM utilizes a normal transformed rank).

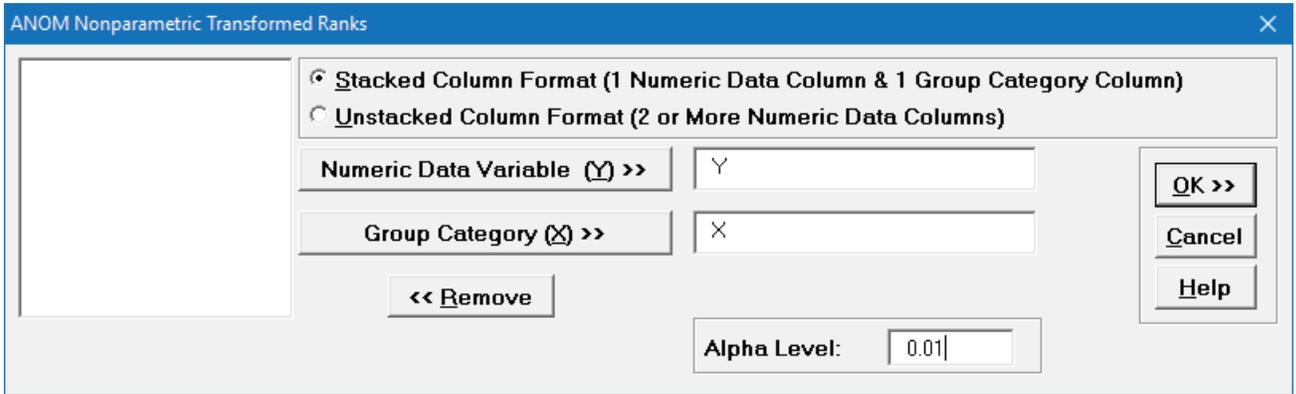
1. Open **Customer Data.xlsx**, click on **Sheet 1** tab.
2. Click **SigmaXL > Graphical Tools > Analysis of Means (ANOM) > ANOM Nonparametric Transformed Ranks**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**.
3. Click **Next**. Ensure that **Stacked Column Format** is checked. Select *Overall Satisfaction*, click **Numeric Data Variable (Y) >>**; select *Customer Type*, click **Group Category (X) >>**. Set **Alpha Level = 0.05**.



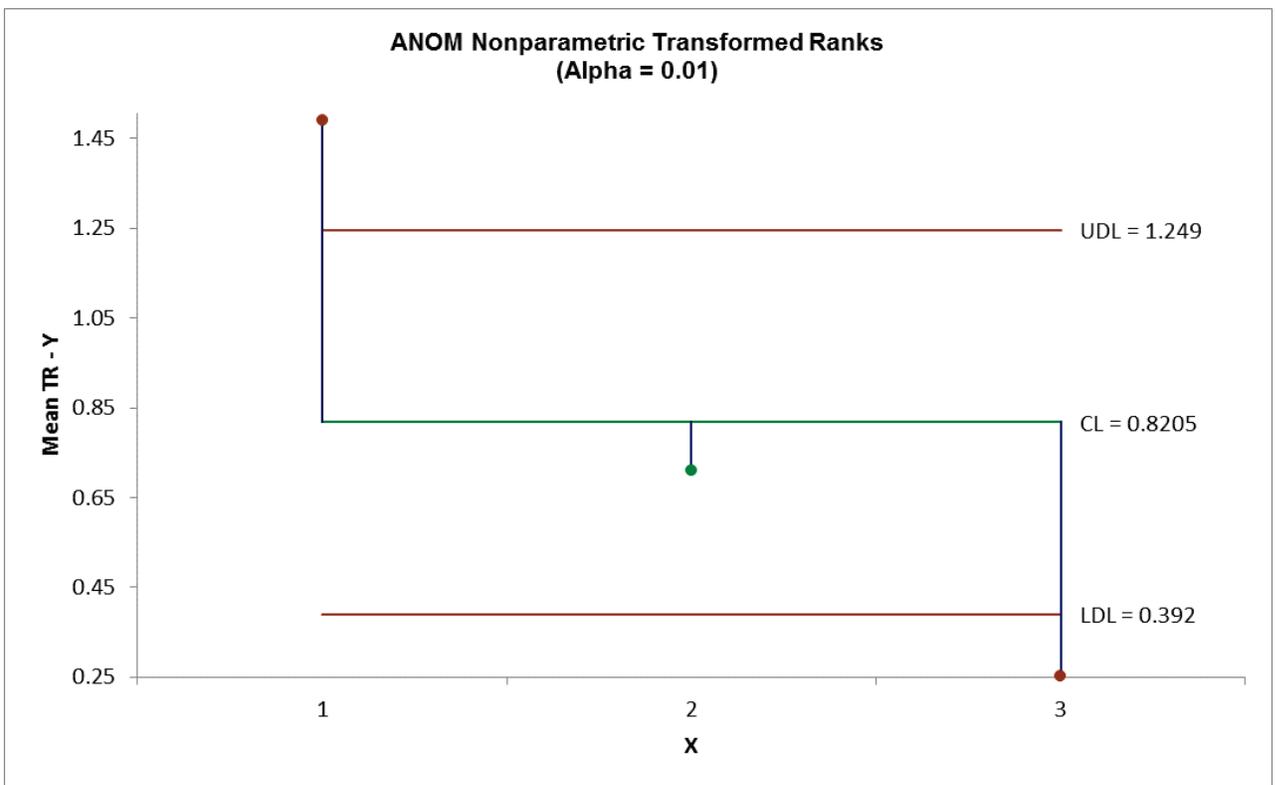
4. Click **OK**. The ANOM Nonparametric Transformed Ranks chart is shown below:



5. Here we see that Customer Type 1 mean transformed rank satisfaction score is significantly below the grand mean and Customer Type 2 is significantly higher. This is consistent with the results that were observed in the Kruskal-Wallis analysis (see [Kruskal-Wallis](#)).
6. The varying decision limits are due to the varying sample sizes for each Customer Type, with smaller sample size giving wider limits in a manner similar to a control chart. If the data are balanced, the decision limit lines will be constant.
7. Open **ANOM Examples.xlsx**, click on the **Exponential Data – TR** tab. This is exponential data, Example 9.16 from the ANOM book with alpha level = 0.01 (used with author permission).
8. Click **SigmaXL > Graphical Tools > Analysis of Means (ANOM) > ANOM Nonparametric Transformed Ranks**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**.
9. Click **Next**. Ensure that **Stacked Column Format** is checked. Select Y, click **Numeric Data Variable (Y) >>**; select X, click **Group Category (X) >>**. Set **Alpha Level = 0.01**.

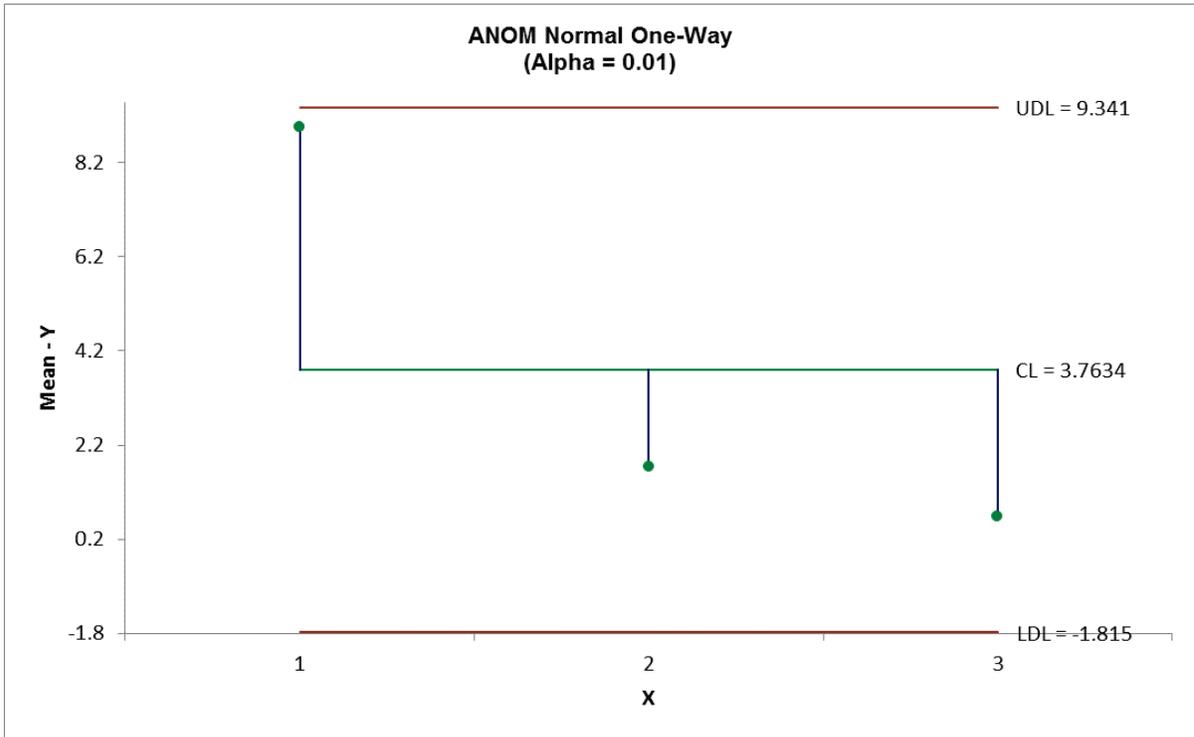


10. Click **OK**. The ANOM Nonparametric Transformed Ranks chart is shown below:



11. X = 1 shows as significantly higher than the grand mean transformed rank and X = 3 is significantly lower (at alpha = 0.01). These results match those given in the ANOM book.

12. By way of comparison, rerun the analysis using ANOM One-Way Normal (with $\alpha = 0.01$).
The results are:

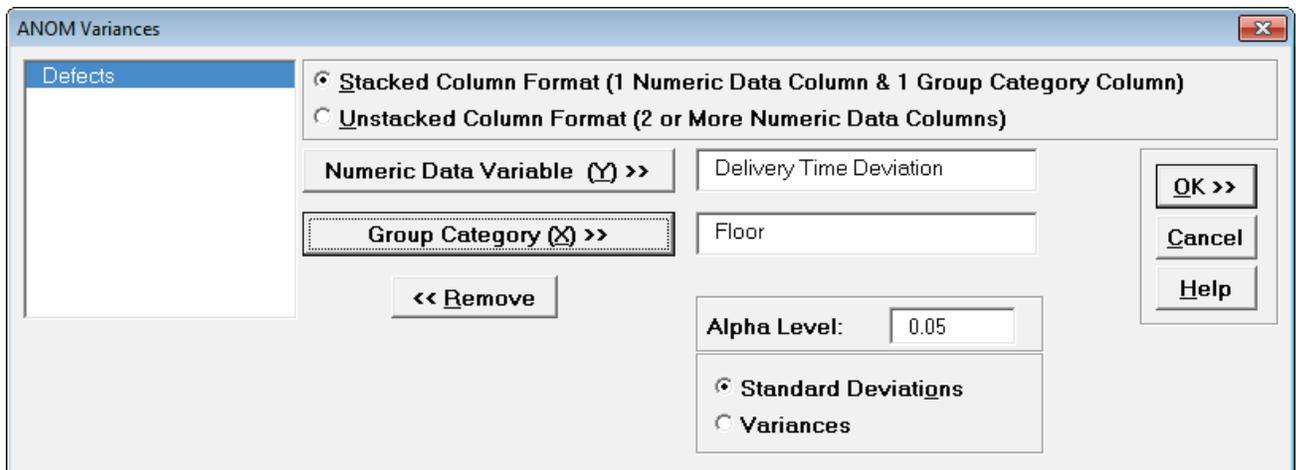


13. ANOM Normal fails to detect the significant differences in mean noted above.

ANOM Variances

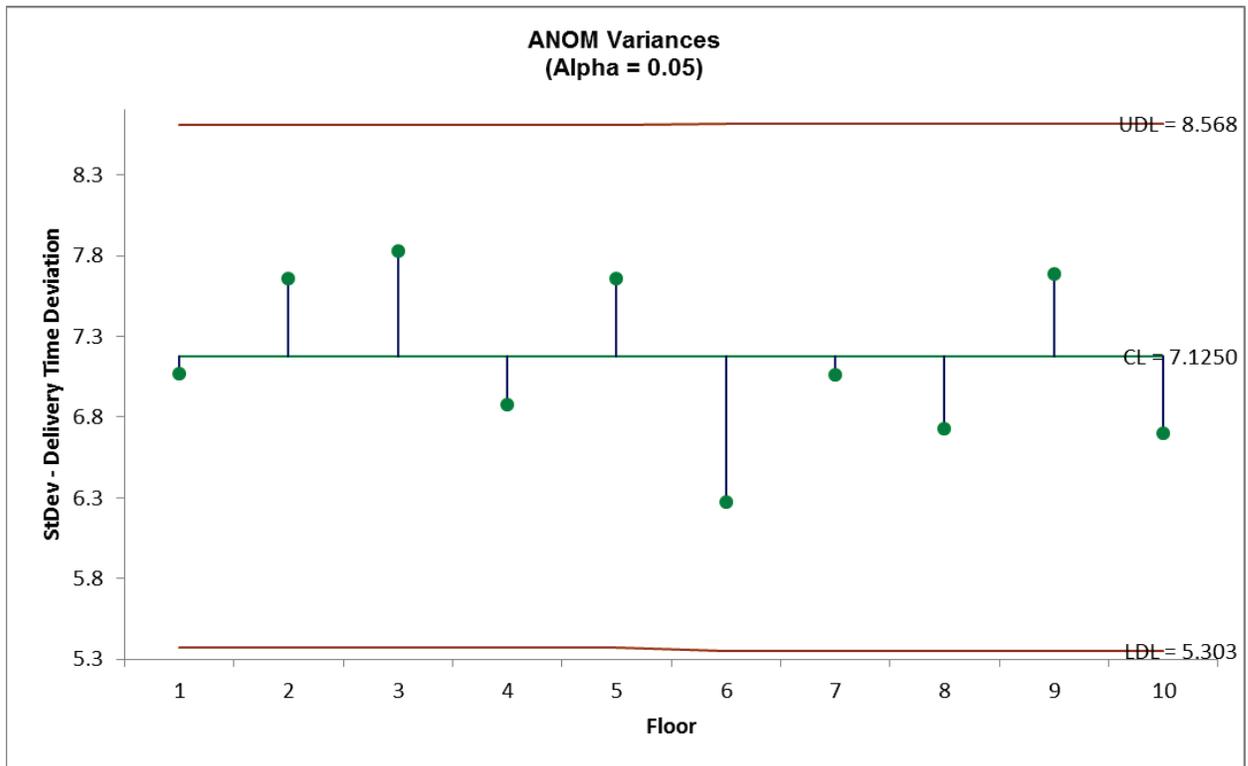
The ANOM Variances chart is complementary to Bartlett's Test for Equal Variance and is also available as an option in the Bartlett's Test dialog. It assumes that the data are normally distributed.

14. Open **Delivery Times.xlsx**, click on **Sheet 1** tab.
15. Click **SigmaXL > Graphical Tools > Analysis of Means (ANOM) > ANOM Variances**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**.
16. Click **Next**. Ensure that **Stacked Column Format** is checked. Select *Delivery Time Deviation*, click **Numeric Data Variable (Y) >>**; select *Floor*, click **Group Category (X) >>**; Select **Standard Deviations**; Set **Alpha Level** = 0.05.



We are analyzing the same normal data used in Bartlett's Test for Equal Variances. See [Bartlett's Test](#).

17. Click **OK**. The ANOM Variances chart is shown below:

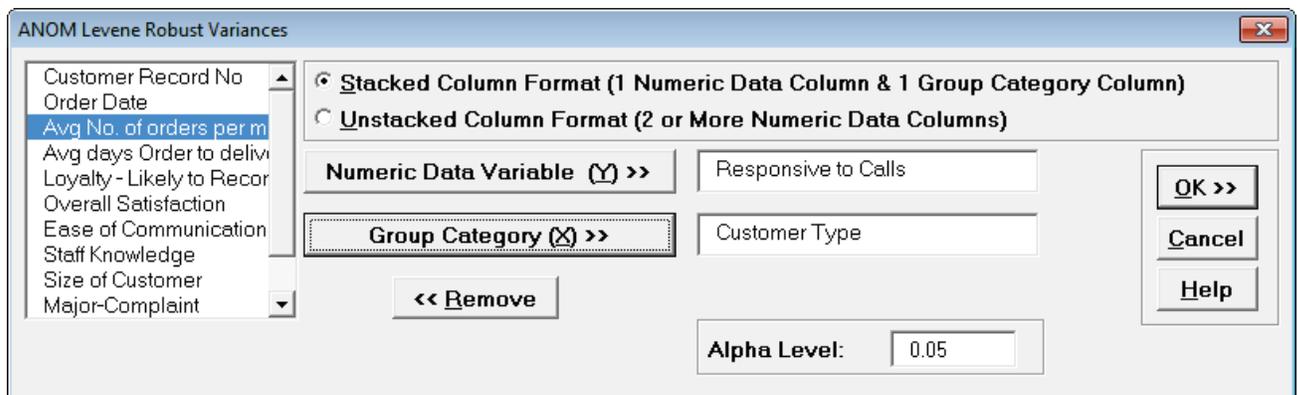


18. The ANOM Variances chart visually shows that none of the group standard deviations are significantly different from the grand mean of all the standard deviations. It is called an ANOM Variances Chart but displays Standard Deviations for ease of interpretation (similar to a Standard Deviation S Control Chart). This does however result in non-symmetrical decision limits. To display Variances rerun the above analysis with option **Variances** selected.

ANOM Levene Robust Variances

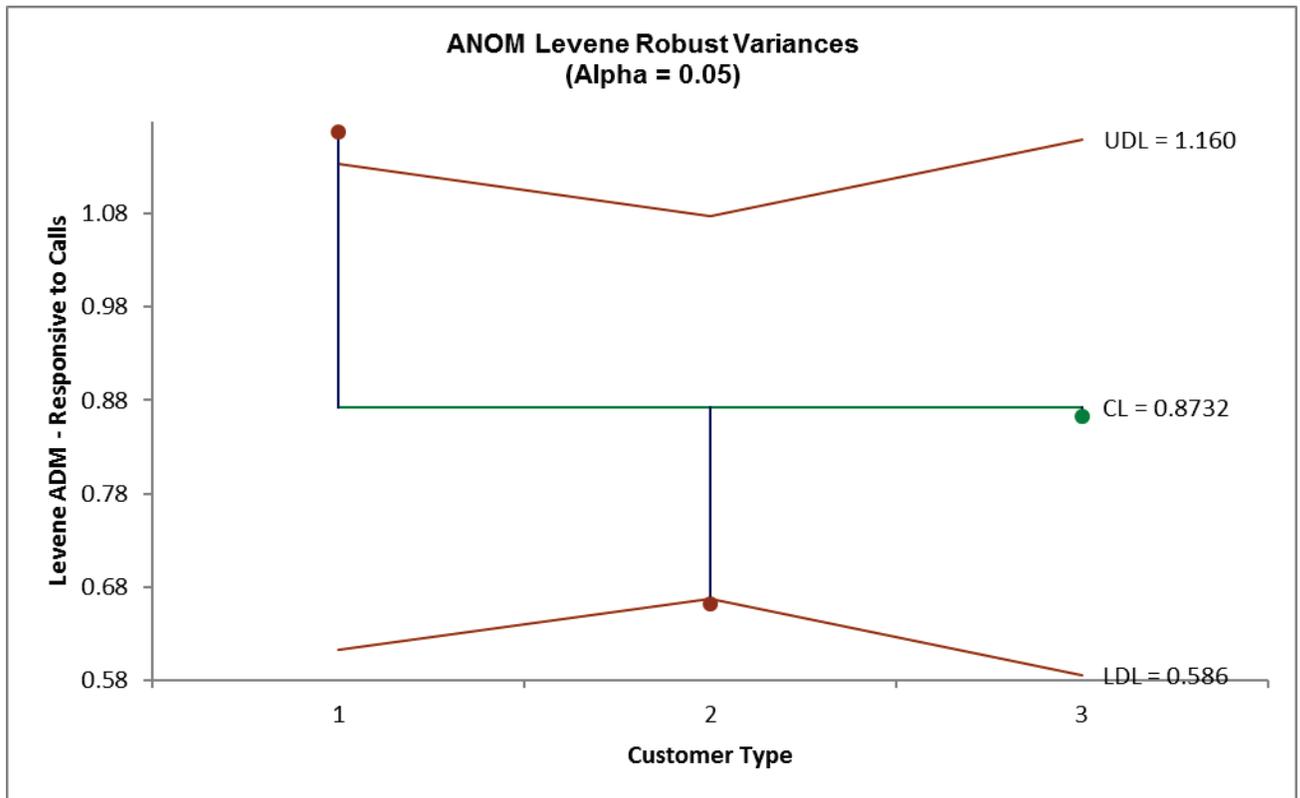
The ANOM Levene Robust Variances chart is complementary to Levene's Test for Equal Variance and is also available as an option in the Levene's Test dialog. Since it uses absolute deviations from the median, it is robust to the assumption of normality.

1. Open **Customer Data.xlsx**, click on **Sheet 1** tab.
2. Click **SigmaXL > Graphical Tools > Analysis of Means (ANOM) > ANOM Levene Robust Variances**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**.
3. Click **Next**. Ensure that **Stacked Column Format** is checked. Select *Responsive to Calls*, click **Numeric Data Variable (Y) >>**; select *Customer Type*, click **Group Category (X) >>**; Set **Alpha Level** = 0.05.



We are analyzing the same nonnormal data used in Levene's Test for Equal Variances. See [Levene's Test](#).

- Click OK. The ANOM Levene Robust Variances chart is displayed:



- The ANOM chart clearly shows Customer Type 1 has significantly higher variance (ADM) than overall and Customer Type 2 has significantly lower variance.
- The varying decision limits are due to the varying sample sizes for each Customer Type, with smaller sample size giving wider limits in a manner similar to a control chart. If the data are balanced the decision limit lines will be constant.

Part Q – Multi-Vari Charts

Multi-Vari Charts

The Multi-Vari chart is a powerful tool to identify dominant Sources of Variation (SOV). The three major “families” of variation are: Within Unit, Between Unit, and Temporal (Over Time). We will look at examples of each type of SOV and then use the Multi-Vari Chart to study Overall Satisfaction in the **Customer Data.xlsx** file.

1. Open **Multi-Vari Data.xlsx**, click Sheet **Within**. Select **SigmaXL > Graphical Tools > Multi-Vari Options**.
2. The charts shown will be updated as options are selected. Note that they are for demonstration purposes and are not based on the **Multi-Vari Data.xlsx** data (sample charts are not displayed in Excel for Mac). Ensure that all general **Options** are selected (**Range Line, Individual Data Points, Min and Max, Standard Deviation Chart**). Select **Mean Options**, ensure that **Show Means, Connect Means, and Group Means** are checked. Ensure that **Save Defaults** is checked. These settings would be typical for a Multi-Vari chart. (The Median options provide the ability to display percentiles as an alternative to the Means).

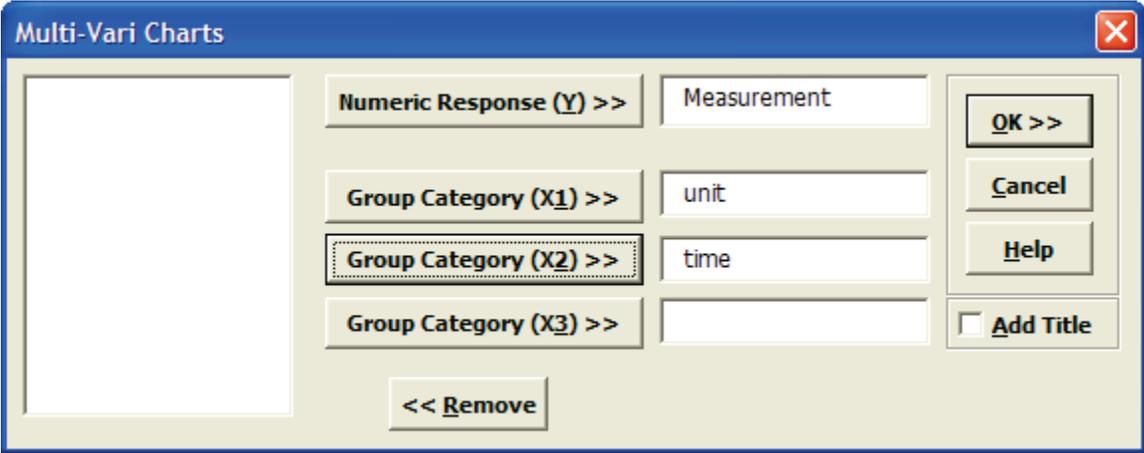
General Options	Mean Options	Median Options
<input checked="" type="checkbox"/> Range Line	<input checked="" type="checkbox"/> Show Means	<input type="checkbox"/> Show Medians
<input checked="" type="checkbox"/> Individual Data Points	<input type="checkbox"/> 95% Confidence Interval	<input type="checkbox"/> 25th, 75th Percentile
<input checked="" type="checkbox"/> Min and Max	<input checked="" type="checkbox"/> Connect Means	<input type="checkbox"/> Connect Medians
<input checked="" type="checkbox"/> Standard Deviation Chart	<input checked="" type="checkbox"/> Group Means	<input type="checkbox"/> Group Medians
	<input type="checkbox"/> Grand Mean	<input type="checkbox"/> Grand Median

Buttons: , , , Save Defaults

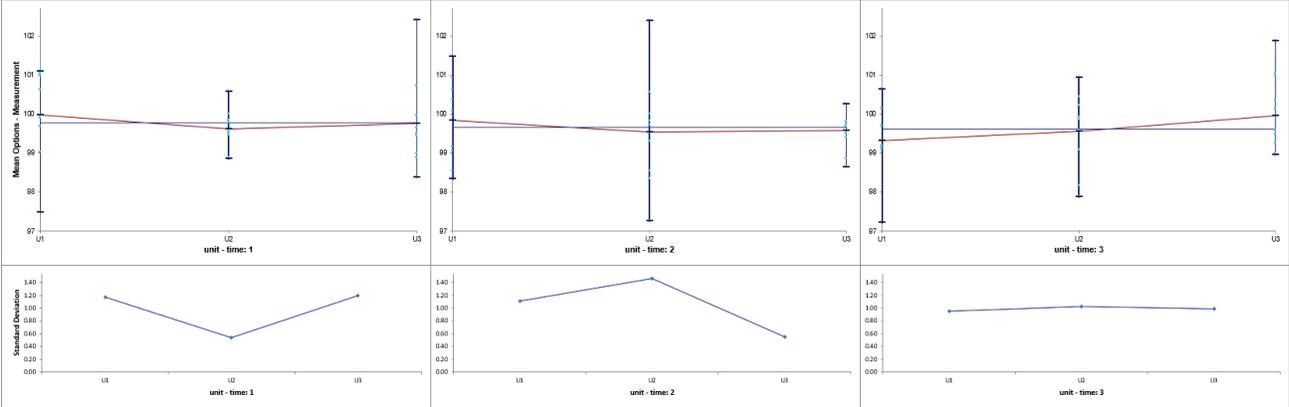
Tip: Multi-Vari Charts can be used to display Confidence Intervals (as we did earlier in Part C). To do this, check the 95% Confidence Interval.

3. Click **Finish**. SigmaXL automatically starts the Multi-Vari Chart procedure (this is equivalent to clicking **SigmaXL > Graphical Tools > Multi-Vari Charts**).
4. Check **Use Entire Data Table**. Click **Next**.
5. Note that the input X’s can be text or numeric but should be discrete. Y’s must be numeric - typically continuous, but can also be count or proportion data.

- 6. Select *Measurement*, click **Numeric Response (Y) >>**; select *unit*, click **Group Category (X1) >>**; select *time*, click **Group Category (X2) >>**.

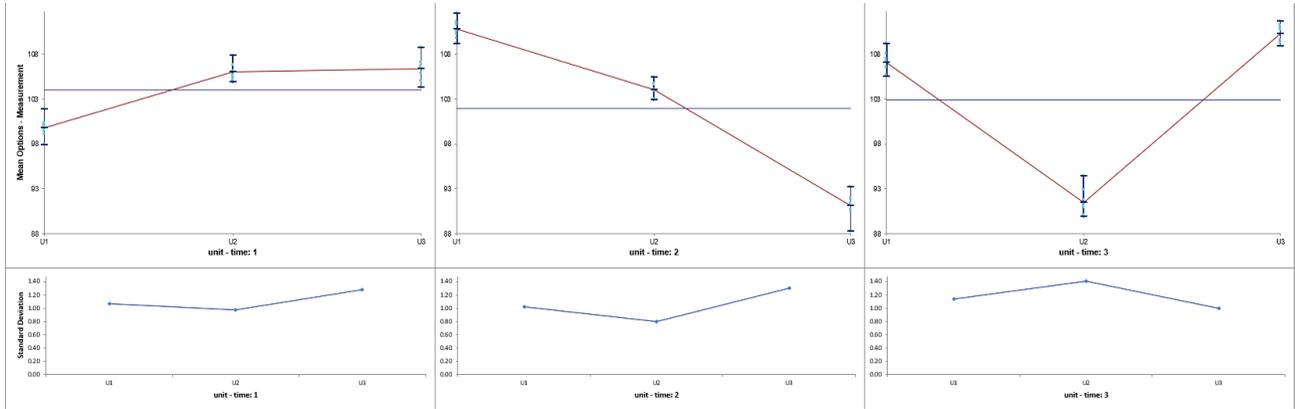


- 7. Click **OK**. Resulting Multi-Vari Chart illustrating dominant “Within Unit” Source of Variation:



- 8. Using **Multi-Vari Data.xlsx**, click Sheet **Between**. Select **SigmaXL > Graphical Tools > Multi-Vari Charts**. Check **Use Entire Data Table**. Click **Next**.
- 9. Select *Measurement*, click **Numeric Response (Y) >>**; select *unit*, click **Group Category (X1) >>**; select *time*, click **Group Category (X2) >>**.
- 10. Click **OK**. Resulting Multi-Vari Chart illustrating dominant “Between Unit” Source of Variation:

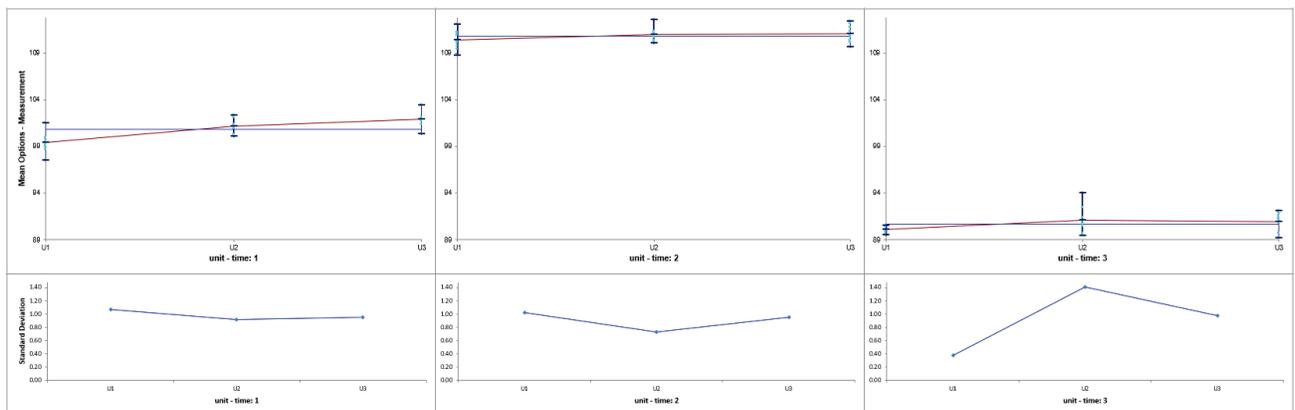
SigmaXL: Analyze Phase Tools



11. Using **Multi-Vari Data.xlsx**, click Sheet **OverTime**. Select **SigmaXL > Graphical Tools > Multi-Vari Charts**. Check **Use Entire Data Table**. Click **Next**.

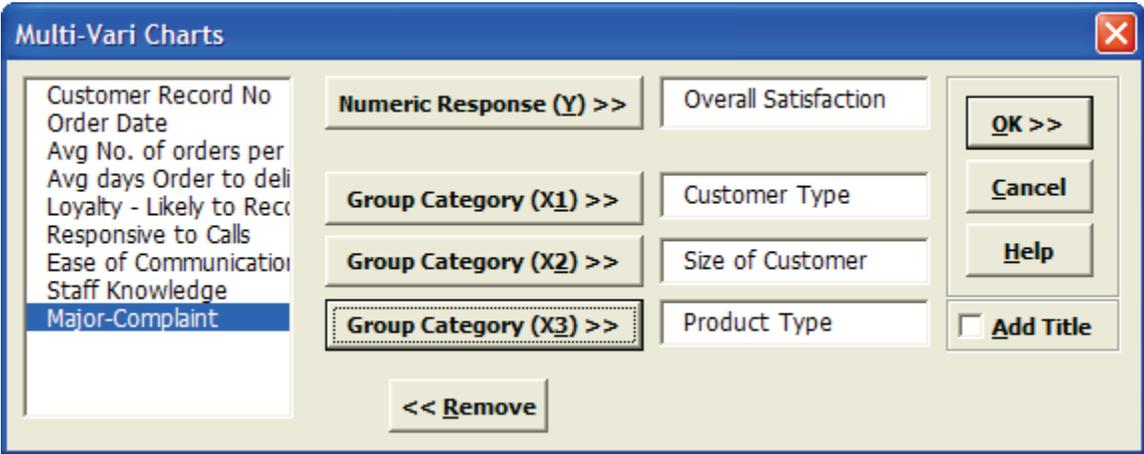
12. Select *Measurement*, click **Numeric Response (Y) >>**; select *unit*, click **Group Category (X1) >>**; select *time*, click **Group Category (X2) >>**.

13. Click **OK**. Resulting Multi-Vari Chart illustrating dominant “Over Time” Source of Variation:

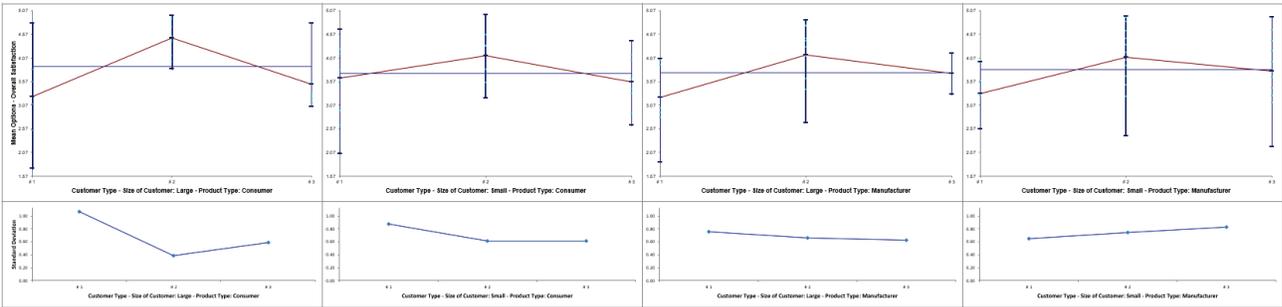


14. Open **Customer Data.xlsx**, click **Sheet 1**. Select **SigmaXL > Graphical Tools > Multi-Vari Charts**.

- 15. Select *Overall Satisfaction*, click **Numeric Response (Y) >>**; select *Customer Type*, click **Group Category (X1) >>**; select *Size of Customer*, click **Group Category (X2) >>**; select *Product Type*, click **Group Category (X3) >>**.



- 16. Click **OK**. Resulting Multi-Vari chart:



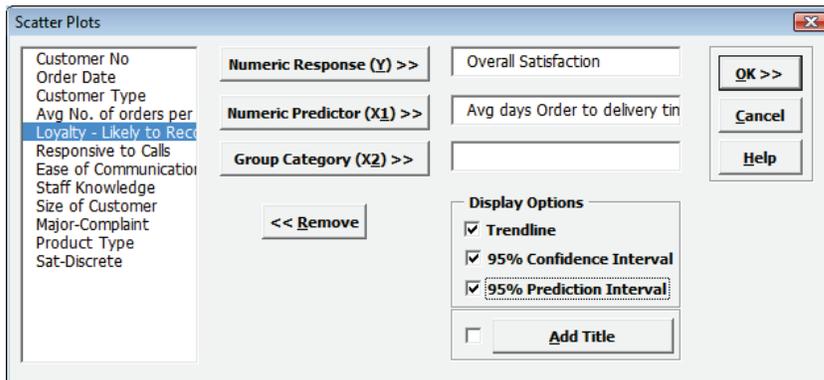
Examining this Multi-Vari chart reveals that the dominant Source of Variation is “within” Customer Type, followed by “between” Customer Type. Furthermore, it would be worthwhile to examine the combination of Customer Type 2, Customer Size Large, and Product Type Consumer.

Other tools that can help us identify potential X factors that may explain some of the large “Within” variability are the Scatter Plot, Scatter Plot Matrix and Correlation Matrix.

Part R – Scatter Plots

Scatter Plots

1. Open **Customer Data.xlsx**. Click **Sheet 1** Tab. Click **SigmaXL > Graphical Tools > Scatter Plots**; if necessary, click **Use Entire Data Table**, click **Next**.
2. Select *Overall Satisfaction*, click **Numeric Response (Y) >>**; select *Avg Days Order Time to Delivery*, click **Numeric Predictor (X1) >>**. Check **Trendline**, **95% Confidence Interval** and **95% Prediction Interval** as shown:

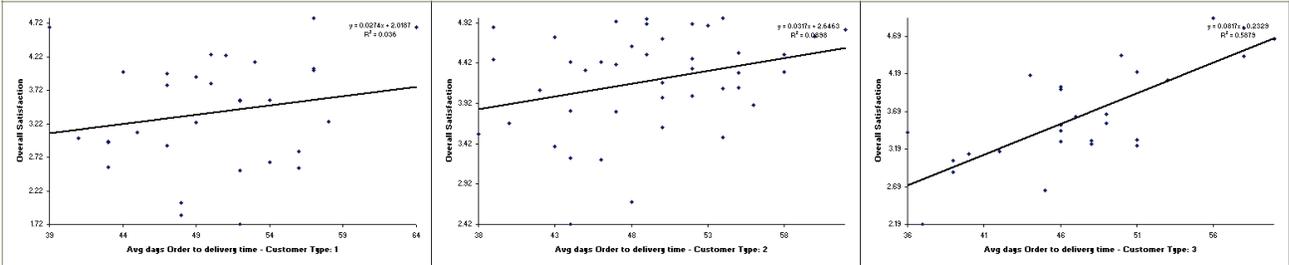


3. Click **OK**. The resulting Scatter Plot is shown with equation, trendline, 95% confidence interval (blue lines – for a given X value this is the 95% confidence interval for predicted mean Overall Satisfaction) and 95% prediction interval (red lines – for a given X value this is the 95% confidence interval for predicted individual values of Overall Satisfaction).



The equation is based on linear regression, using the method of least squares. R-squared * 100 is the percent variation of Y explained by X (here 10.3%).

- 4. Now we want to redo the Scatter Plot and stratify by Customer Type. Press **F3** or click **Recall SigmaXL Dialog** to recall last dialog. (Or, Click **Sheet 1 Tab**; Click **SigmaXL > Graphical Tools > Scatter Plots**; click **Next**.)
- 5. Select *Overall Satisfaction*, click **Numeric Response (Y) >>**; select *Average Days Order Time to Delivery*, click **Numeric Predictor (X1) >>**; select *Customer Type*, click **Group Category (X2) >>**, Uncheck **95% Confidence Interval** and **95% Prediction Interval**. Click **OK**:

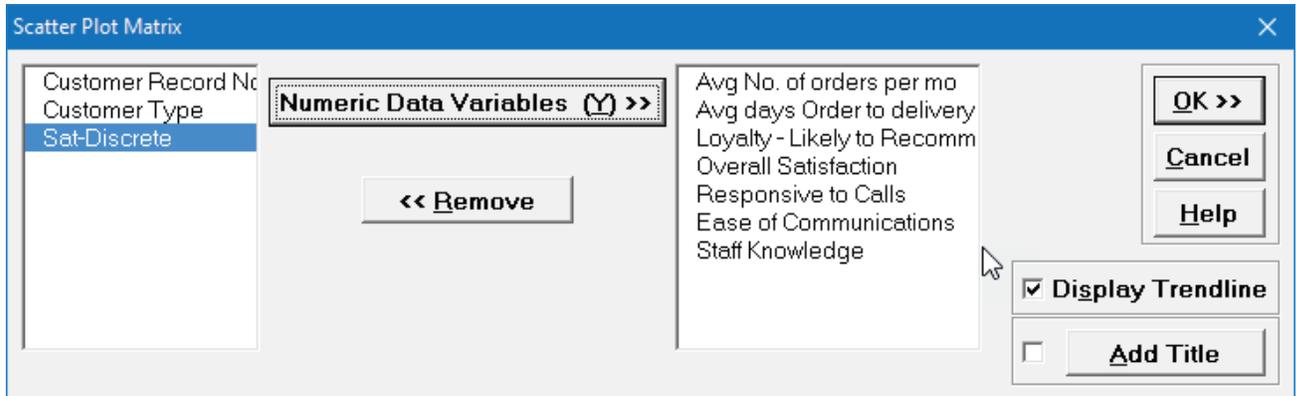


Clearly, according to the analysis, Customer Type 3 is happier if orders take longer! But, does this make sense? Of course not! Customer Sat scores should not increase with Order to Delivery time. What is happening here? This is a coincidental situation. Something else is driving customer satisfaction. We will now look at the Scatter Plot Matrix to help us investigate other factors influencing Customer Satisfaction.

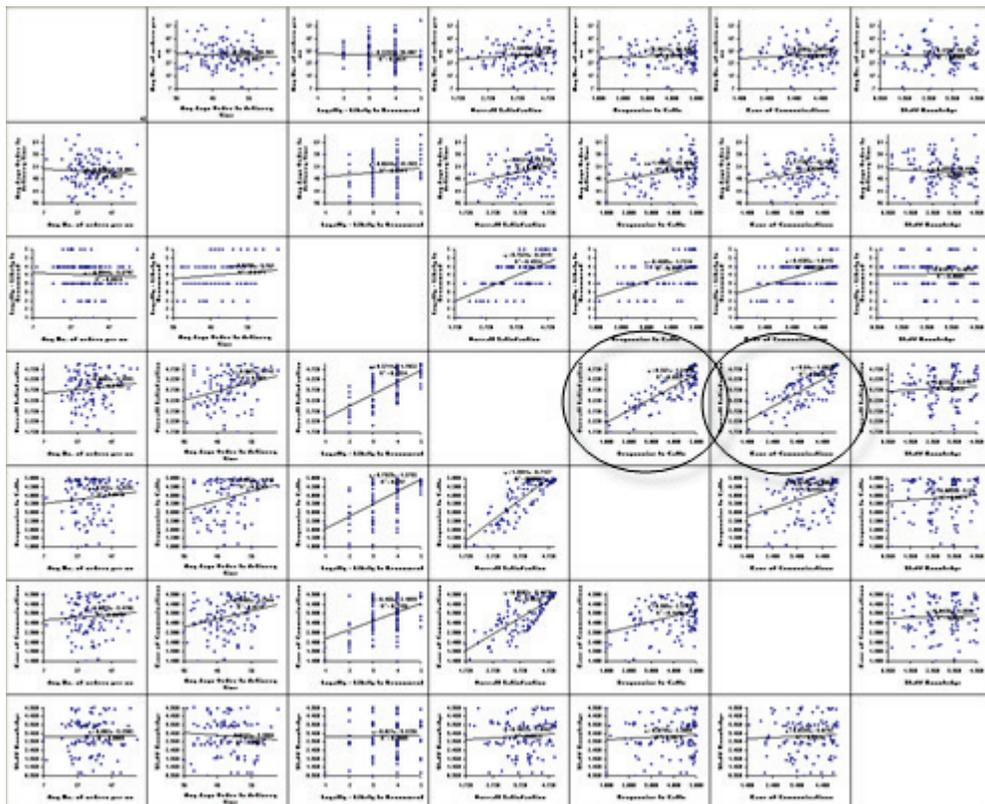
Tip: Be careful when interpreting scatter plots (and other statistical tools): Y versus X correlation or statistical significance does not always mean that we have a causal relationship. Umbrella sales are highly correlated to traffic accidents, but we cannot reduce the rate of traffic accidents by purchasing fewer umbrellas! The best way to validate a relationship is to perform a Design of Experiments (see Improve Phase).

Scatter Plot Matrix

1. Click **Sheet 1** Tab of **Customer Data.xlsx** (or press **F4** to activate last worksheet). Click **SigmaXL** > **Graphical Tools** > **Scatter Plot Matrix**.
2. Ensure that entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
3. Select the variable *Avg No. of orders per month*; shift-click on *Staff Knowledge* and click **Numeric Data Variables (Y) >>** as shown:



4. Click **OK**. Resulting Scatter Plot Matrix:



Of particular interest is Overall Satisfaction versus Responsiveness to Calls and Ease of Communications. These will be explored further with Multiple Linear Regression.

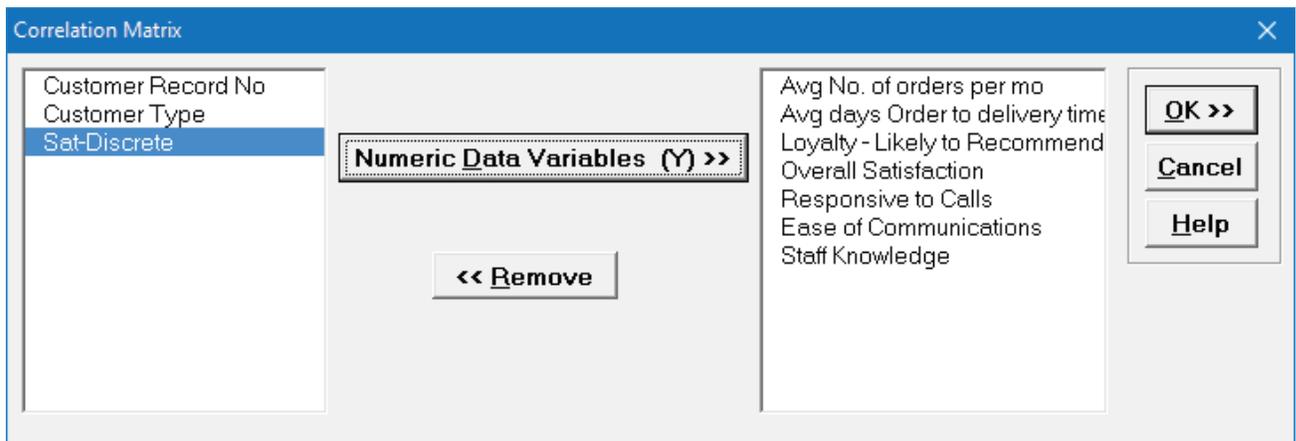
Part S – Correlation Matrix

The correlation matrix complements the scatterplot matrix by quantifying the degree of association. The following table shows the approximate relationship between r , R-squared, and degree of association:

Pearson Correlation Coefficient (r)	R-Squared (%)	Degree of association
$0.9 \leq r \leq 1$	> 80 %	Strong
$0.7 \leq r < 0.9$	50 % to 80 %	Moderate
$ r < 0.7$	< 50 %	Weak
Pearson Probability, $p > 0.05$		None

Correlation Matrix

1. Open **Customer Data.xlsx**. Click **Sheet 1** tab (or press **F4** to activate last worksheet). Click **SigmaXL > Statistical Tools > Correlation Matrix**. Ensure that entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
2. Select the variable *Avg No. of orders per month*; shift-click on *Staff Knowledge* and click **Numeric Data Variables (Y) >>** as shown:



3. Click **OK**. Resulting Correlation Matrix is shown:

Pearson Correlations	Avg No. of orders per mo	Avg days Order to delivery time	Loyalty - Likely to Recommend	Overall Satisfaction	Responsive to Calls	Ease of Communications	Staff Knowledge
Avg No. of orders per mo	1	-0.0518	-0.0491	0.1155	0.1076	0.0885	0.0186
Avg days Order to delivery time		1	0.1307	0.3218	0.2725	0.2681	-0.0781
Loyalty - Likely to Recommend			1	0.6598	0.5805	0.4622	0.0176
Overall Satisfaction				1	0.8262	0.7454	0.0766
Responsive to Calls					1	0.3791	0.0845
Ease of Communications						1	0.0506
Staff Knowledge							1

Pearson Probabilities	Avg No. of orders per mo	Avg days Order to delivery time	Loyalty - Likely to Recommend	Overall Satisfaction	Responsive to Calls	Ease of Communications	Staff Knowledge
Avg No. of orders per mo	1	0.6090	0.6279	0.2527	0.2865	0.3812	0.8541
Avg days Order to delivery time		1	0.1949	0.0011	0.0051	0.0070	0.4398
Loyalty - Likely to Recommend			1	0.0000	0.0000	0.0000	0.8622
Overall Satisfaction				1	0.0000	0.0000	0.4490
Responsive to Calls					1	0.0001	0.4035
Ease of Communications						1	0.6171
Staff Knowledge							1

Spearman Rank Correlations	Avg No. of orders per mo	Avg days Order to delivery time	Loyalty - Likely to Recommend	Overall Satisfaction	Responsive to Calls	Ease of Communications	Staff Knowledge
Avg No. of orders per mo	1	-0.0305	-0.0917	0.1006	0.0738	0.1000	0.0187
Avg days Order to delivery time		1	0.1097	0.3407	0.2489	0.2513	-0.0828
Loyalty - Likely to Recommend			1	0.6167	0.5503	0.4071	-0.0190
Overall Satisfaction				1	0.7782	0.7089	0.0890
Responsive to Calls					1	0.3204	0.0895
Ease of Communications						1	0.0716
Staff Knowledge							1

Spearman Rank Probabilities	Avg No. of orders per mo	Avg days Order to delivery time	Loyalty - Likely to Recommend	Overall Satisfaction	Responsive to Calls	Ease of Communications	Staff Knowledge
Avg No. of orders per mo	1	0.7629	0.3643	0.3192	0.4655	0.3222	0.8532
Avg days Order to delivery time		1	0.2774	0.0005	0.0125	0.0087	0.4127
Loyalty - Likely to Recommend			1	0.0000	0.0000	0.0000	0.8514
Overall Satisfaction				1	0.0000	0.0000	0.3786
Responsive to Calls					1	0.0012	0.3758
Ease of Communications						1	0.4792
Staff Knowledge							1

Correlations highlighted in red are considered significant (P-Values < .05). The corresponding correlation coefficients above the P-Values are also highlighted in red. (Compare these to the Scatterplot Matrix.)

Note that Spearman’s Rank Correlation complements Pearson’s Correlation, in that it provides a robust measure of association. Spearman’s rank is based on correlated ranks, which are not sensitive to outliers or departures from normality.

An automatic normality check is applied, utilizing the powerful Doornik-Hansen bivariate normality test. A yellow highlight recommends Pearson or Spearman correlations be used (but only if it is significant). Pearson is highlighted if the data are bivariate normal, otherwise Spearman is highlighted. Always review the data graphically with the scatterplot matrix as well.

Reference

Doornik, J.A. and Hansen, H. “An Omnibus Test for Univariate and Multivariate Normality,” *Oxford Bulletin Of Economics And Statistics*, 70, Supplement (2008).

Part T – Multiple Regression

Multiple Regression

Multiple Regression analyzes the relationship between one dependent variable (Y) and multiple independent variables (X's). It is used to discover the relationship between the variables and create an empirical equation of the form:

$$Y = b_0 + b_1 * X_1 + b_2 * X_2 + \dots + b_n * X_n$$

This equation can be used to predict a Y value for a given set of input X values. SigmaXL uses the method of least squares to solve for the model coefficients and constant term. Statistical tests of hypothesis are provided for the model coefficients.

1. Open **Customer Data.xlsx**. Click **Sheet 1** Tab (or press **F4** to activate last worksheet). Click **SigmaXL > Statistical Tools > Regression > Multiple Regression**. If necessary, click **Use Entire Data Table**, click **Next**.
2. Select *Overall Satisfaction*, click **Numeric Response (Y) >>**, select *Responsive to Calls* and *Ease of Communications*, click **Continuous Predictors (X) >>**. **Fit Intercept** and **Display Residual Plots** are checked by default.

Note: Fit Intercept is optional. If unchecked the model will not fit an intercept (constant) for the model. This is useful when you have strong *a priori* reasons to believe that $Y = 0$ when the X or X's are equal to 0 and the relationship is linear. An example would be $Y = \text{automobile fuel consumption}$, $X = \text{automobile weight}$.

3. Click **OK**. The resulting Multiple Regression report is shown:

Multiple Regression Model: Overall Satisfaction = (0.493463) + (0.435673) * Responsive to Calls + (0.433346) * Ease of Communications

Model Summary:

R-Square	90.08%
R-Square Adjusted	89.88%
S (Root Mean Square Error)	0.248942

Parameter Estimates:

Predictor Term	Coefficient	SE Coefficient	T	P	VIF	Tolerance
Constant	0.493463	0.116857	4.223	0.0001		
Responsive to Calls	0.435673	0.023710993	18.374	0.0000	1.168	0.856307
Ease of Communications	0.433346	0.029667131	14.607	0.0000	1.168	0.856307

Analysis of Variance for Model:

Source	DF	SS	MS	F	P
Model	2	54.601	27.301	440.53	0.0000
Error	97	6.011	0.061972154		
Total (Model + Error)	99	60.612	0.612246		

Durbin-Watson Test for Autocorrelation in Residuals:

DW Statistic	1.815
P-Value Positive Autocorrelation	0.1744
P-Value Negative Autocorrelation	0.8209

This model of Overall Satisfaction as a function of Responsiveness to Calls and Ease of Communications looks very good with an R-Square value of 90%. Both Predictors are shown to be significant with their respective P-Values < .05. Clearly, we need to focus on these two X factors to improve customer satisfaction.

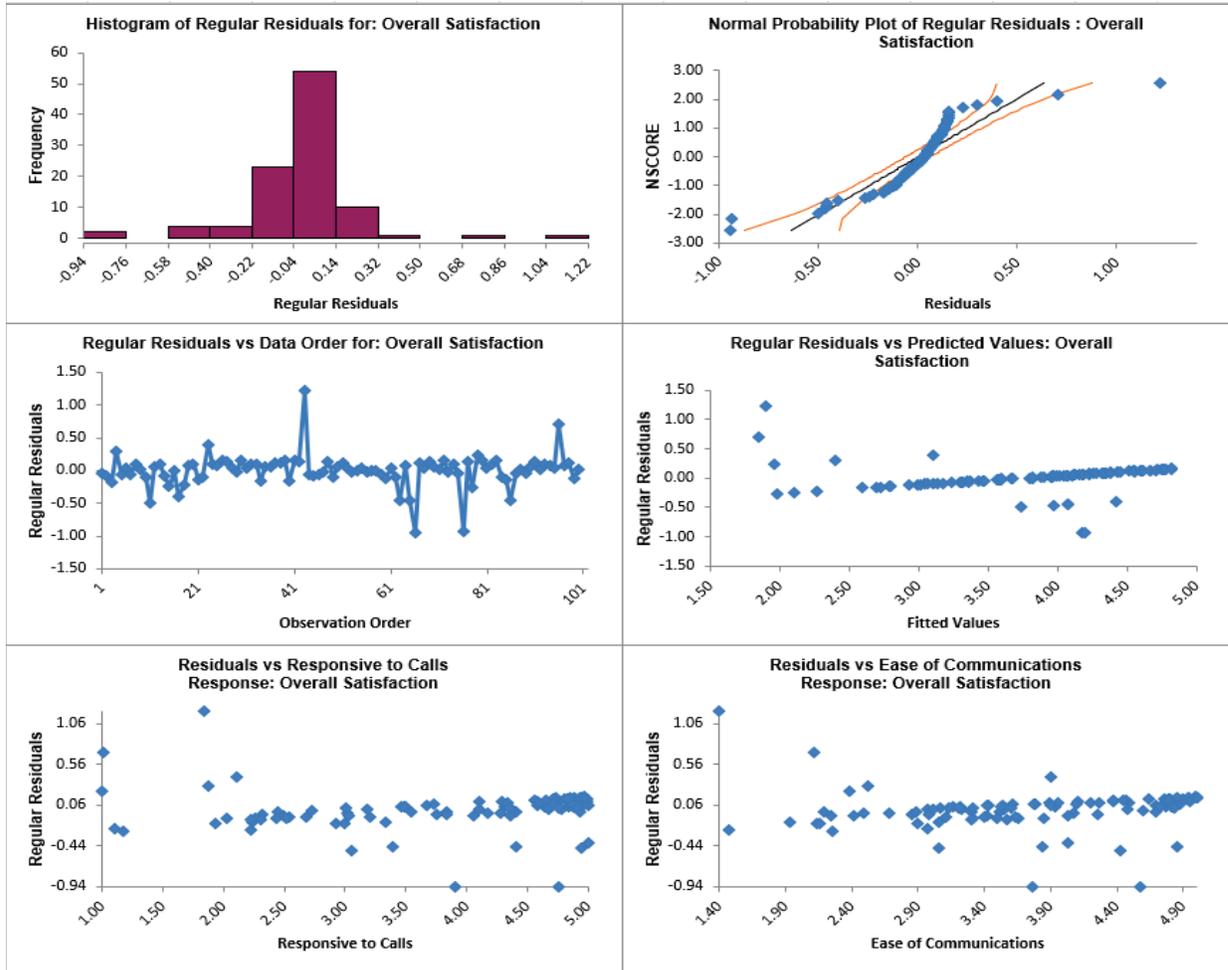
- The variance inflation factor (VIF) and Tolerance scores are used to measure multicollinearity or correlation among predictors. VIF = 1 indicates no relation among predictors (which is highly desirable – you will see examples of this in Design of Experiments); VIF > 1 indicates that the predictors are correlated; VIF > 5 indicates that the regression coefficients are strongly correlated and this will lead to poor estimates of the coefficients. Tolerance = 1/VIF.
- The Durbin-Watson Test is used to determine if the residuals are autocorrelated. If either P-Value is < .05, then there is significant autocorrelation. This is likely due to an external bias factor affecting your process over time (e.g. “warm-up” effect, seasonality). This will also be evident in the plot of residuals versus observation order. Autocorrelation in the residuals is not a problem for this model.
- If the Fit Intercept option is unchecked, the following diagnostics are not reported due to statistical issues associated with the removal of the constant. These are: R-Square, R-Square Adjusted, VIF (Variance Inflation Factor) and Tolerance collinearity diagnostics, DW (Durbin-Watson) autocorrelation report, Anova report for Discrete factors (discrete factors will not be permitted when "Fit Intercept" is unchecked). All other tests including residual diagnostics are reported. Users can run the model with the "Fit Intercept" on to get the above statistics and then run with "Fit Intercept" off.

- A predicted response calculator allows you to enter X values and obtain the predicted response value, including the 95% confidence interval for the long term mean and 95% prediction interval for individual values:

Predicted Response Calculator:

Predictors	Enter Settings:	Predicted Response	Lower 95% CI	Upper 95% CI	Lower 95% PI	Upper 95% PI
Responsive to Calls	5	4.838555471	4.750554344	4.926556597	4.336698591	5.34041235
Ease of Communications	5					

- Click the **Mult Reg Residuals** Sheet and scroll right to view the Residual Plots shown below:

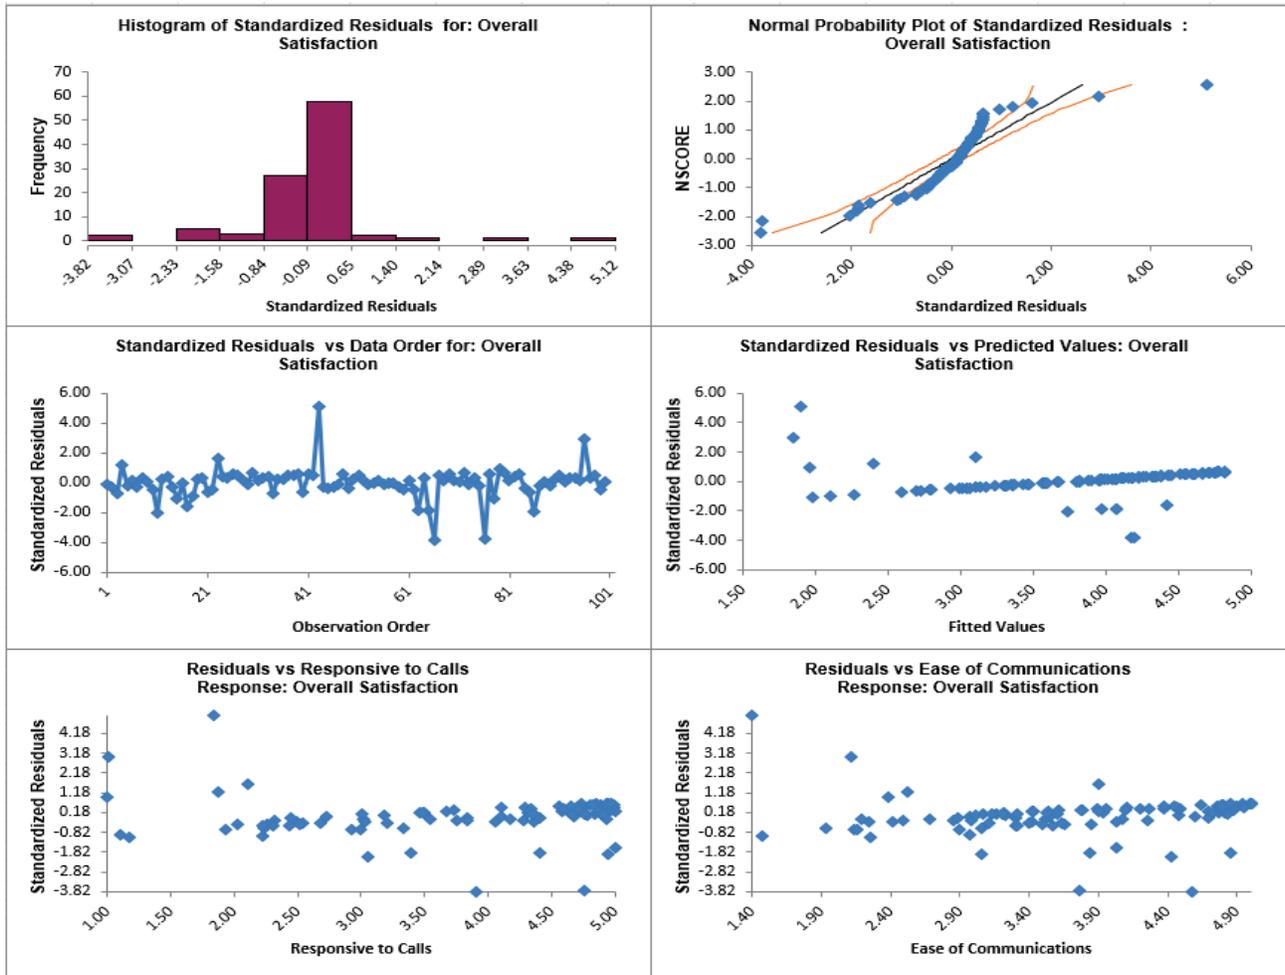


Residuals are the unexplained variation from the regression model ($Y - \hat{Y}$). We expect to see the residuals normally distributed with no obvious patterns in the above graphs. Clearly this is not the case here, with the Residuals versus Predicted Values indicating there is likely some other X factor influencing the Overall Satisfaction. It would be appropriate to consider other factors in the model.

- SigmaXL also provides Standardized Residuals and Studentized (Deleted t) Residuals. The standardized residual is the residual, divided by an estimate of its standard deviation. This makes it easier to detect outliers. Standardized residuals greater than 3 and less than -3 are considered large (these outliers are highlighted in red). Studentized deleted residuals are

computed in the same way that standardized residuals are computed, except that the i^{th} observation is removed before performing the regression fit. This prevents the i^{th} observation from influencing the regression model, resulting in unusual observations being more likely to stand out.

- Click **Recall Last Dialog** (or press **F3**). Change the **Residual** type to *Standardized*. Click **OK**. The resulting Standardized Residual plots are displayed:



- Other diagnostic measures reported but not plotted include Leverage, Cook's Distance (Influence), and DFITS. Leverage is a measure of how far an individual X predictor value deviates from its mean. High leverage points can **potentially** have a strong effect on the estimate of regression coefficients. Leverage values fall between 0 and 1. Cook's distance and DFITS are overall measures of influence. An observation is said to be influential if removing the observation substantially changes the estimate of coefficients. Cook's distance can be thought of as the product of leverage and the standardized residual; DFITS as the product of leverage and the studentized residual.
- The easiest way to identify observations with high leverage and/or influence is to plot the measures on a run chart.

Multiple Regression with Continuous and Categorical Predictors

13. Click **Recall Last Dialog** (or press **F3**). Select *Customer Type*, click **Categorical Predictor (X) >>**. We will not discuss residuals further so you may wish to uncheck **Display Residual Plots**. Keep *Overall Satisfaction* as the **Numeric Response (Y)** and *Responsive to Calls* and *Ease of Communications* as the **Continuous Predictors (X)**:

14. Click **OK**. The resulting Multiple Regression report is shown:

Multiple Regression Model: Overall Satisfaction = (0.552345) + (0.427400) * Responsive to Calls + (0.409625) * Ease of Communications + (0.1

Model Summary:	
R-Square	90.58%
R-Square Adjusted	90.18%
S (Root Mean Square Error)	0.245199

Parameter Estimates:

Predictor Term	Coefficient	SE Coefficient	T	P	VIF	Tolerance
Constant	0.552345	0.120148	4.597	0.0000		
Responsive to Calls	0.427400	0.023788018	17.967	0.0000	1.212	0.825379
Ease of Communications	0.409625	0.031120872	13.162	0.0000	1.325	0.754950
Customer Type_2	0.132728	0.063914154	2.077	0.0406	1.655	0.604180
Customer Type_3	0.023141785	0.065217411	0.354841	0.7235	1.394	0.717173

Analysis of Variance for Categorical (Discrete) Predictors:

Predictor Term	DF	SS	MS	F	P
Customer Type	2	0.299651	0.149825574	2.492006584	0.0881

Analysis of Variance for Model:

Source	DF	SS	MS	F	P
Model	4	54.901	13.725	228.29	0.0000
Error	95	5.712	0.060122608		
Total (Model + Error)	99	60.612	0.612246		

Durbin-Watson Test for Autocorrelation in Residuals:

DW Statistic	1.730
P-Value Positive Autocorrelation	0.0888
P-Value Negative Autocorrelation	0.9137

15. The parameter estimates table now includes Customer Type 2 and Customer Type 3, but where is Customer Type 1? Since Customer Type is a discrete predictor, SigmaXL applies “dummy coding” (i.e. Customer Type 1 corresponds to Coded Variable 1 = 0 and Coded Variable 2 = 0; Customer Type 2 has Coded Variable 1 = 1 and Coded Variable 2 = 0; and Customer Type 3 has Coded Variable 1 = 0 and Coded Variable 2=1). Hence Customer Type 1 becomes the “hidden” or reference value. Other statistical software packages may use a different coding scheme based on -1, 0, +1 instead of 0, 1. The advantage of a 0, 1 coding scheme is the relative ease of interpretation when making predictions with the model.
16. In the **Predicted Response Calculator** enter the settings as shown:

Predicted Response Calculator:

Predictors	Enter Settings:	Predicted Response	Lower 95% CI	Upper 95% CI	Lower 95% PI	Upper 95% PI
Responsive to Calls	5	4.870198931	4.778492731	4.961905131	4.374854056	5.365543806
Ease of Communications	5					
Customer Type 2	1					
Customer Type 3	0					

Note that for Customer Type 1, you would enter **Customer Type 2 = 0** and **Customer Type 3 = 0**; for Customer Type 2, as shown, you entered **Customer Type 2 = 1** and **Customer Type 3 = 0**; for Customer Type 3 you would enter **Customer Type 2 = 0** and **Customer Type 3 = 1**.

17. Note the addition of the Analysis of Variance for Categorical (Discrete) Predictors. The Customer Type P-Value is .08, so we do not have strong evidence to keep this term in the model. However, many practitioners will use an alpha value of 0.1 as a criteria for removal. You are probably wondering why the P-Value for Customer Type is not lower given the results we saw earlier using ANOVA. The change in P-Value is due to the inclusion of *Responsive to Calls* and *Ease of Communications* in the model. We have higher scores for *Ease of Communications* and *Responsive to Calls* with Customer Type 2. Statistically, Customer Type is somewhat correlated to *Responsive to Calls* and *Ease of Communications* (VIF for Customer Type = 1.66 and 1.39).

Part U – Logistic Regression

Binary Logistic Regression

Binary Logistic Regression is used to analyze the relationship between one binary dependent variable (Y) and multiple independent numeric and/or discrete variables (X's). It is used to discover the relationship between the variables and create an empirical equation of the form:

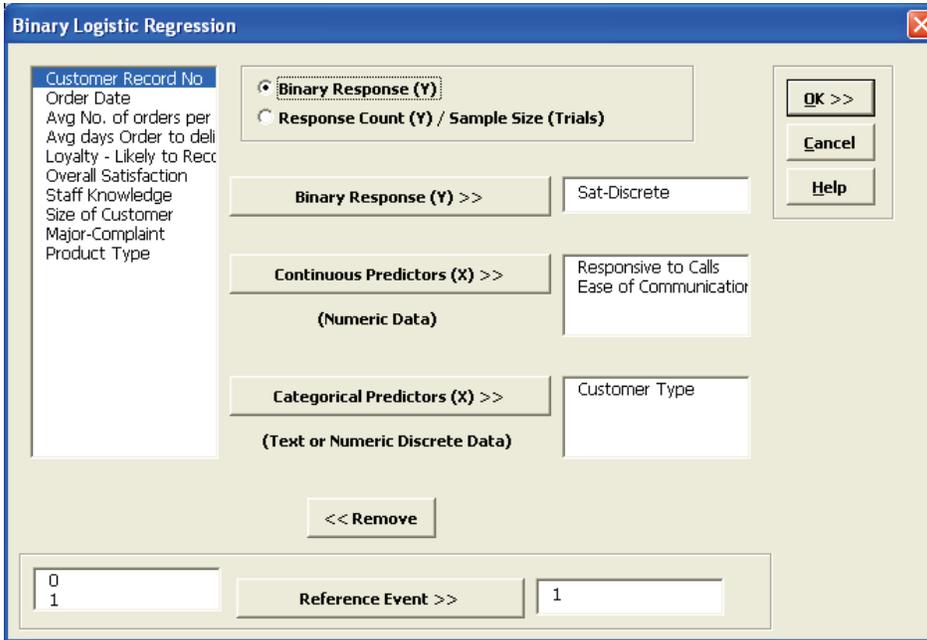
$$\ln(P_Y/(1-P_Y)) = b_0 + b_1*X_1 + b_2*X_2 + \dots + b_n*X_n$$

This equation can be used to predict an event probability Y value for a given set of input X values. SigmaXL uses the method of maximum likelihood to solve for the model coefficients and constant term. Statistical tests of hypothesis and odds-ratios are provided for the model coefficients. The odds-ratios identify change in likelihood of the event for one-unit change in X.

An example application from medical research would be Y=Disease (Yes/No) and X's = Age, Smoker (Yes/No), Number Years of Smoking and Weight. The model coefficient P-Values would indicate which X's are significant and the odds-ratios would provide the relative change in risk for each unit change in X.

We will analyze the familiar Customer Satisfaction data using Y=Discrete Satisfaction where the values have been coded such that an Overall Satisfaction score ≥ 3.5 is considered a 1, and scores < 3.5 are considered a 0. Please note, we are not advising that continuous data be converted to discrete data in actual practice, but simply using the Discrete Satisfaction score for continuity with the previous analysis.

1. Open **Customer Data.xlsx**. Click **Sheet 1** Tab (or press **F4** to activate last worksheet). Click **SigmaXL > Statistical Tools > Regression > Binary Logistic Regression**. If necessary, click **Use Entire Data Table**, click **Next**.
2. Select *Sat-Discrete*, click **Binary Response (Y) >>**; select *Responsive to Calls* and *Ease of Communications*, click **Continuous Predictors (X) >>**; select *Customer Type*, click **Categorical Predictors (X) >>**.



Note that **Response Count (Y)/Sample Size (Trials)** should be used when each record contains both the number of occurrences along with associated sample size. This is common when tracking daily quality data or performing design of experiments where each run contains a response of the number of defects and sample size.

3. Click **OK**. The resulting Binary Logistic report is shown:

Binary Logistic Regression Model: $\ln(P_Y/(1-P_Y)) = (-16.647) + (2.454) * \text{Responsive to Calls} + (2.366) * \text{Ease of Communications} + (0.546089) * \text{Logit Link}$

Response Summary: Sat-Discrete

Value	Count	Proportion	Reference Event
0	33	0.33	
1	67	0.67	x
Total	100		

Parameter Estimates:

Term	Coefficient	SE Coefficient	Z	P	Odds Ratio	Lower 95% Odds Ratio	Upper 95% Odds Ratio
Constant	-16.647	3.837	-4.339	0.0000			
Responsive to Calls	2.454	0.558678	4.392	0.0000	11.630	3.891	34.765
Ease of Communications	2.366	0.673792	3.512	0.0004	10.660	2.846	39.929
Customer Type_2	0.546089	1.076518129	0.507273	0.6120	1.726	0.209321	14.240
Customer Type_3	-1.131	0.990522	-1.141	0.2537	0.322850	0.046328807	2.250

Wald Estimates for Categorical (Discrete) Predictors:

Term	Chi-Square	DF	P
Customer Type	2.630	2	0.2685

Model Summary and Goodness-of-Fit Statistics:

Log-Likelihood	-19.883
Test that all slope coefficients are equal to zero:	
Likelihood Ratio Chi-Square (G)	87.070
DF	4
P-Value	0.0000
McFadden's Pseudo R-Square	68.65%

Observed and Predicted Outcomes:

Observed Outcome	Predicted Outcome		Row Total
	$\hat{Y} = 0$	$\hat{Y} = 1$	
Y = 0	29	4	33
Y = 1	3	64	67
Column Total	32	68	100
Percent Correctly Predicted:	93.00%		

4. The Likelihood Ratio P-Value $< .05$ tells us that the model is significant. The P-Values for the coefficients in the Parameter Estimates table confirm that Responsive to Calls and Ease of Communications are significant.
5. The P-Value in Wald Estimates for Categorical (Discrete) Predictors table tells us that Customer Type is not significant here.
Tip: Significance for categorical predictors should be based on the Wald Estimates not the P-Values given in the Parameter Estimates table.
6. Note that Customer Type 1 is not displayed in the Parameter Estimates table. This is the “hidden” reference value for Customer Type. Categorical predictors must have one level selected as a reference value. SigmaXL sorts the levels alphanumerically and selects the first level as the reference value.
7. Now we will rerun the binary logistic regression but remove *Customer Type* as a predictor. Press **F3** or click **Recall SigmaXL Dialog** to recall last dialog. Remove *Customer Type* by highlighting *Customer Type* and double-clicking (or press the **Remove** button).

The screenshot shows the "Binary Logistic Regression" dialog box. On the left is a list of variables: Customer Record No, Order Date, Avg No. of orders per, Avg days Order to deli, Loyalty - Likely to Recc, Overall Satisfaction, Staff Knowledge, Size of Customer, Major-Complaint, Product Type, and Customer Type. The "Binary Response (Y)" radio button is selected, and "Sat-Discrete" is entered in the field next to it. The "Continuous Predictors (X)" field contains "Responsive to Calls" and "Ease of Communication". The "Categorical Predictors (X)" field is empty. At the bottom, the "Reference Event" section shows "0" and "1" with "1" selected. Buttons for "OK >>", "Cancel", "Help", and "<< Remove" are present.

8. Click **OK**. The resulting Binary Logistic report is shown:

SigmaXL: Analyze Phase Tools

Binary Logistic Regression Model: $\ln(P_y/(1-P_y)) = [-17.703] + [2.462] * \text{Responsive to Calls} + [2.589] * \text{Ease of Communications}$
 Logit Link

Response Summary: Sat-Discrete

Value	Count	Proportion	Reference Event
0	33	0.33	
1	67	0.67	*
Total	100		

Parameter Estimates:

Term	Coefficient	SE Coefficient	Z	P	Odds Ratio	Lower 95% Odds Ratio	Upper 95% Odds Ratio
Constant	-17.703	3.947	-4.486	0.0000			
Responsive to Calls	2.462	0.559880	4.397	0.0000	11.726	3.913	35.133
Ease of Communications	2.589	0.668700	3.871	0.0001	13.314	3.590	49.375

Model Summary and Goodness-of-Fit Statistics:

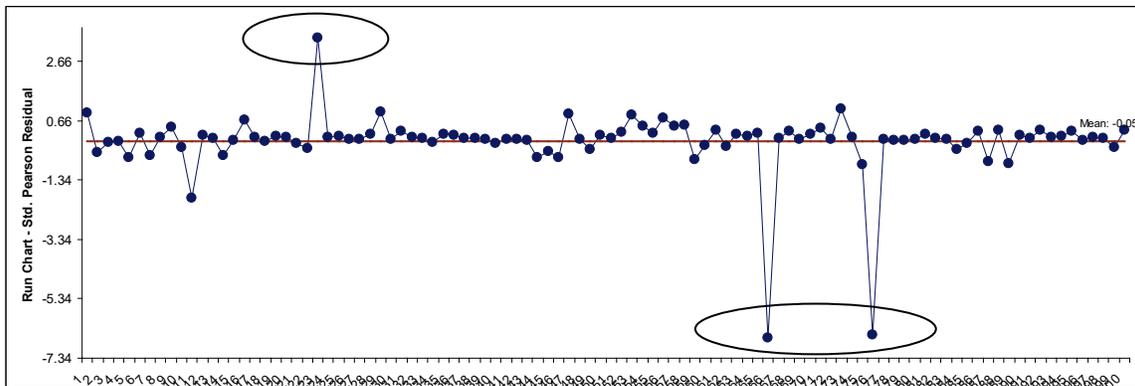
Log-Likelihood	-21.275
Test that all slope coefficients are equal to zero:	
Likelihood Ratio Chi-Square (G)	84.285
DF	2
P-Value	0.0000
McFadden's Pseudo R-Square	66.45%
Goodness-of-Fit Tests (P-Value < .05 indicates Lack-of-Fit):	
Pearson Residuals Chi-Square	112.63
DF	97
P-Value	0.1325
Deviance Residuals Chi-Square	42.551
DF	97
P-Value	1.0000
Hosmer-Lemeshow Chi-Square	20.187
DF	8
P-Value	0.0097
Measures of Association	
Concordant	2123
Discordant	86
Ties	2
Total	2211
Concordant Percent	96.02
Discordant Percent	3.89
Ties Percent	0.09
Goodman-Kruskal Gamma	0.922137
Somers' D	0.921303
Kendall's Tau-a	0.411515

Observed and Predicted Outcomes:

Observed Outcome	Predicted Outcome		Row Total
	$\hat{y} = 0$	$\hat{y} = 1$	
Y = 0	30	3	33
Y = 1	1	66	67
Column Total	31	69	100
Percent Correctly Predicted:	96.00%		

- The Odds Ratios in the Parameter Estimates table tell us that for every unit increase in Responsive to Calls we are 11.7 times more likely to obtain a satisfied customer. For every unit increase in Ease of Communications we are 13.3 times more likely to obtain a satisfied customer.
- McFadden's Pseudo R-Square mimics the R-square found in linear regression. This value varies between 0 and 1 but is typically much lower than the traditional R-squared value. A value less than 0.2 indicates a weak relationship; 0.2 to 0.4 indicates a moderate relationship; greater than 0.4 indicates a strong relationship. Here we have an R-square value of 0.66 indicating a strong relationship. This is also confirmed with the Percent Correctly Predicted value of 96%.

11. The Pearson, Deviance and Hosmer-Lemeshow Goodness of Fit tests are used to confirm if the binary logistic model fits the data well. P-Values < .05 for any of these tests indicate a significant lack of fit. Here the Hosmer-Lemeshow test is indicating lack of fit. Residuals analysis will help us to see where the model does not fit the data well.
12. The measures of association are used to indicate the relationship between the observed responses and the predicted probabilities. Larger values for Goodman-Kruskal Gamma, Somers' D and Kendall's Tau-a indicate that the model has better predictive ability.
13. The residuals report is given on the Sheet **Binary Logistic Residuals**. Three types of residuals are provided: Pearson, Standardized Pearson and Deviance. The Standardized Pearson Residual is most commonly used and is shown here plotted on a Run Chart. To create, click **SigmaXL > Graphical Tools > Run Chart**, check **Use Entire Data Table** to select the *Residuals* data, click **Next**, select *Std. Pearson Residual* as the **Numeric Data Variable (Y)**. Click **OK**.



Any Standardized Pearson Residual value that is less than -3 or greater than +3 is considered extreme and should be investigated. There are 3 such outliers here: rows 24, 67, and 77 in the residuals table. The +3.4 value indicates that the predicted event probability was low (.08) but the actual result was a 1. The -6.6 value indicates that the predicted event probability was high (.98) but the actual result was a 0. The large negative residuals have high Responsive to Calls and Ease of Communications but dissatisfied customers. The reasons for these discrepancies should be explored further but we will not do so here.

14. Reselect the **Binary Logistic** sheet. Scroll over to display the Event Probability calculator:

Response Event Probability:

Predictors	Enter Settings:	Predicted Event Probability
Responsive to Calls		
Ease of Communications		

This calculator provides a predicted Event Probability for given values of X (in this case the probability of a satisfied customer). Enter the values 3,3; 4,4; 5,5 as shown:

Response Event Probability:

Predictors	Enter Settings:	Predicted Event Probability
Responsive to Calls	3	0.072336996
Ease of Communications	3	

Response Event Probability:

Predictors	Enter Settings:	Predicted Event Probability
Responsive to Calls	4	0.924088278
Ease of Communications	4	

Response Event Probability:

Predictors	Enter Settings:	Predicted Event Probability
Responsive to Calls	5	0.999474064
Ease of Communications	5	

If Responsive to Calls and Ease of Communications are both equal to 3, the probability of a satisfied customer is only .07 (7%); if Responsive to Calls and Ease of Communications are both equal to 5, the probability of a satisfied customer is .9995 (99.95%)

15. Note that if the calculator includes predictors that are categorical (discrete), enter a 0 or 1 to denote the selected level as shown below (using the original analysis which included Customer Type):

Response Event Probability:

Predictors	Enter Settings:	Predicted Event Probability
Responsive to Calls	3.02	0.719225458
Ease of Communications	4.07	
Customer Type_2	1	
Customer Type_3	0	

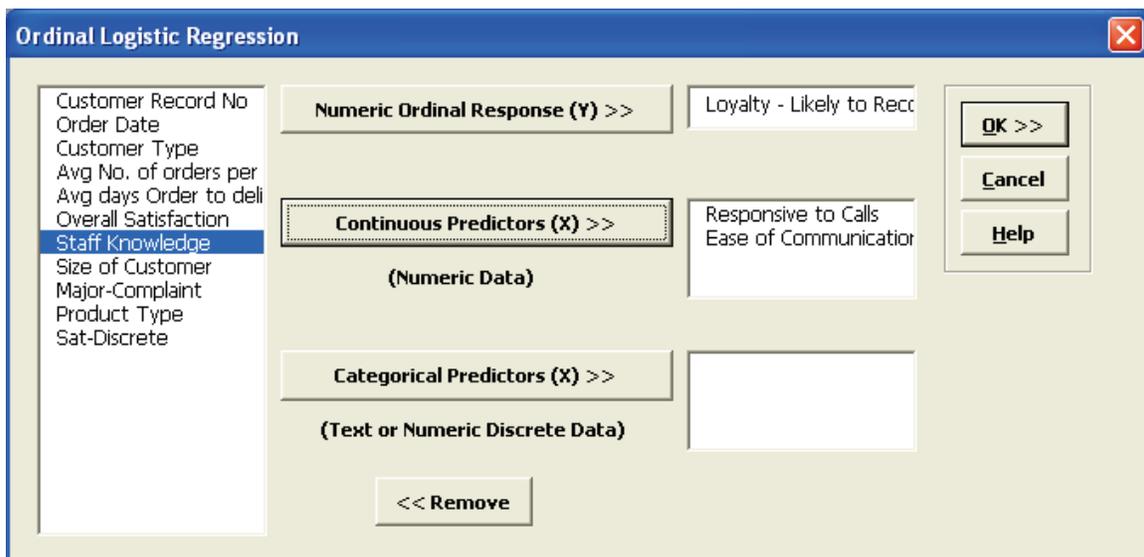
If we wanted to select Customer Type 1, enter a 0 for both **Customer Types 2 and 3**. Customer Type 1 is the hidden reference value.

Ordinal Logistic Regression

Ordinal Logistic Regression is used to analyze the relationship between one ordinal dependent variable (Y) and multiple independent continuous and/or discrete variables (X's).

We will analyze the Customer Satisfaction data using Y=Loyalty –Likely to Recommend score which contains ordinal integer values from 1 to 5, where 5 indicates that the customer is very likely to recommend us and 1 indicates that they are very likely to not recommend us.

1. Open **Customer Data.xlsx**. Click **Sheet 1** Tab (or press **F4** to activate last worksheet). Click **SigmaXL > Statistical Tools > Regression > Ordinal Logistic Regression**. If necessary, click **Use Entire Data Table**, click **Next**.
2. Select *Loyalty – Likely to Recommend*, click **Numeric Ordinal Response (Y) >>**; select *Responsive to Calls* and *Ease of Communications*, click **Continuous Predictors (X) >>**.



3. Click **OK**. The resulting Ordinal Logistic report is shown:

SigmaXL: Analyze Phase Tools

Ordinal Logistic Regression Model: $\ln(\text{Py}/(1-\text{Py})) = \text{Constant} + (-1.011305417) * \text{Responsive to Calls} + (-0.790025) * \text{Ease of Communications}$
 Logit Link

Response Summary: Loyalty - Likely to Recommend

Value	Count	Proportion
1	2	0.02
2	10	0.1
3	33	0.33
4	43	0.43
5	12	0.12
Total	100	

Parameter Estimates:

Term	Coefficient	SE Coefficient	Z	P	Odds Ratio	Lower 95% Odds Ratio	Upper 95% Odds Ratio
Constant 1	1.661	1.110374236	1.496	0.1348			
Constant 2	4.138	1.049364975	3.944	0.0001			
Constant 3	6.721	1.199	5.607	0.0000			
Constant 4	9.519	1.351	7.048	0.0000			
Responsive to Calls	-1.01130542	0.208727	-4.845	0.0000	0.363744	0.241615	0.547605
Ease of Communications	-0.790025	0.253997	-3.110	0.0019	0.453833	0.275861	0.746626

Model Summary and Goodness-of-Fit Statistics:

Log-Likelihood	-105.24
Test that all slope coefficients are equal to	
Likelihood Ratio Chi-Square (G)	47.850
DF	2
P-Value	0.0000
McFadden's Pseudo R-Square	0.185221
Goodness-of-Fit Tests	
Pearson Residuals Chi-Square	279.41
DF	394
P-Value	1.0000
Deviance Residuals Chi-Square	210.49
DF	394
P-Value	1.0000
Measures of Association	
Concordant	2702
Discordant	694
Ties	11
Total	3407
Concordant Percent	79.307
Discordant Percent	20.370
Ties Percent	0.322865
Goodman-Kruskal	
Gamma	0.591284
Somers' D	0.589375
Kendall's Tau-a	0.405657

Observed and Predicted Outcomes:

Observed Outcome	Predicted Outcomes					Row Total
	Y = 1	Y = 2	Y = 3	Y = 4	Y = 5	
Y = 1	0	2	0	0	0	2
Y = 2	0	2	6	2	0	10
Y = 3	0	2	14	17	0	33
Y = 4	0	1	10	32	0	43
Y = 5	0	0	0	12	0	12
Column Total	0	7	30	63	0	100
Percent Correctly Predicted:	48.00%					

- The Likelihood Ratio P-Value < .05 tells us that the model is significant. The low P-Values for the coefficients confirm that Responsive to Calls and Ease of Communications are significant.
- The Odds Ratios tell us that for every one-unit increase in Responsive to Calls the chance of a Loyalty score of 1 versus 2 (or 2 versus 3, etc.) is reduced by a multiple of 0.36. This is not very intuitive but will be easy to see when we use the Response Outcome Probability calculator.
- McFadden's Pseudo R-Square value is 0.185 indicating that this is a weak (but close to moderate) degree of association. This is also confirmed with the Percent Correctly Predicted value of 48%.
- The Pearson and Deviance Goodness of Fit (GOF) tests are used to confirm if the ordinal logistic model fits the data well. P-Values < .05 would indicate a significant lack of fit. Given that the GOF P-Values are greater than .05, we conclude that there is no significant lack of fit.

8. Scroll across to the Response Outcome Probability calculator. This calculator provides predicted outcome (event) probabilities for given values of X (in this case the probability of a satisfied customer). Enter the values 3,3; 4,4; 5,5 as shown:

Response Outcome Probability:

Predictors	Enter Settings:	Outcome	Predicted Cumulative Probability	Predicted Probability for each Level
Responsive to Calls	3	1	0.023126734	0.023126734
Ease of Communications	3	2	0.2199757	0.196848965
		3	0.788637058	0.568661358
		4	0.983934474	0.195297416
		5	1	0.016065526

Response Outcome Probability:

Predictors	Enter Settings:	Outcome	Predicted Cumulative Probability	Predicted Probability for each Level
Responsive to Calls	4	1	0.003892909	0.003892909
Ease of Communications	4	2	0.044483294	0.040590386
		3	0.381166233	0.336682938
		4	0.90999329	0.528827057
		5	1	0.09000671

Response Outcome Probability:

Predictors	Enter Settings:	Outcome	Predicted Cumulative Probability	Predicted Probability for each Level
Responsive to Calls	5	1	0.000644733	0.000644733
Ease of Communications	5	2	0.007626511	0.006981777
		3	0.092294817	0.084668306
		4	0.625327323	0.533032506
		5	1	0.374672677

9. Referring to the Predicted Probability for each Level, if **Responsive to Calls** and **Ease of Communications** are both equal to 3, we would expect to see typical loyalty scores of 3 (57%) with some at 2 (20%) and 4 (20%); if **Responsive to Calls** and **Ease of Communications** are both equal to 4, we would expect typical loyalty scores of 4 (53 %) with some at 3 (34%) and 5 (9%); if **Responsive to Calls** and **Ease of Communications** are both equal to 5, we would expect typical loyalty scores of 4 (53 %) with some at 3 (8.5%) and 5 (38%).

SigmaXL: Improve Phase Tools: Design of Experiments (DOE)

Part A – Overview of Basic Design of Experiments (DOE) Templates

The DOE templates are similar to the other SigmaXL templates: simply enter the inputs and resulting outputs are produced immediately. The DOE templates provide common 2-level designs for 2 to 5 factors. These basic templates are ideal for training, but use **SigmaXL > Design of Experiments > 2-Level Factorial/Screening Designs** to accommodate up to 19 factors with randomization, replication and blocking.

Click **SigmaXL > Design of Experiments > Basic DOE Templates** to access these templates:

- Two-Factor, 4-Run, Full-Factorial
- Three-Factor, 4-Run, Half-Fraction, Res III
- Three-Factor, 8-Run, Full-Factorial
- Four-Factor, 8-Run, Half-Fraction, Res IV
- Four-Factor, 16-Run, Full-Factorial
- Five-Factor, 8-Run, Quarter-Fraction, Res III
- Five-Factor, 16-Run, Half-Fraction, Res V

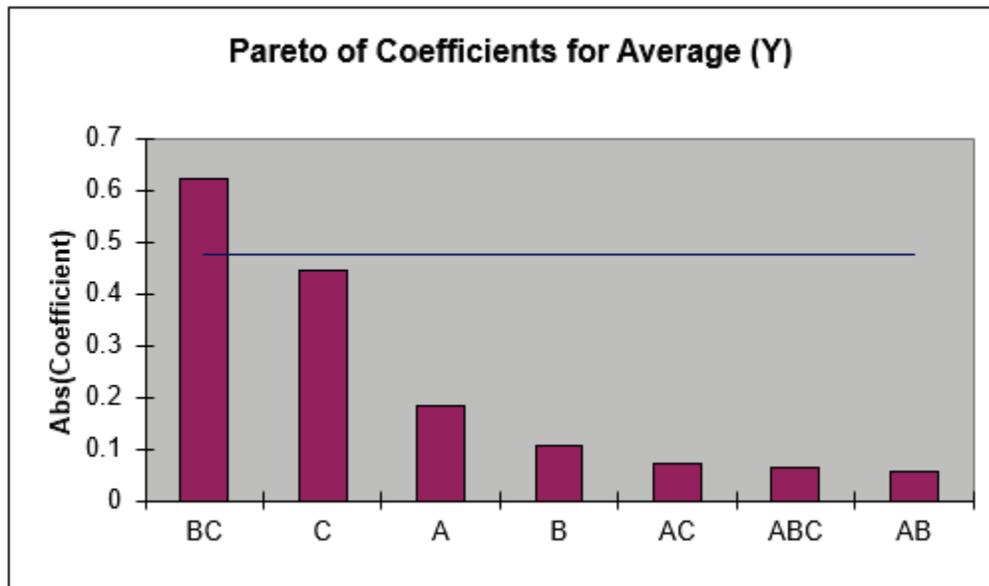
After entering the template data, main effects and interaction plots may be created by clicking **SigmaXL > Basic DOE Templates > Main Effects & Interaction Plots**. The DOE template must be the active worksheet.

DOE Templates are protected worksheets by default, but this may be modified by clicking **SigmaXL > Help > Unprotect Worksheet**.

Advanced analysis is available, but this requires that you unprotect the DOE worksheet. The following example shows how to use Excel's Equation Solver and SigmaXL's Multiple Regression in conjunction with a DOE template. **Caution:** If you unprotect the worksheet, do not change the worksheet title (e.g. **Three-Factor, Two-Level, 8-Run, Full-Factorial Design of Experiments**). This title is used by the Main Effects & Interaction Plots to determine appropriate analysis. Also, do not modify any cells with formulas.

Part B – Three Factor Full Factorial Example Using DOE Template

1. Open the file **DOE Example - Robust Cake.xlsx**. This is a Robust Cake Experiment adapted from the Video "Designing Industrial Experiments" by Box, Bisgaard and Fung.
2. The response is Taste Score (on a scale of 1-7 where 1 is "awful" and 7 is "delicious").
3. The five Outer Array Reps have different Cooking Time and Temperature Conditions. The outer array was a two-factor, full-factorial plus center point, hence 5 replications.
4. The goal is to Maximize Mean and Minimize StDev of the Taste Score.
5. The X factors are Flour, Butter, and Egg. Actual low and high settings are not given in the video, so we will use coded -1 and +1 values. We are looking for a combination of Flour, Butter, and Egg that will not only taste good, but consistently taste good over a wide range of Cooking Time and Temperature conditions.
6. Scroll down to view the Pareto of Abs. Coefficients for Average (Y).



7. The BC (Butter * Egg) interaction is clearly the dominant factor. The bars above the 95% confidence blue line indicate the factors that are statistically significant; in this case only BC is significant. Keep in mind that this is an initial analysis. Later, we will show how to do a more powerful Multiple Regression analysis on this data. (Also, the Rule of Hierarchy states that if an interaction is significant, we must include the main effects in the model used.)

8. The significant BC interaction is also highlighted in red in the table of Effects and Coefficients:

Calculation of Effects and Coefficients for Average (Y):

	Constant	A	B	C	AB	AC	BC	ABC
Avg(Avg(Y)) @ +1:		4.865	4.79	5.125	4.74	4.755	4.06	4.62
Avg(Avg(Y)) @ -1:		4.5	4.575	4.24	4.625	4.61	5.305	4.745
Effect (Delta):		0.365	0.215	0.885	0.115	0.145	-1.245	-0.125
Coefficient (Delta/2):	4.6825	0.1825	0.1075	0.4425	0.0575	0.0725	-0.6225	-0.0625
SE Coefficient:	0.2325	0.2325	0.2325	0.2325	0.2325	0.2325	0.2325	0.2325
T-value	20.1397849	0.7849	0.4624	1.9032	0.24731	0.31183	-2.6774	-0.2688
P-value	9.6928E-20	0.4383	0.6469	0.066	0.80625	0.75719	0.01161	0.7898

9. The R-Square value is given as 27%. This is very poor for a Designed Experiment. Typically, we would like to see a minimum of 50%, with > 80% desirable.

R-Square:	27.03%
R-Square Adj.:	11.06%
S	1.4705

The reason for the poor R-square value is the wide range of values over the Cooking Temperature and Time conditions. In a robust experiment like this, it is more appropriate to analyze the mean response as an individual value rather than as five replicate values. The Standard Deviation as a separate response will also be of interest.

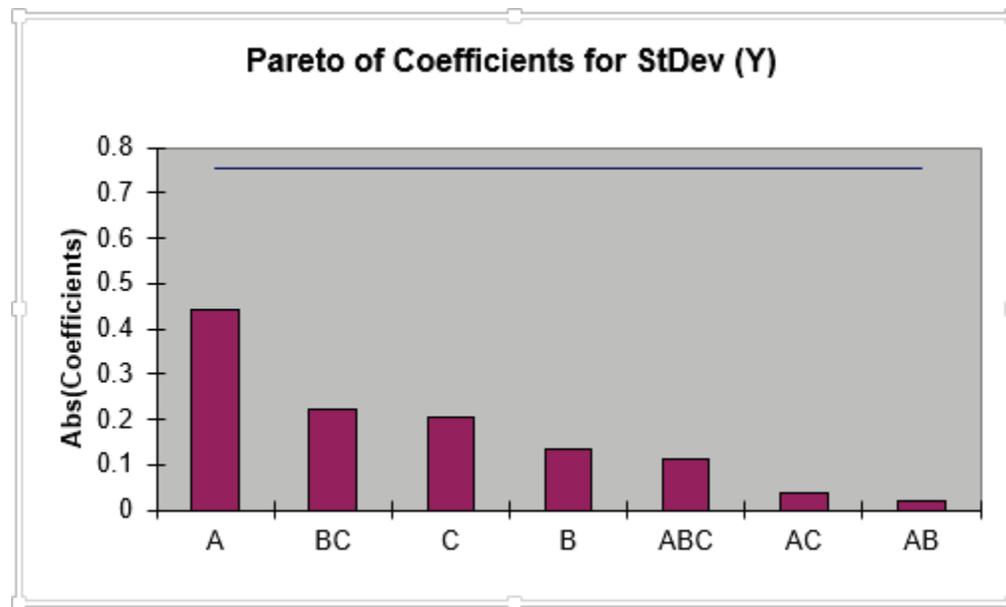
10. If the Responses are replicated, SigmaXL draws the blue line on the Pareto Chart using an estimate of experimental error from the replicates. If there are no replicates, an estimate called Lenth's Pseudo Standard Error is used.
11. If the 95% Confidence line for coefficients were to be drawn using Lenth's method, the value would be 0.409 as given in the table:

Lenth's Pseudo Standard Error (PSE) Analysis for Unreplicated Data:

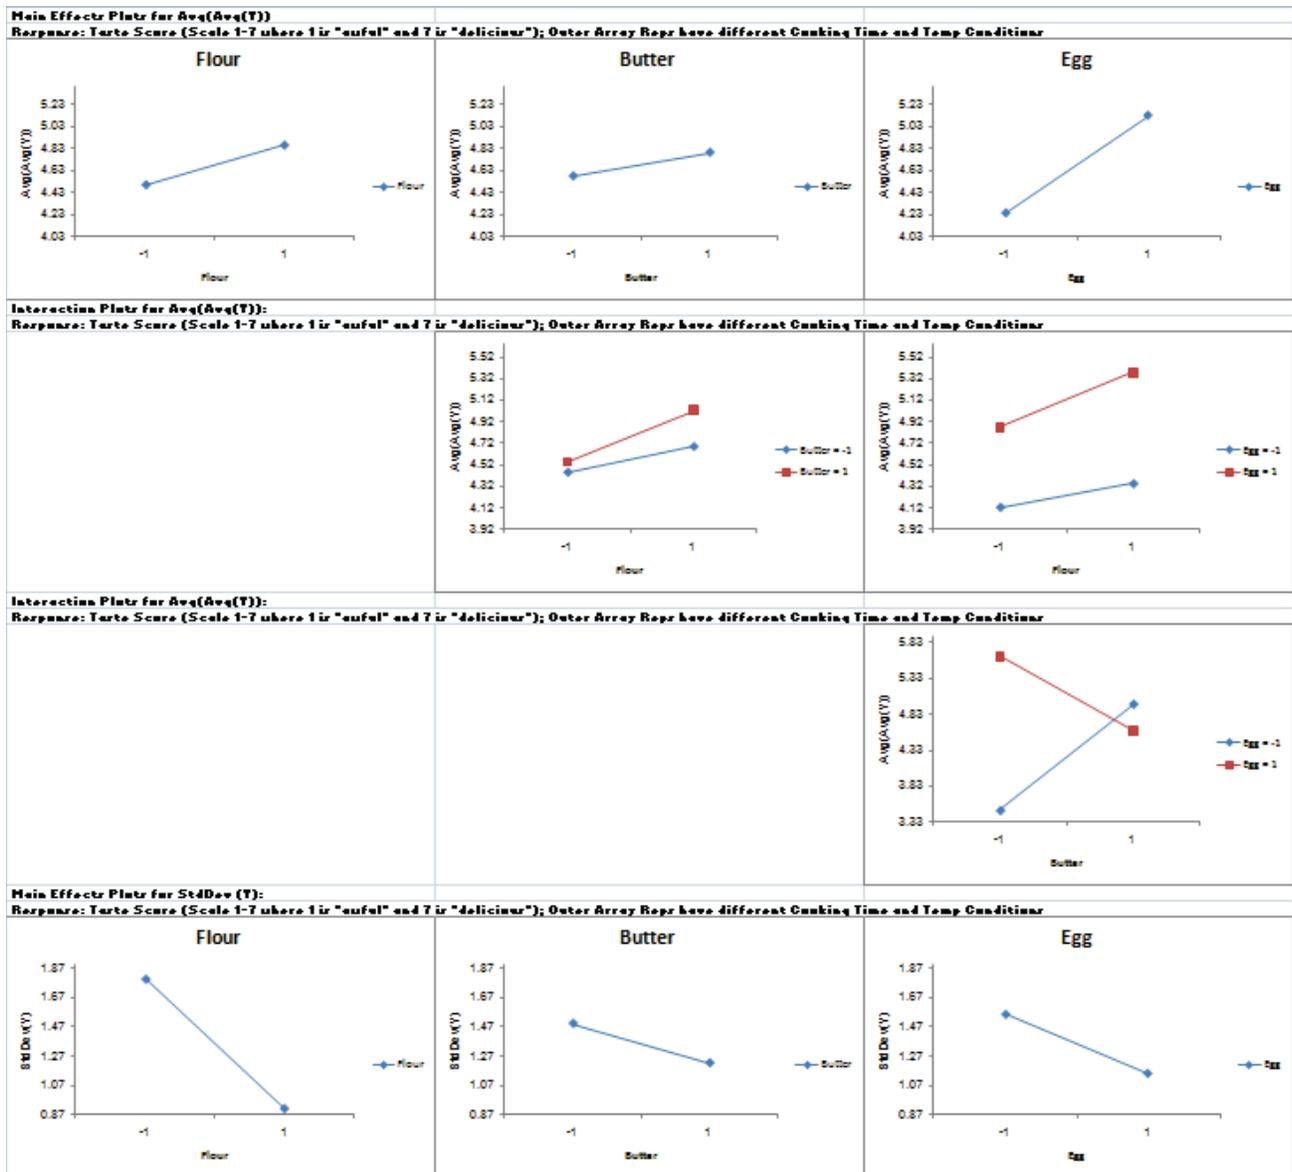
Lenth's PSE for Coefficients:	0.10875
Lenth's Margin of Error for Coefficients (95% Conf. Level):	0.40935
Lenth's Margin of Error for Effects (95% Conf. Level):	0.8187

This would show factor C as significant.

12. Scroll down to view the Pareto of Coefficients for StDev(Y).



13. The A (Flour) main effect is clearly the dominant factor, but it does not initially appear to be statistically significant (based on Lenth's method). Later, we will show how to do a more powerful Regression analysis on this data.
14. The Pareto chart is a powerful tool to display the relative importance of the main effects and interactions, but it does not tell us about the direction of influence. To see this, we must look at the main effects and interaction plots. Click **SigmaXL > Design of Experiments > Basic DOE Templates > Main Effects & Interaction Plots**. The resulting plots are shown below:



15. The Butter*Egg two-factor interaction is very prominent here. Looking at only the Main Effects plots would lead us to conclude that the optimum settings to maximize the average taste score would be Butter = +1, and Egg = +1, but the interaction plot tells a very different story. The correct optimum settings to maximize the taste score is Butter = -1 and Egg = +1.

16. Since Flour was the most prominent factor in the Standard Deviation Pareto, looking at the Main Effects plots for StDev, we would set Flour = +1 to minimize the variability in taste scores. The significance of this result will be demonstrated using Regression analysis.

17. Click on the Sheet **Three-Factor 8-Run DOE**. At the **Predicted Output for Y**, enter *Flour* = 1, *Butter* = -1, *Egg* = 1 as shown:

Factor	Factor Name	Low	High
A	Flour	-1	1
B	Butter	-1	1
C	Egg	-1	1

Predicted Output for Y:

Factor Name	Enter Actual Factor Setting - uncoded	Factor setting coded	Y-hat:	S-hat:
	1	1	5.9	0.68191
	-1	-1		
	1	1		

The predicted average (Y-hat) taste score is 5.9 with a predicted standard deviation (S-hat) of 0.68. Note that this prediction equation includes all main effects, two-way interaction, and the three-way interaction.

Multiple Regression and Excel Solver (Advanced Topics):

18. In order to run Multiple Regression analysis, we will need to unprotect the worksheet. Click **SigmaXL > Help > Unprotect Worksheet**.

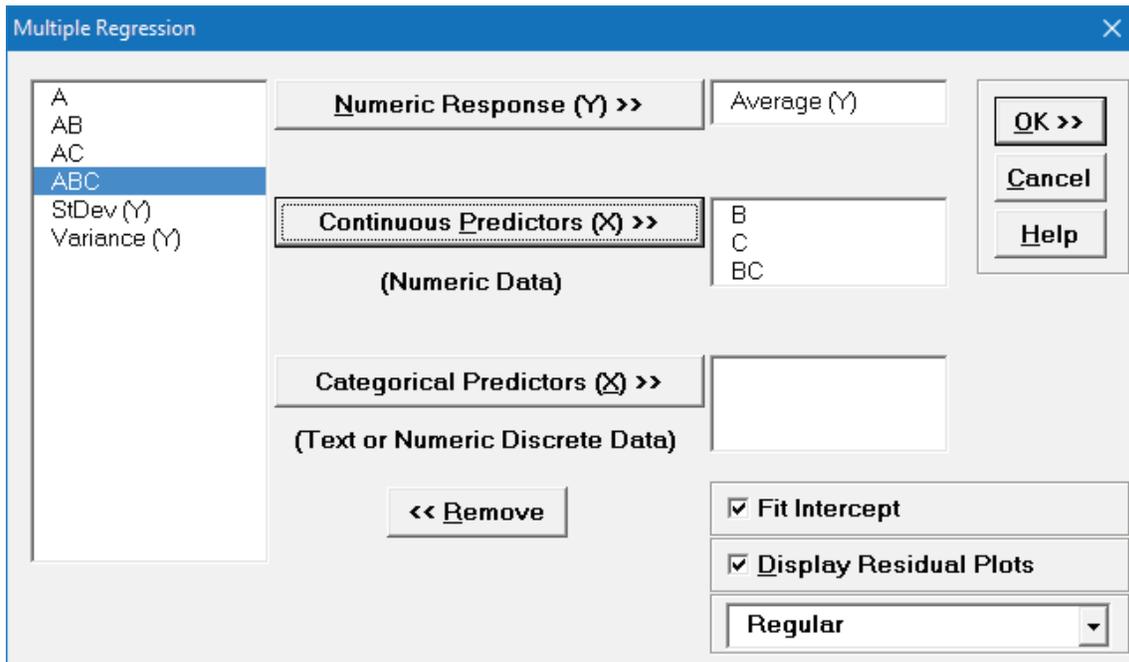
19. In the **Coded Design Matrix**, highlight columns A to ABC, and the calculated responses as shown:

Coded Design Matrix:

Standard Run Order	Actual Run Order	A	B	C	AB	AC	BC	ABC	Average (Y)	StDev (Y)	Variance (Y)
1	6	-1	-1	-1	1	1	1	-1	3.52	2.4499	6.002
2	2	1	-1	-1	-1	-1	1	1	3.5	1.38022	1.905
3	7	-1	1	-1	-1	1	-1	1	4.74	1.47919	2.188
4	1	1	1	-1	1	-1	-1	-1	5.2	0.93541	0.875
5	5	-1	-1	1	1	-1	-1	1	5.38	1.45499	2.117
6	8	1	-1	1	-1	1	-1	-1	5.9	0.68191	0.465
7	3	-1	1	1	-1	-1	1	-1	4.36	1.81742	3.303
8	4	1	1	1	1	1	1	1	4.86	0.66558	0.443

20. Click **SigmaXL > Statistical Tools > Regression > Multiple Regression**. Click **Next**.

21. Select *Average (Y)*, click **Numeric Response (Y) >>**; holding the CTRL key, select *B, C,* and *BC*; click **Continuous Predictors (X) >>** as shown:



22. Click **OK**. The resulting regression report is shown:

Multiple Regression Model: $Average (Y) = (4.683) + (0.1075) * B + (0.442500) * C + (-0.622500) * BC$

Model Summary:

R-Square	92.85%
R-Square Adjusted	87.50%
S (Root Mean Square Error)	0.302572

Parameter Estimates:

Predictor Term	Coefficient	SE Coefficient	T	P	VIF	Tolerance
Constant	4.683	0.106975464	43.772	0.0000		
B	0.1075	0.106975464	1.005	0.3718	1	1.00000
C	0.442500	0.106975464	4.136	0.0144	1	1.00000
BC	-0.622500	0.106975464	-5.819	0.0043	1	1.00000

Analysis of Variance for Model:

Source	DF	SS	MS	F	P
Model	3	4.759	1.586	17.327	0.0093
Error	4	0.366200	0.09155		
Pure Error	4	0.366200	0.09155		
Total (Model + Error)	7	5.125	0.732164		

Durbin-Watson Test for Autocorrelation in Residuals:

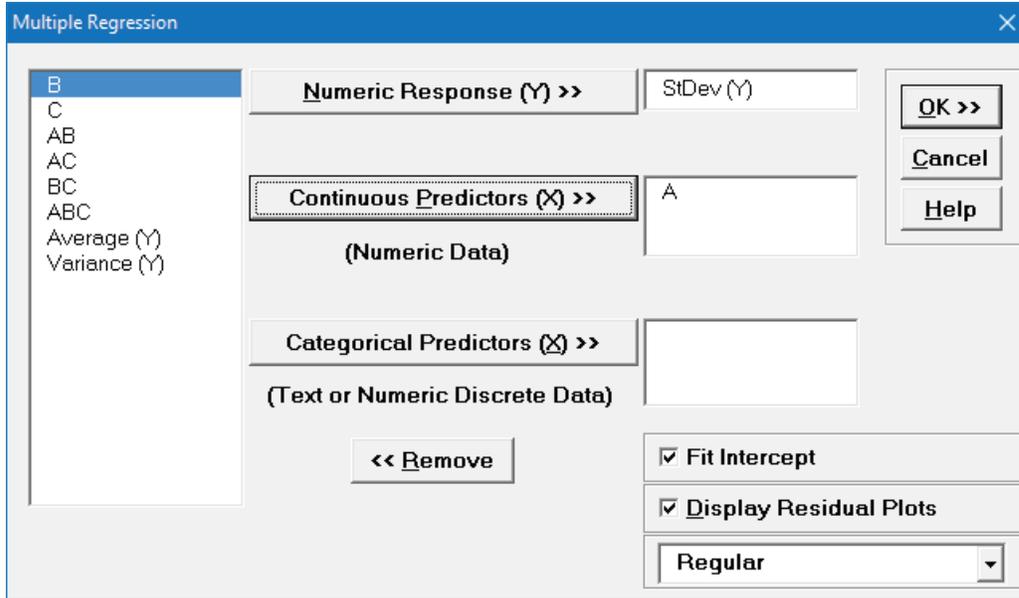
DW Statistic	3.498
P-Value Positive Autocorrelation	0.9670
P-Value Negative Autocorrelation	0.0000

Note that the R-square value of 92.85% is much higher than the earlier result of 27%. This is due to our modeling the mean response value rather than considering all data in the outer array. Note also that the C main effect now appears as significant.

23. Click on the Sheet **Three-Factor 8-Run DOE**.

24. With the **Coded Design Matrix** highlighted as before, click **SigmaXL > Statistical Tools > Regression > Multiple Regression**. Click **Next**.

25. Select *StDev (Y)*, click **Numeric Response (Y) >>**; select A, click **Continuous Predictors (X) >>** as shown:



26. Click **OK**. The resulting regression report is shown:

Multiple Regression Model: $StDev (Y) = (1.358) + (-0.442296) * A$

Model Summary:

R-Square	61.54%
R-Square Adjusted	55.13%
S (Root Mean Square Error)	0.403733

Parameter Estimates:

Predictor Term	Coefficient	SE Coefficient	T	P	VIF	Tolerance
Constant	1.358	0.142741	9.514	0.0001		
A	-0.442296	0.142741	-3.099	0.0212	1	1.00000

Analysis of Variance for Model:

Source	DF	SS	MS	F	P
Model	1	1.565	1.565	9.601	0.0212
Error	6	0.978003	0.163000		
Pure Error	6	0.978003	0.163000		
Total (Model + Error)	7	2.543	0.363287		

Durbin-Watson Test for Autocorrelation in Residuals:

DW Statistic	1.071225365
P-Value Positive Autocorrelation	0.1278
P-Value Negative Autocorrelation	0.9821

Note that Factor A (Flour) now shows as a statistically significant factor affecting the Standard Deviation of Taste Score.

27. Now we will use Excel’s Equation Solver to verify the optimum settings determined using the Main Effects and Interaction Plots.

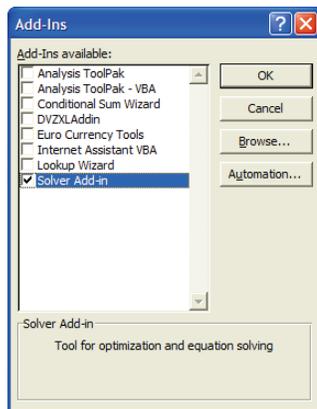
28. Click on the Sheet **Three-Factor 8-Run DOE**. At the **Predicted Output for Y**, enter 1 for *Flour*. We are setting this as a constraint, because Flour = +1 minimizes the Standard Deviation. Reset the *Butter* and *Egg* to 0 as shown:

Factor	Factor Name	Low	High
A	Flour	-1	1
B	Butter	-1	1
C	Egg	-1	1

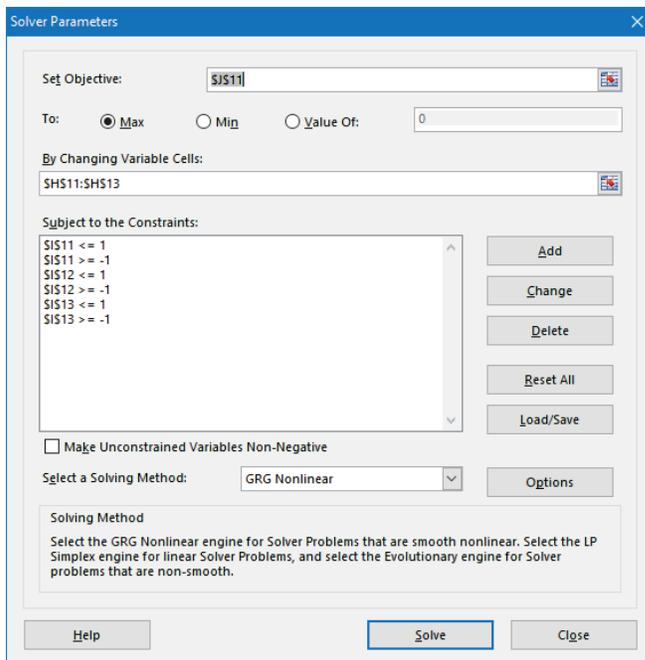
Predicted Output for Y:

Factor Name	Enter Actual Factor Setting - uncoded	Factor setting coded	Y-hat:	S-hat:
	1	1	4.865	0.91578
	0	0		
	0	0		

29. In Excel, click **File > Options > Add-Ins > Manage Excel Add-Ins**, click **Go**. Ensure that the **Solver Add-in** is checked. If the Solver Add-in does not appear in the Add-ins available list, you will need to re-install Excel to include all add-ins.



30. Click **OK**. Click **Data > Solver**. Set the **Solver Parameters** as shown:



Cell J11 is the Y-hat, predicted average taste score. Solver will try to maximize this value. Cells H11 to H13 are the Actual Factor Settings to be changed. Cells I11 to I13 are the Coded Factor settings where the following constraints are given: I11=1; I12 >= -1; I12 <= 1; I13 >= -1; I13 <=1.

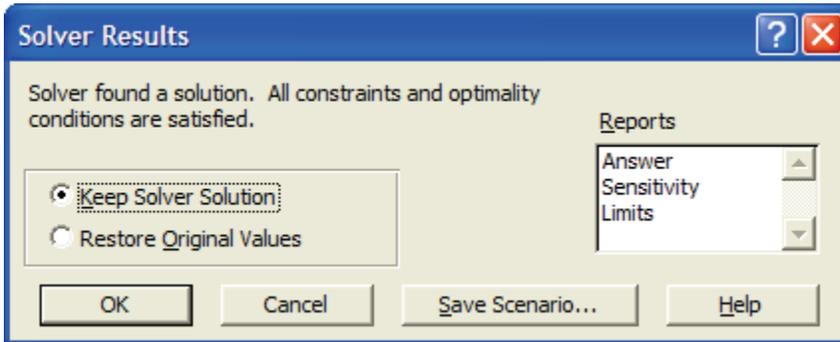
31. Click **Solve**. The solver results are given in the **Predicted Output for Y** as *Butter* = -1 and *Egg* = 1.

Factor	Factor Name	Low	High
A	Flour	-1	1
B	Butter	-1	1
C	Egg	-1	1

Predicted Output for Y:

Factor Name	Enter Actual Factor Setting - uncoded	Factor setting coded	Y-hat:	S-hat:
	1	1	5.9	0.68191
	-1	-1		
	1	1		

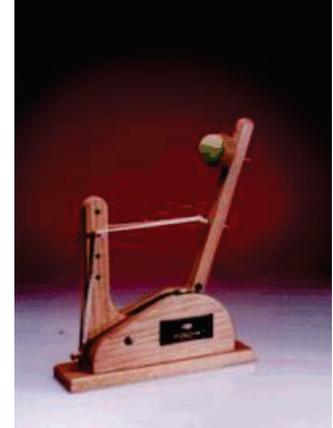
32. Solver indicates that a solution is found:



33. Click **OK** to keep the solution.

Part C – Design and Analysis of Catapult Full Factorial Experiment

Catapults are frequently used in Six-Sigma or Design of Experiments training. They are a powerful teaching tool and make the learning fun. If you have access to a catapult, we recommend that you perform the actual experiment and use your own data. Of course, you can also follow along using the data provided. The response variable (Y) is distance, with the goal being to consistently hit a target of 100 inches.



1. Click **SigmaXL > Design of Experiments > 2-Level Factorial/Screening > 2-Level Factorial/Screening Designs**.
2. The Number of X Factors can be 2 to 19. Using process knowledge, we will limit ourselves to 3 factors: *Pull Back Angle*, *Stop Pin* and *Pin Height*. Pull Back will be varied from 160 to 180 degrees, Stop Pin will be positions 2 and 3 (count from the back), and Pin Height will be positions 2 and 3 (count from the bottom).
3. Select **Number of Factors = 3**.
4. The available designs are then given as: **4-Run, 2^{**}(3-1), 1/2 Fraction, Res III** and **8-Run, 2^{**}3, Full-Factorial**. If we had more than 5 factors, a Resolution III or Plackett-Burman Screening design would typically be used. Here we will choose the **8-Run, 2^{**}3, Full-Factorial** design.

2-Level Factorial/Screening Design of Experiments

Number of Factors: 3

Select Design: 8-Run, 2**3, Full-Factorial

Number of Replicates: 1

Power Information (based on # of runs and replicates):

These power calculations assume 3 center points:
 Very Low Power to detect Effect = 1*StDev (1-Beta < 0.5);
 Very Low Power to detect Effect = 2*StDev (1-Beta < 0.5);
 Medium Power to detect Effect = 3*StDev (0.8 <= 1-Beta < 0.95).

Number of Blocks: 1

Number of Center Points per Block: 0

Randomize Runs

Factor Names and Level Settings:

	Name	Low	High
A:	A	-1	1
B:	B	-1	1
C:	C	-1	1

Number of Responses: 1

Response Name

Y1: Y1

OK>> Cancel Help Reset

Notes: Design Generators and Aliasing of Effects will be reported for Fractional Factorial designs. When the Number of Factors is 9 or higher, Factor Name "I" is not used, in order to avoid confusion with the Fractional Factorial Defining Relation "I".

The Power Information is presented to assist the user with selection of number of runs and replicates for the design, so that one can see the trade-off between experimental cost and sensitivity to detect Effects of interest.

Power (1-Beta) < 0.5 is considered as *Very Low Power* to detect Effect.

Power (1-Beta) >= 0.5 and < 0.8 is considered as *Low Power* to detect Effect.

Power (1-Beta) >= 0.8 and < 0.95 is considered as *Medium Power* to detect Effect.

Power (1-Beta) >= 0.95 and < 0.99 is considered as *High Power* to detect Effect.

Power (1-Beta) >= 0.99 is considered as *Very High Power* to detect Effect.

The power calculations require an estimate of experimental error, so when Replicates = 1 an assumption of 3 center points is used in order to give an approximate reference level of power (regardless of the value for Number of Center Points per Block).

5. This design currently shows the following:

These power calculations assume 3 center points:

Very Low Power to detect Effect = 1*StDev (1-Beta < 0.5);

Very Low Power to detect Effect = 2*StDev (1-Beta < 0.5);

Medium Power to detect Effect = 3*StDev (0.8 <= 1-Beta < 0.95).

We would like to have medium power to detect an Effect = 2*StDev. Change the **Number of Replicates** to 2. The Power Information is now:

Very Low Power to detect Effect = $1 \times \text{StDev}$ ($1 - \text{Beta} < 0.5$);
 Medium Power to detect Effect = $2 \times \text{StDev}$ ($0.8 \leq 1 - \text{Beta} < 0.95$);
 Very High Power to detect Effect = $3 \times \text{StDev}$ ($1 - \text{Beta} \geq 0.99$).

We will therefore choose two replicates. The number of replicates will always be a tradeoff between the desired power and the cost of the experimental runs.

Specify 2 or more blocks if there are constraints such as the number of runs per day or some other known external “nuisance” variable (like 2 different catapults or 2 operators). Here we will keep **Blocks** = 1 (i.e. no Blocking).

Center Points are useful to provide an estimate of experimental error with unreplicated designs, and allow detection of curvature. Typically, 3 to 5 center points are used. Here we will not use center points because we have replicated the design twice and do not expect significant curvature in the distance response. Furthermore, center points could not be set for Pin Height and Stop Pin (without drilling additional holes!).

6. Complete the **Factor Names, Level Settings** and **Response Name** as shown:

2-Level Factorial/Screening Design of Experiments

Number of Factors: 3

Select Design: 8-Run, 2**3, Full-Factorial

Number of Replicates: 2

Power Information (based on # of runs and replicates):
 Very Low Power to detect Effect = $1 \times \text{StDev}$ ($1 - \text{Beta} < 0.5$);
 Medium Power to detect Effect = $2 \times \text{StDev}$ ($0.8 \leq 1 - \text{Beta} < 0.95$);
 Very High Power to detect Effect = $3 \times \text{StDev}$ ($1 - \text{Beta} \geq 0.99$).

Number of Blocks: 1

Number of Center Points per Block: 0

Randomize Runs

Factor Names and Level Settings:

	Name	Low	High
A:	Pull Back Angle	160	180
B:	Stop Pin	2	3
C:	Pin Height	2	3

Number of Responses: 1

Response Name
 Y1: Distance

OK>>
 Cancel
 Help
 Reset

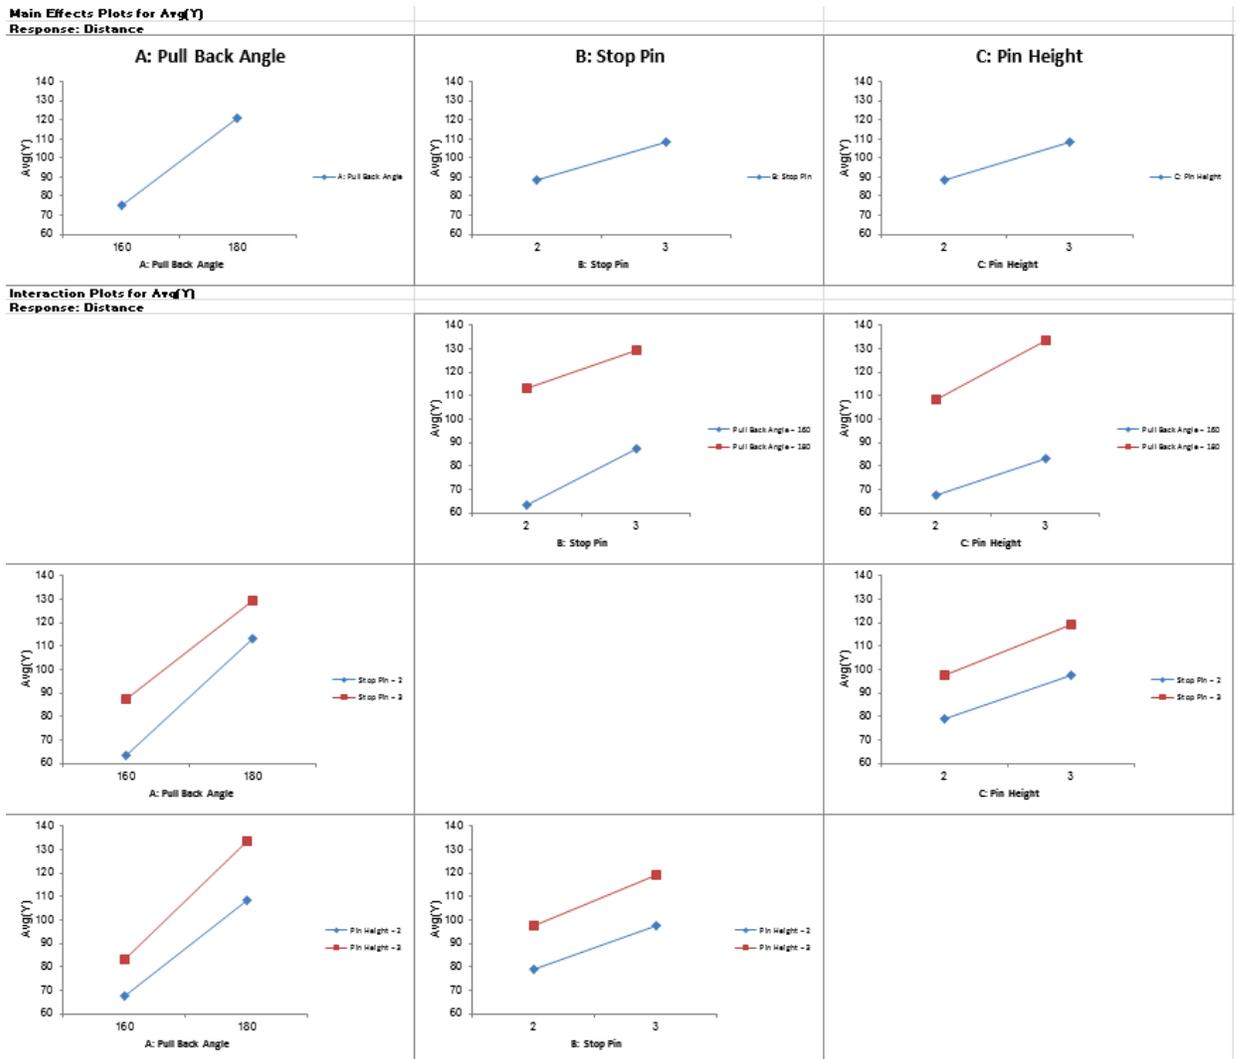
7. Click **OK**. The following worksheet is produced:

Design Type:	3 Factor, 8-Run, 2**3, Full-Factorial
Number of Replicates:	2
Number of Blocks:	1
Number of Center Points per Block:	0
Number of Responses:	1

Run Order	Std. Order	Center Points	Blocks	A: Pull Back Angle	B: Stop Pin	C: Pin Height	Distance
1	14	1	1	180	2	3	
2	4	1	1	180	3	2	
3	12	1	1	180	3	2	
4	9	1	1	160	2	2	
5	1	1	1	160	2	2	
6	5	1	1	160	2	3	
7	2	1	1	180	2	2	
8	15	1	1	160	3	3	
9	10	1	1	180	2	2	
10	7	1	1	160	3	3	
11	11	1	1	160	3	2	
12	13	1	1	160	2	3	
13	16	1	1	180	3	3	
14	6	1	1	180	2	3	
15	8	1	1	180	3	3	
16	3	1	1	160	3	2	

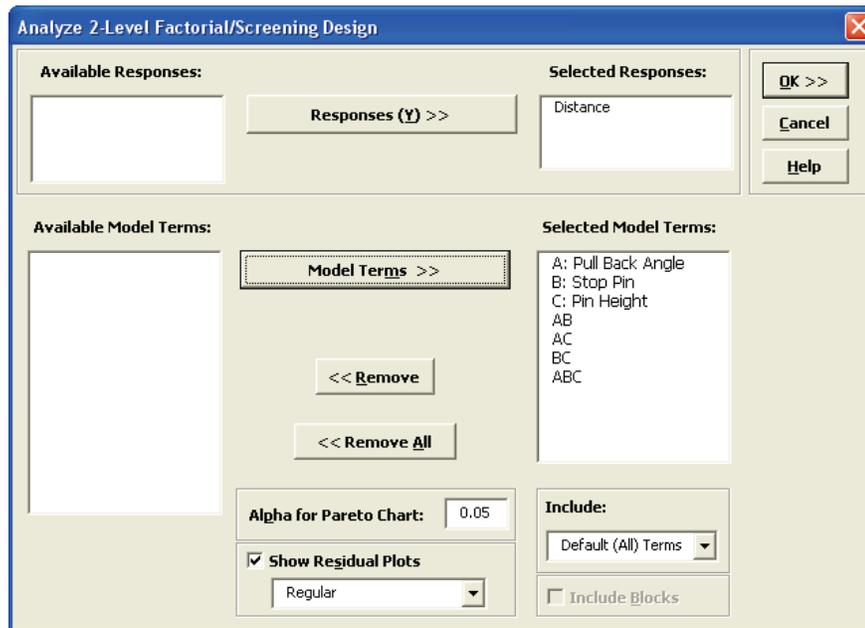
8. You can enter information about the experiment in the fields provided. If you have access to a catapult, perform the experimental runs in the given randomized sequence, and enter the distance values in the Distance column.
9. If you are not able to perform the catapult experiment, open the file **Catapult DOE V6.xlsx**.

10. Before we begin the regression analysis, we will have a quick look at the Main Effects and Interaction Plots. Click **SigmaXL > Design of Experiments > 2 Level Factorial/Screening > Main Effects & Interaction Plots**. The resulting plots are shown below:



Pull Back Angle is the dominant factor having the steepest slope. We can also see that the interaction terms are weak with the almost parallel lines.

11. Click **SigmaXL > Design of Experiments > 2-Level Factorial/Screening > Analyze 2-Level Factorial/Screening Design**.



12. We will use the default analyze settings (all terms in the model) to start. Click **OK**. The resulting Analysis report is shown:

Design of Experiments Analysis

DOE Multiple Regression Model: Distance = (98.3125) + (22.8125) * A: Pull Back Angle + (10.0625) * B: Stop Pin + (10.1875) * C: Pin

Title:	Catapult Example (SigmaXL Version 6)
Date:	
Name of Experimenter:	
Notes:	

Design Type:	3 Factor, 8-Run, 2**3, Full-Factorial
Number of Replicates:	2
Number of Blocks:	1
Number of Center Points per Block:	0
Response:	Distance

Model Summary:

R-Square	99.93%
R-Square Adjusted	99.86%
S (Root Mean Square Error)	1.030776406

Parameter Estimates (Coded Units):

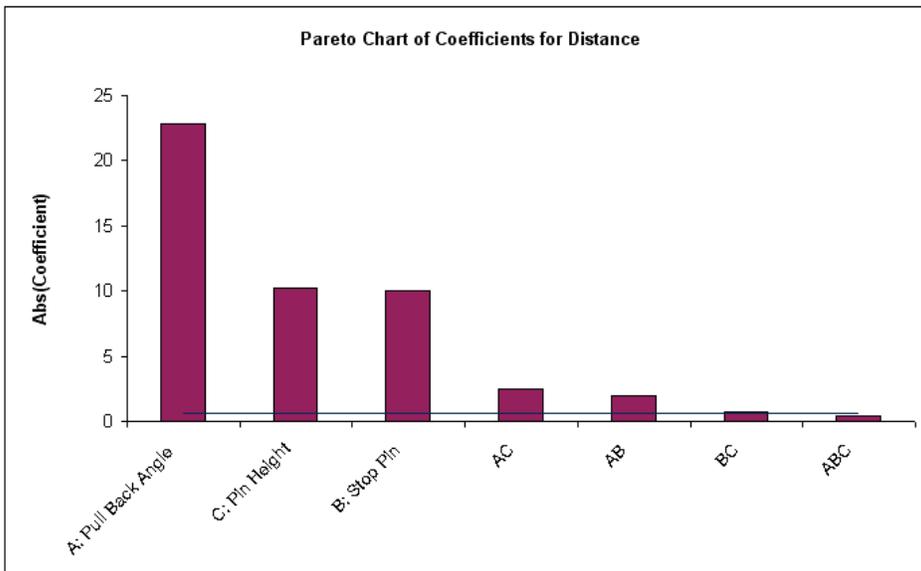
Term	Coefficient	SE Coefficient	T	P	VIF	Tolerance
Constant	98.3125	0.257694102	381.51	0.0000		
A: Pull Back Angle	22.8125	0.257694102	88.526	0.0000	1	1
B: Stop Pin	10.0625	0.257694102	39.048	0.0000	1	1
C: Pin Height	10.1875	0.257694102	39.533	0.0000	1	1
AB	-1.9375	0.257694102	-7.519	0.0001	1	1
AC	2.4375	0.257694102	9.459	0.0000	1	1
BC	0.6875	0.257694102	2.668	0.0285	1	1
ABC	0.4375	0.257694102	1.698	0.1280	1	1

Analysis of Variance for Model:

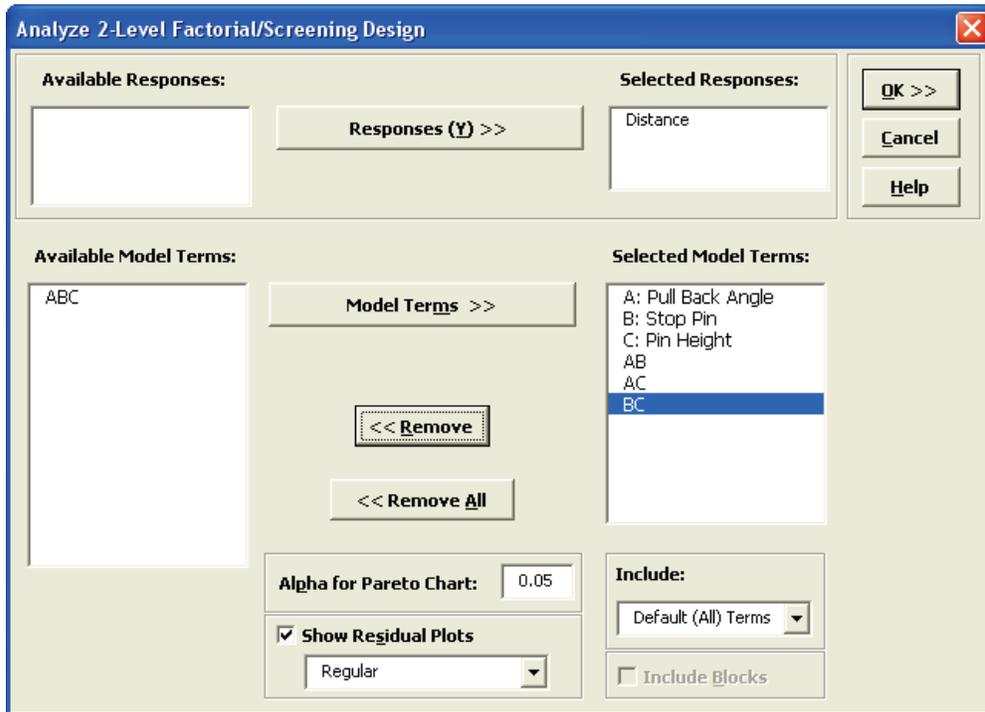
Source	DF	SS	MS	F	P
Model	7	11773	1681.8	1582.9	0.0000
Error	8	8.500	1.0625		
Pure Error	8	8.500	1.0625		
Total (Model + Error)	15	11781	785.43		

Durbin-Watson Test for Autocorrelation in Residuals:

DW Statistic	2.176
P-Value Positive Autocorrelation	0.5491
P-Value Negative Autocorrelation	0.2336



- 13. The model looks very good with an R-Square value of 99.9%! The standard deviation (experimental error) is only 1.03 inches. Clearly *Pull Back Angle* is the most important predictor (X factor), but all the main effects and two-way interaction are significant. However, the three-way interaction is not significant, so it should be removed from the model.
- 14. Click **Recall Last Dialog** (or press **F3**).
- 15. Remove the *ABC* interaction term as shown:



- 16. Click **OK**. The revised report is shown below:

Design of Experiments Analysis

DOE Multiple Regression Model: Distance = (98.3125) + (22.8125) * A: Pull Back Angle + (10.0625) * B: Stop Pin + (10.1875) * C: Pin

Title:	Catapult Example (SigmaXL Version 6)
Date:	
Name of Experimenter:	
Notes:	

Design Type:	3 Factor, 8-Run, 2**3, Full-Factorial
Number of Replicates:	2
Number of Blocks:	1
Number of Center Points per Block:	0
Response:	Distance

Model Summary:

R-Square	99.90%
R-Square Adjusted	99.84%
S (Root Mean Square Error)	1.133

Parameter Estimates (Coded Units):

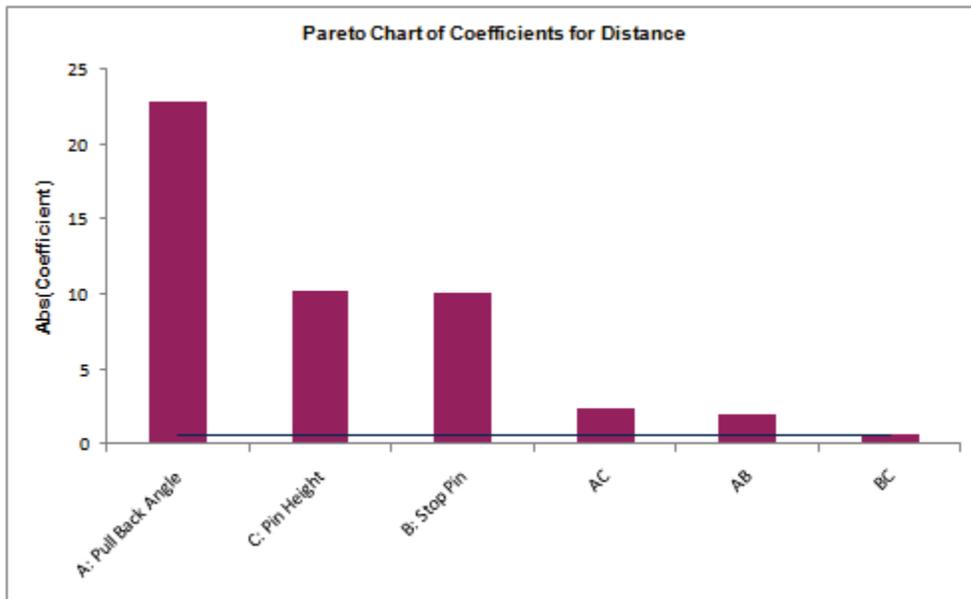
Term	Coefficient	SE Coefficient	T	P	VIF	Tolerance
Constant	98.3125	0.283363969	346.95	0.0000		
A: Pull Back Angle	22.8125	0.283363969	80.506	0.0000	1	1
B: Stop Pin	10.0625	0.283363969	35.511	0.0000	1	1
C: Pin Height	10.1875	0.283363969	35.952	0.0000	1	1
AB	-1.9375	0.283363969	-6.837	0.0001	1	1
AC	2.4375	0.283363969	8.602	0.0000	1	1
BC	0.6875	0.283363969	2.426	0.0382	1	1

Analysis of Variance for Model:

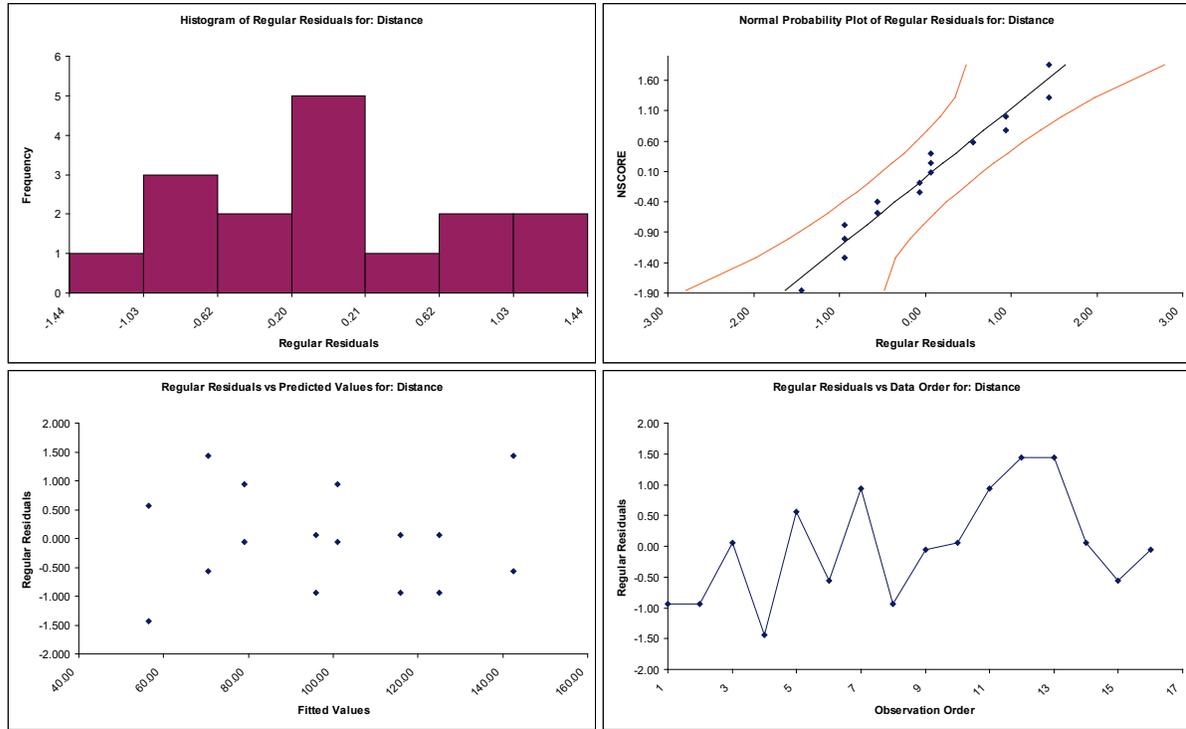
Source	DF	SS	MS	F	P
Model	6	11770	1961.6	1526.9	0.0000
Error	9	11.563	1.285		
Lack of Fit	1	3.063	3.063	2.882	0.1280
Pure Error	8	8.500	1.0625		
Total (Model + Error)	15	11781	785.43		

Durbin-Watson Test for Autocorrelation in Residuals:

DW Statistic	1.609
P-Value Positive Autocorrelation	0.1339
P-Value Negative Autocorrelation	0.7654



17. All the terms in the model are now significant, and there is no evidence of lack of fit (P-Value for lack-of-fit is 0.128 which is $> .05$).
18. The Residual Plots also look very good:



19. Scroll up to the **Predicted Response Calculator**. Enter the predicted values shown. These initial settings were determined by trial and error.

Predicted Response Calculator:

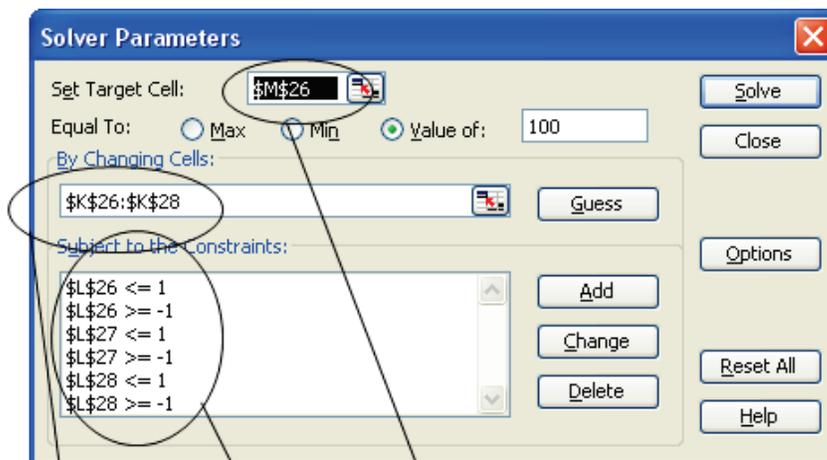
Predictors	Enter Actual Settings:	Coded Settings	Predicted Response	Lower 95% CI	Upper 95% CI	Lower 95% PI	Upper 95% PI
A:	179.5	0.95	99.946875	98.28672343	101.6070266	96.89229107	103.0014589
B:	2	-1					
C:	2	-1					

Note:

1. Enter settings for predictors. If the predictors are block terms, specify a 0 or 1 for each level.
2. Do not insert or delete rows or columns in this worksheet.

Note: CI is the 95% confidence interval for the long term mean and PI is the 95% prediction interval for individual values.

20. Excel's **Solver** may also be used to get a more exact solution:



Predicted Response Calculator:

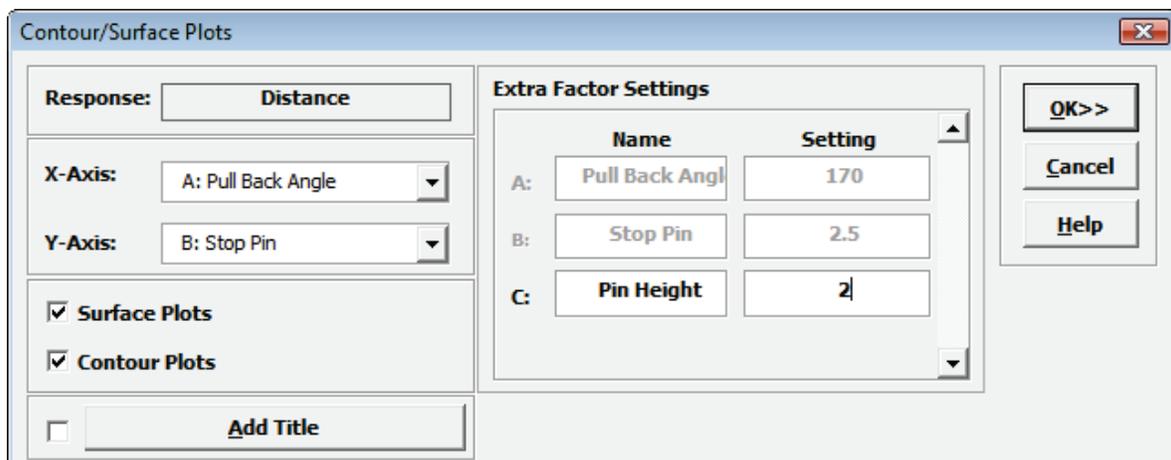
Predictors	Enter Actual Settings:	Coded Settings	Predicted Response	Lower 95% CI	Upper 95% CI	Lower 95% PI	Upper 95% PI
A:	179.52381	0.952380997	100.000001	98.28672343	101.6070266	96.89229107	103.0014589
B:	2	-1					
C:	2	-1					

Note:

1. Enter settings for predictors. If the predictors are block terms, specify a 0 or 1 for each level.
2. Do not insert or delete rows or columns in this worksheet.

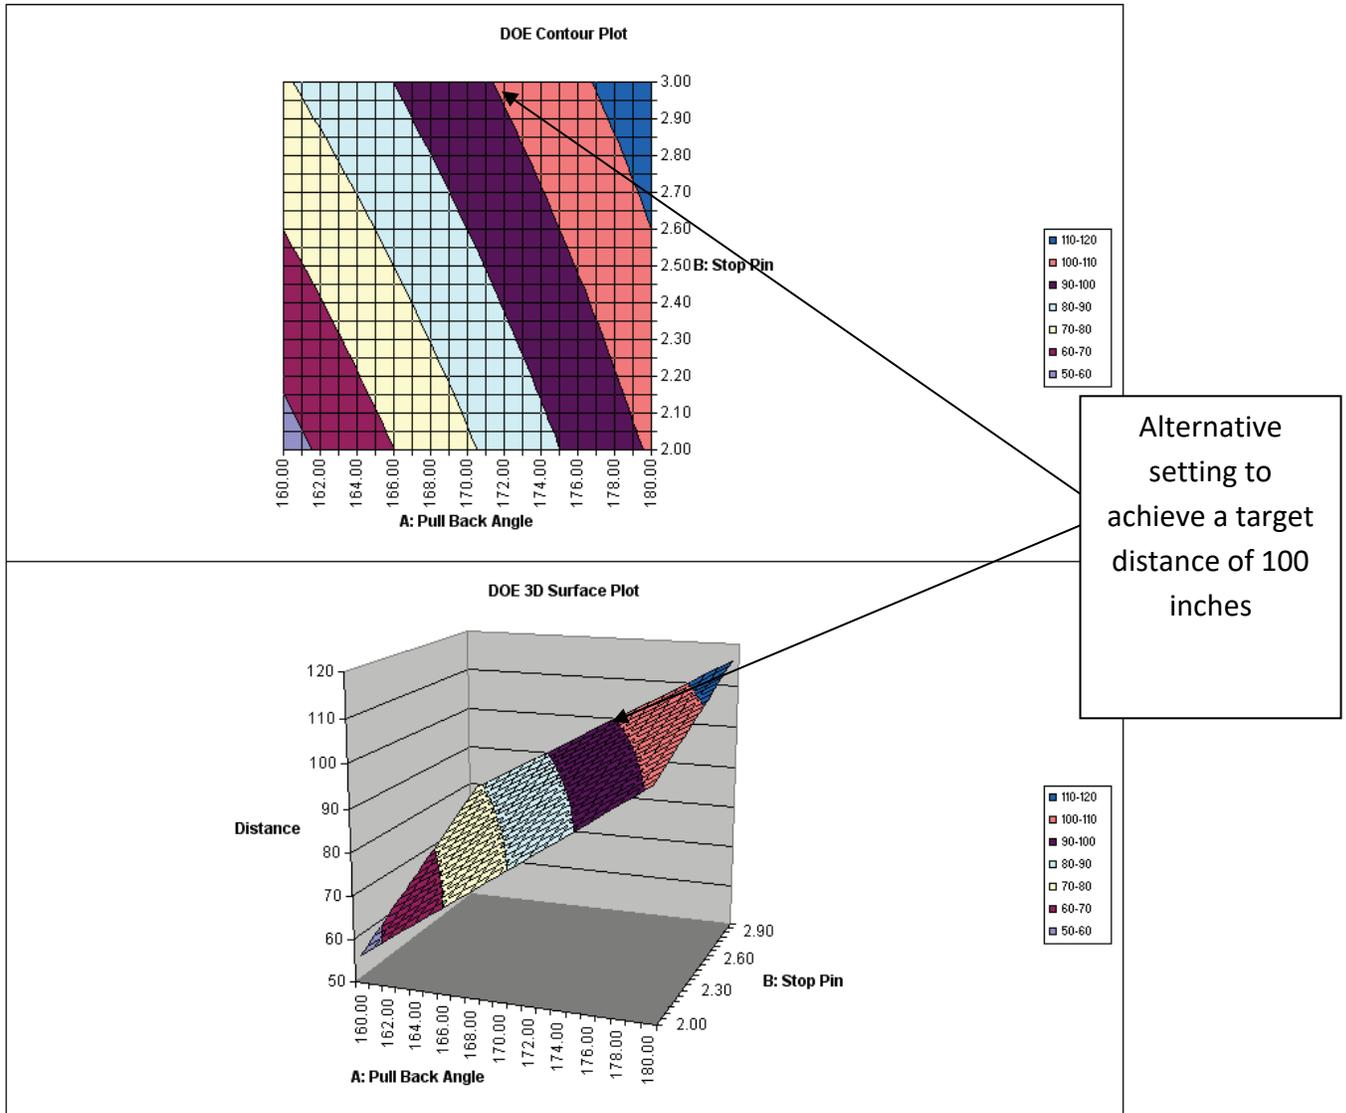
21. The model prediction must then be confirmed with actual experimental runs at the given settings of Pull Back Angle = 179.5, Stop Pin = 2, and Pin Height = 2.

22. Alternative settings to achieve the target distance may be obtained with Contour/Surface Plots. Click **SigmaXL > Design of Experiments > 2-Level Factorial/Screening > Contour/Surface Plots**. Set the **Pin Height** to 2 as shown (after clicking OK, you can use **Recall SigmaXL Dialog** to create another Contour/Surface plot with **Pin Height** set to 3):



23. Click **OK**. The following Contour and Surface Plots are displayed (with Pin Height = 2). Note the contour line with Catapult target distance = 100 inches. Although pin settings are discrete, they appear as continuous, so this will be a constraint in our selection of alternative settings. In addition to Pull Back Angle = 179.5, Stop Pin = 2, Pin Height = 2, we see that Pull Back Angle = 171.4, Stop Pin = 3, Pin Height = 2 is also a valid setting. Alternative setting options are valuable in a designed experiment because they allow you to select lowest cost optimum settings, or settings that are easier to control.

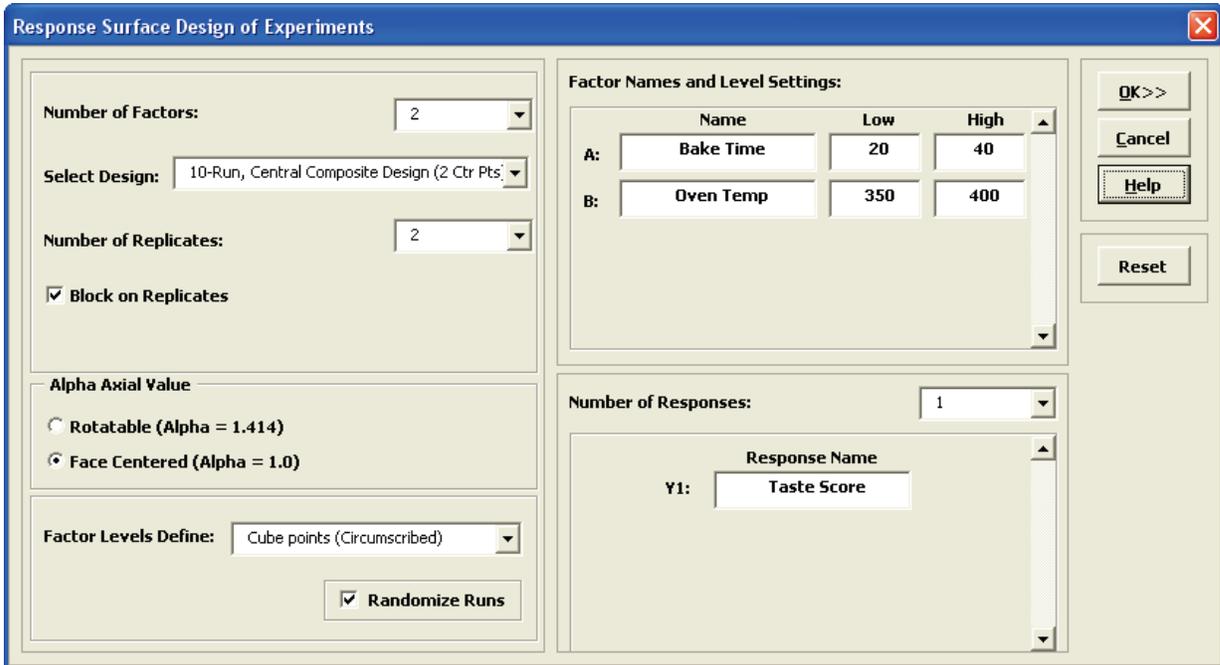
DOE Multiple Regression Model: Distance = (98.3125) + (22.8125) * A: Pull Back Angle + (10.0625) * B: Stop Pin + (10.1875) * C: Pin Height + (-1.9375) * AB + (2.4375) * AC + (0.6875) * BC



Tip: Use the contour/surface in conjunction with the predicted response calculator to determine optimal settings.

Part D – Design and Analysis of Response Surface Experiment – Cake Bake

1. We will illustrate the use of response surface methods using a layer cake baking experiment. The response variable is Taste Score (on a scale of 1-7 where 1 is "awful" and 7 is "delicious"). Average scores for a panel of tasters have been recorded. The X Factors are A: Bake Time (20 to 40 minutes) and B: Oven Temperature (350 to 400 F). The experiment goal is to find the settings that maximize taste score. Other factors such as pan size need to be taken into consideration. In this experiment, these are held constant.
2. Click **SigmaXL > Design of Experiments > Response Surface > Response Surface Designs**.
3. The Number of X Factors can be 2 to 5. Select **Number of Factors = 2**.
4. The available designs are sorted by number of runs. Increasing the number of runs allows for uniform precision or blocking, but we will select the design with fewest runs, the **10-Run, Central Composite Design (2 Ctr Pts)**.
5. As discussed earlier with the catapult experiment, it is a good practice to replicate an experiment if affordable to do so. Here we will select **Number of Replicates = 2** and check **Block on Replicates**. Blocking on replicates allows us to perform the experiment over a two week period, with each block corresponding to week number.
6. Change the **Alpha Axial Value** option to **Face Centered (Alpha = 1.0)**. This simplifies the design to a 3-level design, rather than a 5-level design with $\alpha = 1.414$. (The trade-off is that we lose the desirable statistical property of rotatability for prediction).
Tip: This alpha is the distance from the center point to the end of the axial (star) point. Do not confuse this with **Alpha for Pareto Chart** which is the P-Value used to determine statistical significance. Unfortunately, the term "alpha" has been chosen by statisticians to define two completely different things.
7. Enter **Factor Names and Level Settings** and **Response Name** as shown:



8. Click **OK**. The following worksheet is produced:

Response Surface Methods Design Worksheet

Title:	
Date:	
Name of Experimenter:	
Notes:	

Design Type:	2 Factor, 10-Run, Central Composite Design (2 Ctr Pts)
Alpha Axial Value:	Face Centered (Alpha = 1.0)
Factor Levels Define:	Cube points (Circumscribed)
Number of Replicates:	2
Block on Replicates:	Yes
Number of Responses:	1

Run Order	Std. Order	Center Points	Blocks	A: Bake Time	B: Oven Temp	Taste Score
1	9	0	1	30	375	
2	1	1	1	20	350	
3	7	1	1	30	350	
4	10	0	1	30	375	
5	5	1	1	20	375	
6	3	1	1	20	400	
7	2	1	1	40	350	
8	6	1	1	40	375	
9	4	1	1	40	400	
10	8	1	1	30	400	
11	14	1	2	40	400	
12	20	0	2	30	375	
13	12	1	2	40	350	
14	13	1	2	20	400	
15	16	1	2	40	375	
16	19	0	2	30	375	
17	18	1	2	30	400	
18	15	1	2	20	375	
19	17	1	2	30	350	
20	11	1	2	20	350	

Do not add or delete rows in this worksheet.

9. Open the file **RSM Example – Cake Bake** to obtain response values.

10. Click **SigmaXL > Design of Experiments > Response Surface > Analyze Response Surface Design**.

11. We will use the default analyze settings (all terms in the model, including the block term) to start. Click **OK**. The resulting Analysis report is shown:

SigmaXL: Improve Phase Tools: DOE

Response Surface Methods Analysis

RSM Regression Model: Taste Score = (6.537142857) + (-1.3) * A: Bake Time + (-0.966666667) * B: Oven Temp + (-0.4625) * AB + (-0.8892

Title:	Cake Bake RSM Example
Date:	
Name of Experimenter:	
Notes:	

Design Type:	2 Factor, 10-Run, Central Composite Design (2 Ctr Pts)
Alpha Axial Value:	Face Centered (Alpha = 1.0)
Factor Levels Define:	Cube points (Circumscribed)
Number of Replicates:	2 - Center Points are excluded from the model
Block on Replicates:	Yes
Response	Taste Score

Model Summary:

R-Square	99.00%
R-Square Adjusted	98.54%
S (Root Mean Square Error)	0.186698

Parameter Estimates (Coded Units):

Term	Coefficient	SE Coefficient	T	P	VIF	Tolerance
Constant	6.537142857	0.089258824	73.238	0.0000		
A: Bake Time	-1.3	0.053895135	-24.121	0.0000	1	1
B: Oven Temp	-0.966666667	0.053895135	-17.936	0.0000	1	1
AB	-0.4625	0.06600779	-7.007	0.0000	1	1
AA	-0.889285714	0.086424485	-10.290	0.0000	1.028571429	0.972222222
BB	-1.139285714	0.086424485	-13.182	0.0000	1.028571429	0.972222222
Blocks 2	-0.11	0.083493984	-1.317	0.2104	1	1

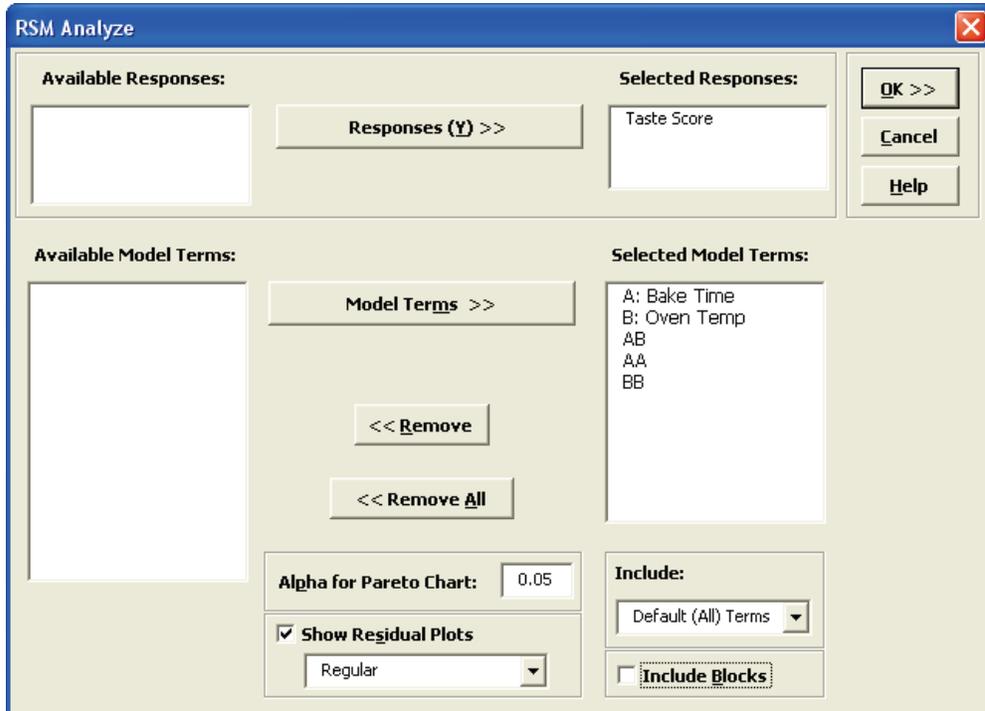
Analysis of Variance for Model:

Source	DF	SS	MS	F	P
Blocks	1	0.0605	0.0605	1.736	0.2104
Model	5	44.852	8.970	257.35	0.0000
Error	13	0.453131	0.03485623		
Lack of Fit	11	0.373131	0.033921	0.848025	0.6564
Pure Error	2	0.08	0.04		
Total (Model + Error)	19	45.366	2.388		

Durbin-Watson Test for Autocorrelation in Residuals:

DW Statistic	2.720
P-Value Positive Autocorrelation	0.9407
P-Value Negative Autocorrelation	0.0315

12. The model looks very good with an R-Square value of 99.0%. The standard deviation (experimental error) is only 0.19 on a 1 to 7 taste scale. All of the model terms are statistically significant ($P < .05$), but the Block term is not, so it should be removed from the model. Note that AA and BB denote the quadratic model terms.
13. Click **Recall Last Dialog** (or press **F3**).
14. Uncheck **Include Blocks** as shown:



15. Click **OK**. The revised report is shown below:

Model Summary:

R-Square	98.87%
R-Square Adjusted	98.46%
S (Root Mean Square Error)	0.191541

Parameter Estimates (Coded Units):

Term	Coefficient	SE Coefficient	T	P	VIF	Tolerance
Constant	6.482142857	0.080940813	80.085	0.0000		
A: Bake Time	-1.3	0.055293102	-23.511	0.0000	1	1
B: Oven Temp	-0.966666667	0.055293102	-17.483	0.0000	1	1
AB	-0.4625	0.067719943	-6.830	0.0000	1	1
AA	-0.889285714	0.088666218	-10.029589	0.0000	1.028571429	0.972222222
BB	-1.139285714	0.088666218	-12.849	0.0000	1.028571429	0.972222222

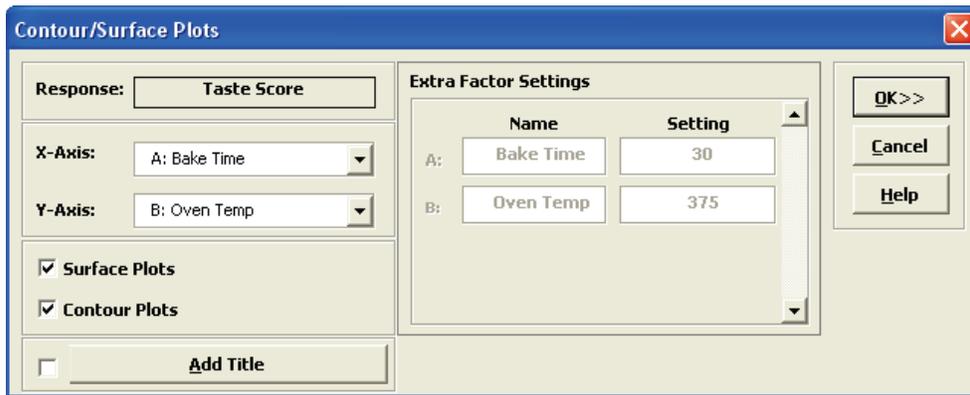
Analysis of Variance for Model:

Source	DF	SS	MS	F	P
Model	5	44.852	8.970	244.50	0.0000
Error	14	0.513631	0.03668793		
Lack of Fit	3	0.248631	0.08287698	3.440	0.0555
Pure Error	11	0.265000	0.02409091		
Total (Model + Error)	19	45.366	2.388		

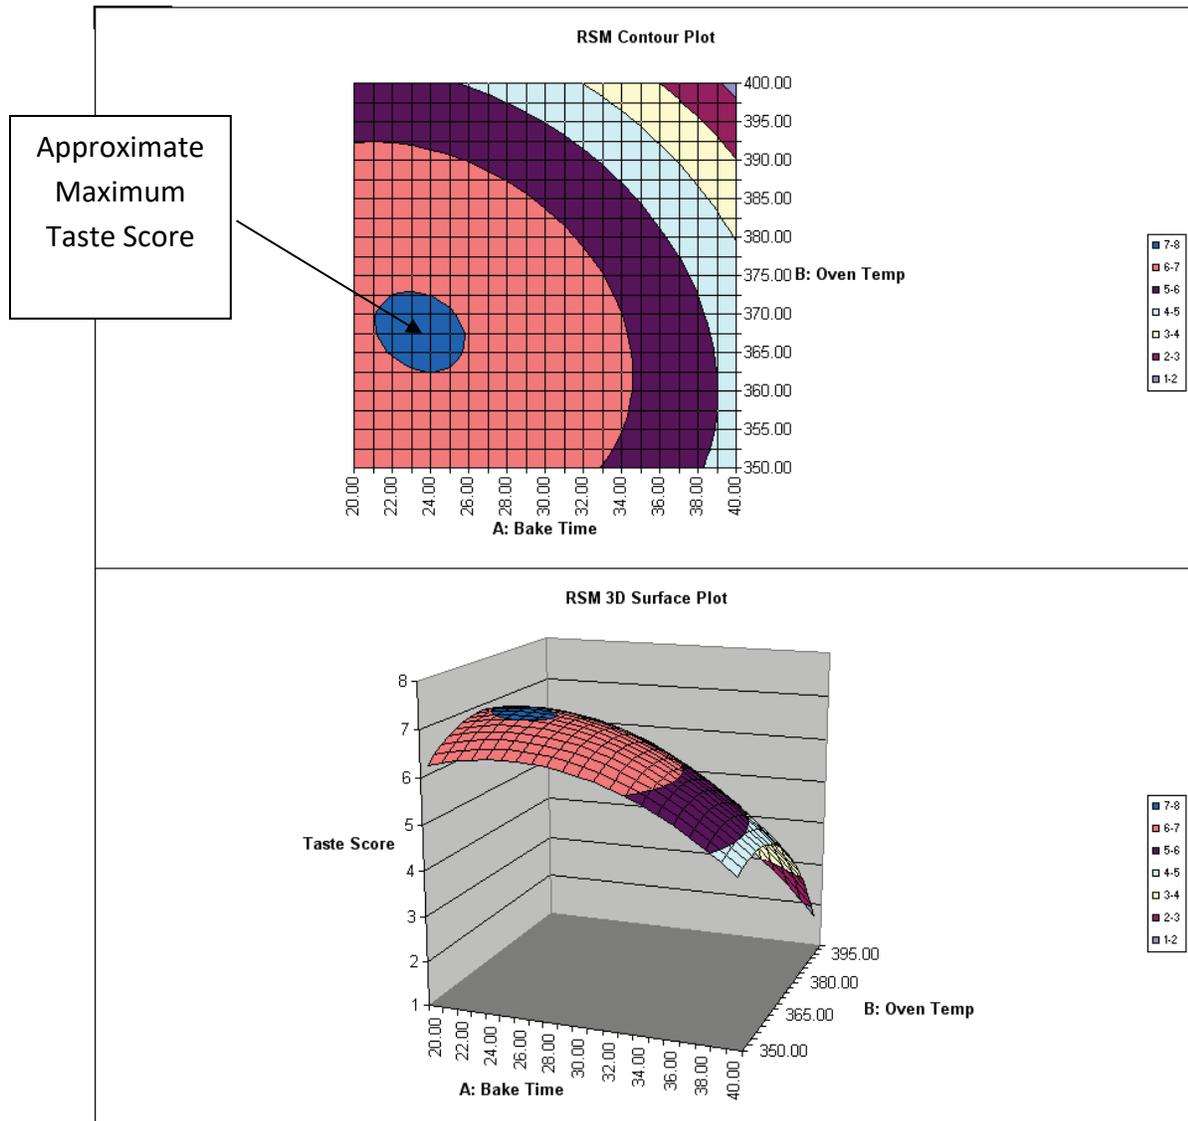
Durbin-Watson Test for Autocorrelation in Residuals:

DW Statistic	2.687
P-Value Positive Autocorrelation	0.9597
P-Value Negative Autocorrelation	0.0798

16. To create a contour and surface plot, click **SigmaXL > Design of Experiments > Response Surface > Contour/Surface Plots**.



17. Click **OK**. The following Contour and Surface Plots are displayed.

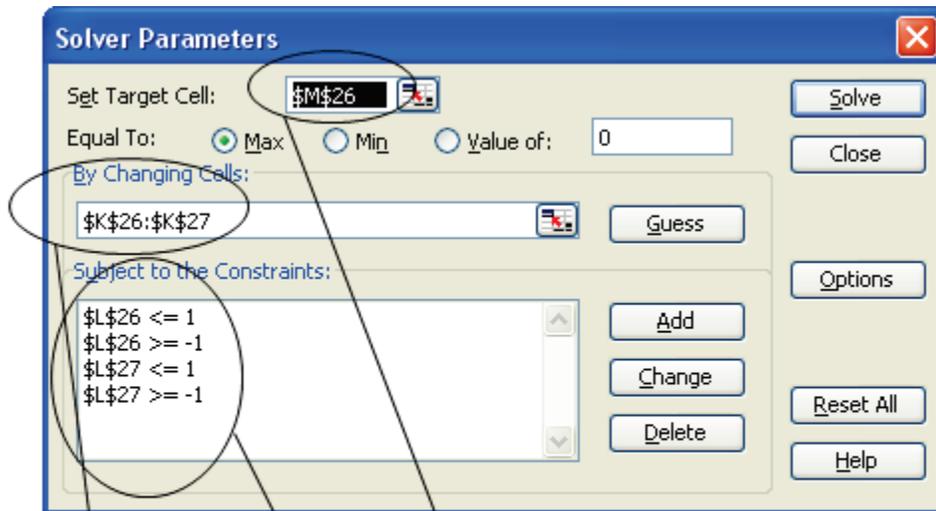


18. Go to **Analyze – 2 Factor RSM** tab. Scroll across to the **Predicted Response Calculator**. Enter the predicted values shown. These initial settings were determined from the contour plot as an estimate to yield a maximum taste score.

Predicted Response Calculator:

Predictors	Enter Actual Settings:	Coded Settings	Predicted Response	Lower 95% CI	Upper 95% CI	Lower 95% PI	Upper 95% PI
A:	23	-0.7	7.047089524	6.883297028	7.21088202	6.604826626	7.489352422
B:	368	-0.28					

19. Excel's **Solver** may also be used to get a more exact solution:



Predicted Response Calculator:

Predictors	Enter Actual Settings:	Coded Settings	Predicted Response	Lower 95% CI	Upper 95% CI	Lower 95% PI	Upper 95% PI
A:	23.44815482	-0.655184518	7.048786514	6.883297028	7.21088202	6.604826626	7.489352422
B:	367.7184597	-0.29126161					

20. Although the model is predicting values that exceed the maximum taste score of 7, this is expected to give us the best possible settings for cook time and bake temperature. To quote the eminent statistician George Box, "All models are wrong, some are useful". Additional experimental runs carried out at Time = 23.4 minutes and Temperature = 367.7 F. confirm that these are ideal settings.

Part E – Basic Taguchi DOE Templates

Introduction – Taguchi Methods

Dr. Genichi Taguchi was a Japanese engineer and quality consultant who codified his consulting techniques into a formal methodology. This generic engineering tool, called the Taguchi method, is used for system studies. The purpose of these studies is to reduce system variability while simultaneously decreasing costs and increasing productivity.

The features of the Taguchi method include:

- extensive use of experimental design
- separation of factors by role (control and noise)
- use of measures of variability as responses
- dual objectives of process centering and noise minimization
- use of loss functions for economic justification of each application

The following steps outline the Taguchi Design of Experiments method:

1. Identify what are the control factors and what are the noise factors. Noise factors are costly and difficult to control, while an ideal control factor is easy to control precisely. Noise factors can include environmental factors, component deterioration, and process variation.
2. Identify the levels for the control and noise factors.
3. Identify the responses of interest and determine a quality criterion for each response. Does the system require that the response match a specified value, be as large as possible, or approach zero?
4. Construct a design for the control factors (inner array) and a design for the noise factors (outer array). The inner-array is selected based on the number of control factors. Typically, interactions are assumed to be negligible, but if they are to be included, then a larger design may be required. The outer array simulates the various conditions that the noise factors would produce in reality. For each run in the inner array, all runs in the outer array are carried out.
5. Conduct the design of experiments. Typically, you complete all the runs in the outer array before proceeding to the next run of the inner array. The runs for the inner array and outer array are (ideally) randomized separately.
6. Evaluate the performance statistic for each run of the inner array. These measures become the responses for the inner array. The performance statistics include Mean(Y), StDev(Y) or Ln(StDev), and Taguchi's signal-to-noise ratios. Some practitioners prefer to use Ln(StDev) as a variance stabilizing transformation. Taguchi signal-to-noise ratios include Nominal is Best, Larger is Better and Smaller is Better. Signal-to-noise ratios are always maximized.

7. From the Pareto of Deltas (Main Effects), Pareto of ANOVA SS (Sum-of-Squares) % Contribution, Main Effects & Interaction Plots and Prediction Equation model, determine the new set points for the control factors. Taguchi's two step optimization first finds control settings that maximize the SN Ratio (and/or minimize the StDev), then if available, factors that move the mean to Target without affecting the SN Ratio.
8. Confirm that the new settings meet target, dispersion and loss criterion with follow-up experimental runs.

Taguchi Orthogonal Arrays are a cookbook of designs that are similar to Full & Fractional-Factorial and Plackett-Burman designs. For example, the L4 design is 2-level, 4 runs with up to 3 factors; L8 is 2-level, 8 runs with up to 7 factors; L9 is 3-level, 9 runs with up to 4 factors. Taguchi designs use 2-level coded values of 1, 2 instead of the orthogonal coding -1, +1 and 3-level coded values of 1, 2, 3 instead of -1, 0, +1.

Interactions are typically assumed to be negligible compared to main effects, but some designs permit the analysis of all interactions or aliased interactions. Selection of aliased interactions is more economical than all interactions, but they should be used with caution. Process knowledge, engineering or theory are used to make the selection and assume that the chosen interaction is dominant and the others are negligible. Aliased interactions are often associated with the largest main effects. Confirmation runs should always be used to validate the model.

For further reading, see:

Basic:

Fowlkes, W.Y.; Creveling, C.M. (2006) *Engineering Methods for Robust Product Design: Using Taguchi Methods in Technology and Product Development*, Prentice Hall.

Ross, P.J. (1996) *Taguchi Techniques for Quality Engineering*, 2nd Edition, McGraw-Hill, New York, NY.

Roy, R.K. (2010) *A Primer on the Taguchi Method*, 2nd Edition, Society of Manufacturing Engineers, Dearborn, MI.

Advanced:

Taguchi, G.; Chowdhury, S.; Wu, Y. (2005) *Taguchi's Quality Engineering Handbook*, John Wiley, Hoboken, NJ.

Overview of Taguchi DOE Templates

The Taguchi DOE templates are similar to the other SigmaXL templates: simply enter the inputs and resulting outputs are produced immediately.

Click **SigmaXL > Design of Experiments > Basic Taguchi DOE Templates** or **SigmaXL > Templates and Calculators > Basic Taguchi DOE Templates** to access these templates:

- Taguchi L4 (2 Level)
 - Two-Factor (with Two-Way Interaction)
 - Three-Factor

- Taguchi L8 (2 Level)
 - Three-Factor (with Two-Way Interactions)
 - Four to Six-Factor (with Aliased Two-Way Interactions)
 - Seven-Factor

- Taguchi L9 (3 Level)
 - Two-Factor (with Two-Way Interaction)
 - Four-Factor

- Taguchi L12 (2 Level): Eleven Factor

- Taguchi L16 (2 Level)
 - Five-Factor (with Two-Way Interactions)
 - Eight to Fourteen-Factor (with Aliased Two-Way Interactions)
 - Fifteen-Factor

- Taguchi L18 (2/3 Level)
 - Three-Factor (with Two-Way Interactions)
 - Eight-Factor (with A*B Interaction)

- Taguchi L27 (3 Level)
 - Three-Factor (with Two-Way Interactions)
 - Thirteen-Factor

Note: SigmaXL Basic Taguchi DOE templates do not include L32 (2 Level), L36 (2/3 Level), L54 (2/3 Level), designs with more than 3 Levels, or a signal factor for dynamic characteristics.

Template Features:

- Levels are discrete categorical so may be numeric or text
- Fill in the blanks template, charts automatically update
- Predicted Response Calculator and Charts for Mean, Standard Deviation (or Ln Standard Deviation) and Signal-to-Noise Ratio
- Available Signal-to-Noise Ratios:
 - Nominal is Best
 - Nominal is Best (Variance Only)
 - Nominal is Best (Mean Square Deviation with Target)
 - Larger is Better
 - Smaller is Better

Signal-to-Noise Ratios	Formula
SN: Nominal is Best	$10 \cdot \log_{10}(\bar{Y}^2/S^2)$
SN: Nominal is Best (Variance Only)	$-10 \cdot \log_{10}(S^2)$
SN: Nominal is Best (MSD with Target)	$-10 \cdot \log_{10}(\text{Sum}((Y-T)^2)/n)$
SN: Larger is Better	$-10 \cdot \log_{10}(\text{Sum}(1/Y^2)/n)$
SN: Smaller is Better	$-10 \cdot \log_{10}(\text{Sum}(Y^2)/n)$

- Up to 27 Replications for Outer Array (i.e., support up to L27 Outer Array)
- Pareto of Deltas (Effects) and ANOVA SS (Sum-of-Squares) % Contribution (for Main Effects and Two-Way Interactions)
- Main Effects Plot and Interaction Plots (if applicable)
- For designs with aliased interactions a drop-down list of available aliased interactions is provided. This is much easier to use than Linear Graphs.
- Column assignments to Orthogonal Array are optimized to ensure maximum design resolution.

Template Notes:

1. Select desired Signal-to-Noise Ratio to maximize. The SN formula is displayed.
2. For Nominal is Best, "SN: Nominal is Best" is recommended for non-negative data. Use "SN: Nominal is Best (Variance Only)" if the data has a mixture of negative and positive values.
3. If selection is "SN: Nominal is Best MSD with Target", enter Target value (MSD denotes Mean Square Deviation); if Target = 0, this is equivalent to Smaller Is Better. For use of MSD in SN Ratio, see Ranjit Roy (2010) "A Primer on the Taguchi Method," Second Edition, Society of Manufacturing Engineers.
4. Larger is Better requires positive data. Smaller is Better requires non-negative data (target is zero).
5. Select desired Standard Deviation Response: StDev(Y) or Ln(StDev(Y)).
6. Enter Factor Names and Factor Levels. Levels are discrete categorical, so may be numeric or text.
7. If applicable, select Aliased Interactions. Enter Run Number (if runs have been randomized) and Outer Array Response values.
8. Selection of Aliased Interaction assumes that the chosen interaction is dominant and the others are negligible. It is often associated with the largest main effects. Confirmation runs should be used to validate model. Please refresh selection if message appears.
9. For Taguchi L8 Five to Six Factors and L16 Nine to Fourteen Factors, not all possible interactions are available in the drop-down list (they are aliased with main effects). If an interaction of interest is not available in any of the drop-down lists, please modify the Factor Names so that they are assigned to columns shown in the list and then select that interaction.
10. The typical default display of Yrep1 to Yrep9 accommodates an outer array up to L9. Unhide columns to display Yrep10 to Yrep27, which permits an outer array up to L27.
11. Delta is the difference between the maximum and minimum average response for the levels of each factor in the inner array. Two-way interactions are not included. ANOVA SS (Sum-of-Squares) % Contribution includes two-way interactions.
12. To compute a predicted response, enter Actual Factor Setting using the drop-down selection. Factors are treated as categorical. If Factor Setting Coded is "FALSE", please refresh selection to match level settings. Note, Excel Solver cannot be used with this calculator.
13. The Interaction Plot X-axis/Legend factors may be switched by clicking on the chart, then select Excel Design > Switch Row/Column.
14. DOE Templates are protected worksheets by default, but this may be modified by clicking **SigmaXL > Help > Unprotect Worksheet**. Alternatively, you can click Excel **File > Info > Unprotect** or **Home > Format > Unprotect Sheet**.

15. Chart Y-axis scaling is automatic. To modify, double-click on the Y axis and adjust Minimum and Maximum.
16. The Orthogonal Array or Dummy Coding may be analyzed using Multiple Regression Analysis, adding and removing terms from the model as necessary. Terms that are removed will be used to estimate error for p-values. In particular, regression analysis of 3 Level designs with interactions (L9 Two Factor, L18 Three Factor and L27 Three Factor) should use Dummy Coding.
17. Dummy Coding uses Level 1 as the reference value, so Level 1 does not appear in the array. The Predicted Output calculator formula coefficients and ANOVA Sum-of-Squares (SS) % Contribution are computed using Dummy Coding regression. This is consistent with SigmaXL's use of Dummy Coding regression for categorical factors in other tools.
18. Do not modify any other part of the template.
19. Caution: The use of Cell Autocomplete with Flash Fill in Excel 2016 or higher may result in a crash (this is a known Microsoft issue). If this occurs, please turn off Automatic Flash Fill (File > Options > Advanced, uncheck "Automatically Flash Fill" under "Enable Autocomplete for cell value").
20. If there are further stability issues, please turn off Automatic Recalculate (Formulas > Calculation Options, Select Manual). Use "Calculate Now" to refresh calculations after entry of Factor Names, Levels, Outer Array data or as needed.

Taguchi L8 (2 Level) Three Factor – Robust Cake Example

This is the same Robust Cake Experiment adapted from the Video *Designing Industrial Experiments*, by Box, Bisgaard and Fung, analyzed earlier using a Three Factor Full-Factorial. Here we will use the Taguchi L8 Three Factor.

1. Click **SigmaXL > Help > Template Examples > Taguchi Examples > L8 Three Factor – Robust Cake** to open the example file. If you wish to start with a blank template and populate the values, click **SigmaXL > Design of Experiments > Basic Taguchi DOE Templates > Taguchi L8 (2 Level) > Three Factor (with Two-Way Interactions)**.
2. The response is Taste Score (on a scale of 1-7 where 1 is "awful" and 7 is "delicious").
3. The five **Outer Array** values (**Yrep1** to **Yrep5**) have different Cooking Time and Temperature Conditions. Note, the outer array was a two-factor, full-factorial plus center point, hence 5 replications, so was not a typical Taguchi outer array.
4. The goal is to maximize the Taste Score (with minimum variation), so we will maximize the Signal-to-Noise Ratio **SN:Larger is Better**, selected from the list as shown:

Signal-to-Noise Ratio:	SN: Larger is Better
Standard Deviation Response:	SN: Nominal is Best SN: Nominal is Best (Variance Only) SN: Nominal is Best (MSD with Target)
Factor	SN: Larger is Better SN: Smaller is Better

Signal-to-Noise Ratio:	SN: Larger is Better
Standard Deviation Response:	StDev (Y)

SN Formula: $-10 \cdot \log_{10}(\text{Sum}(1/Y^2)/n)$

5. The **Inner Array** controllable factors are Flour, Butter and Egg. Actual low and high settings are not given in the video, so we will use coded 1 and 2 values for **Level 1** and **Level 2**. Note that the Taguchi L8 standard sort order (left to right) is different from Full-Factorial (right to left), so we will enter **Factor Name** as Egg, Butter, Flour to adjust for this difference. (In the Full-Factorial it was Flour, Butter, Egg).

Factor	Factor Name	Level 1	Level 2
A	Egg	1	2
B	Butter	1	2
C	Flour	1	2

6. We are looking for a combination of Egg, Butter, and Flour that will not only taste good, but consistently taste good over a wide range of Cooking Time and Temperature conditions.

7. The runs were randomized and order recorded in the column **Run No.** The **Outer Array** Taste Score values are entered as shown:

Inner Array:

Std. Order	Run No.	Egg	Butter	Flour
1	6	1	1	1
2	2	1	1	2
3	7	1	2	1
4	1	1	2	2
5	5	2	1	1
6	8	2	1	2
7	3	2	2	1
8	4	2	2	2

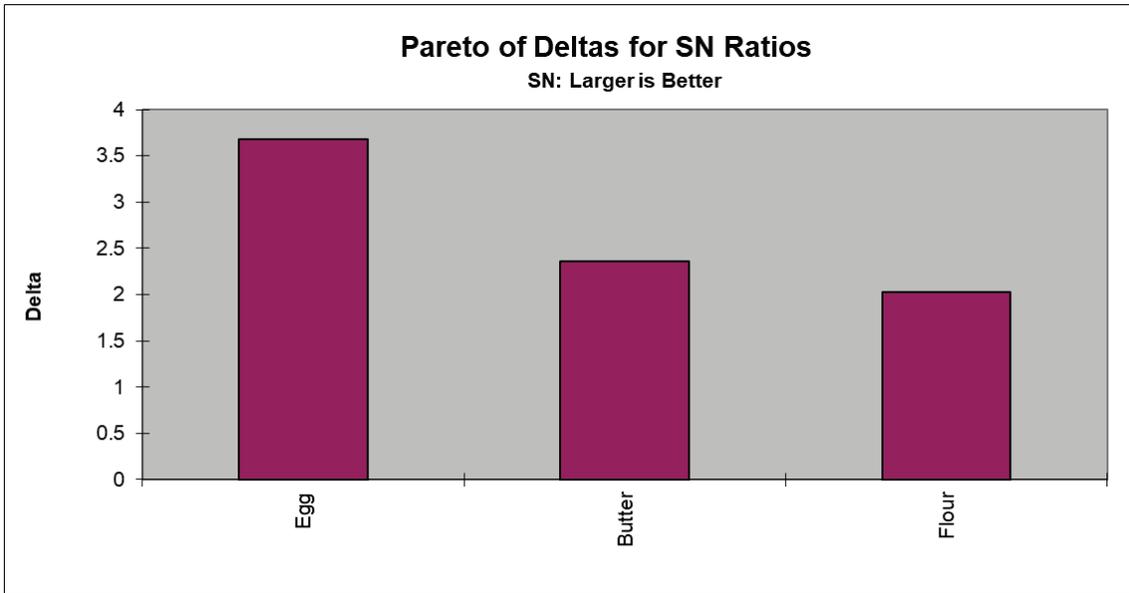
Outer Array:

Yrep1	Yrep2	Yrep3	Yrep4	Yrep5
3.1	1.1	5.7	6.4	1.3
3.2	3.8	4.9	4.3	1.3
5.3	3.7	5.1	6.7	2.9
4.1	4.5	6.4	5.8	5.2
5.9	4.2	6.8	6.5	3.5
6.9	5	6	5.9	5.7
3	3.1	6.3	6.4	3
4.5	3.9	5.5	5	5.4

Scroll to the right to view the calculated statistics:

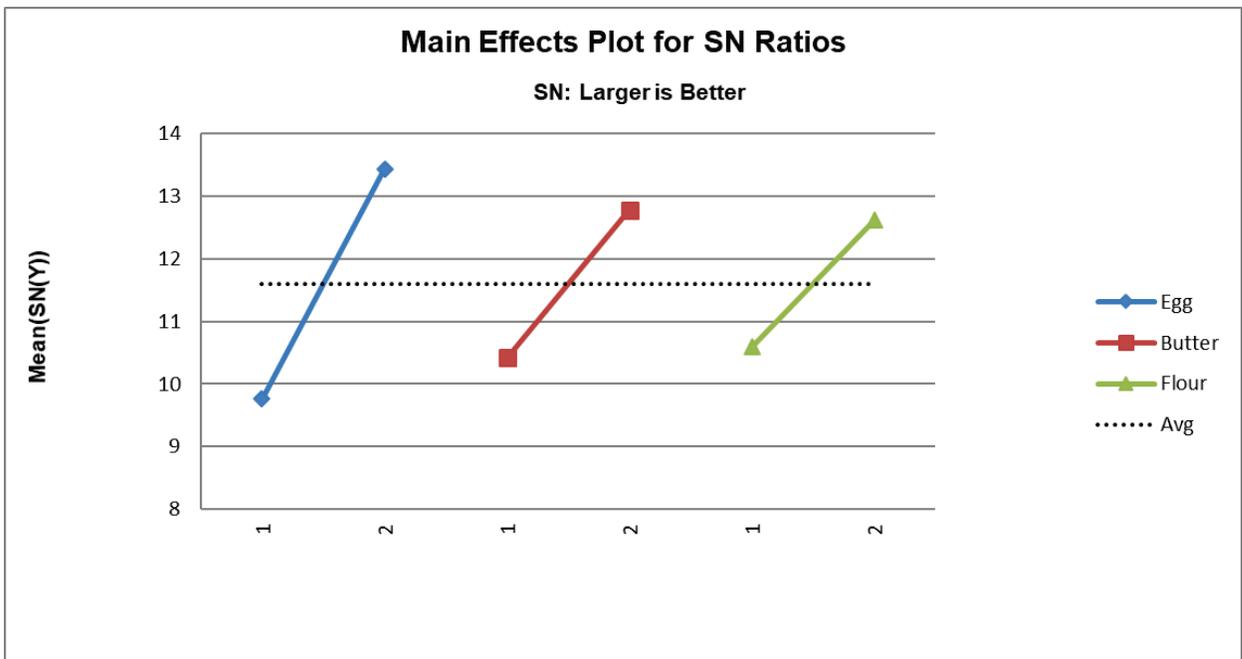
Mean (Y)	StDev (Y)	SN: Larger is Better
3.52	2.449897957	5.010245078
3.5	1.380217374	7.673306706
4.74	1.479188967	12.39162469
5.2	0.935414347	13.97953162
5.38	1.454991409	13.71929378
5.9	0.681909085	15.27864445
4.36	1.817415748	11.23909439
4.86	0.665582452	13.51678365

- We will focus our analysis on the Signal-To-Noise Ratio and review the Mean and StDev afterward as a check on the results. Scroll right and down to view the **Pareto of Deltas for SN Ratios**:

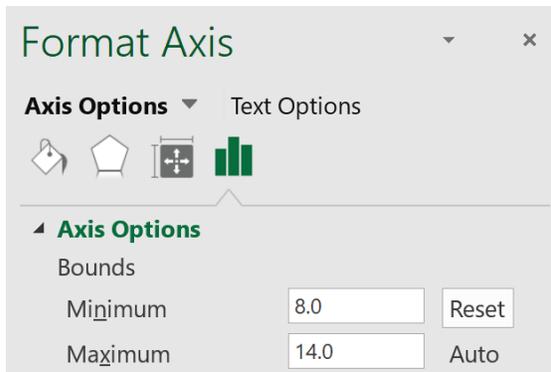


Delta is the difference between the maximum and minimum average response for the levels of each factor in the inner array. It appears that Egg is the dominant factor followed by Butter.

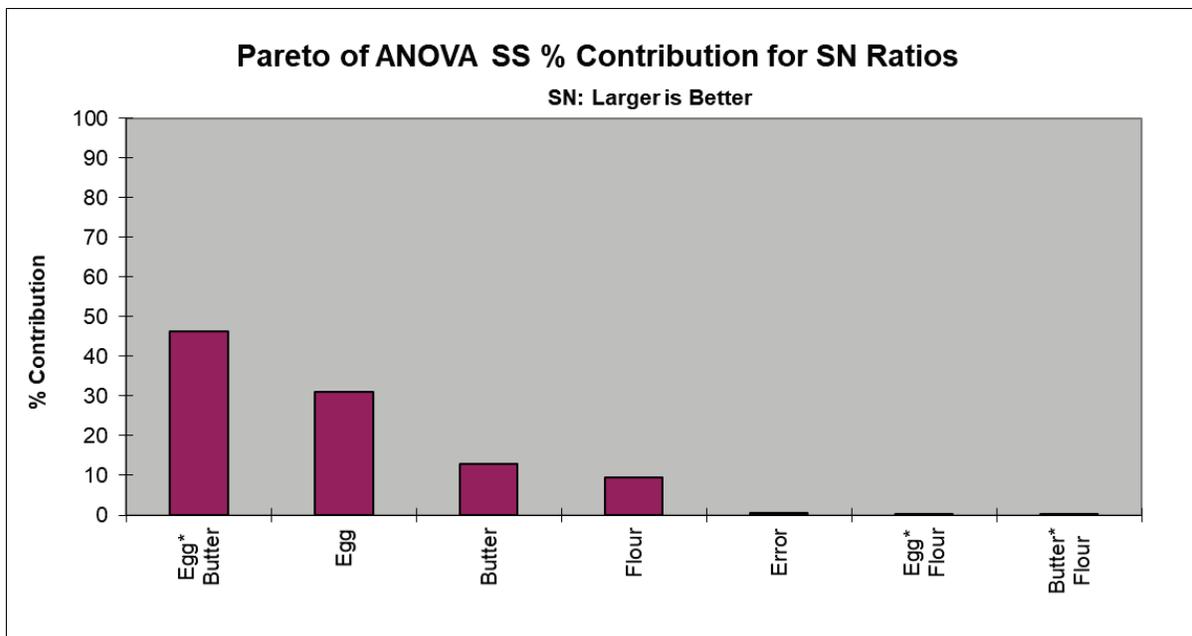
- Scroll down to look at the **Main Effects Plot for SN Ratios**:



It initially appears that the maximum Taste Score occurs with Egg = 2, Butter = 2 and Flour = 2. Note: The Y Axis has been adjusted to get a better scale (double click on the Y-Axis to open Excel Format Axis dialog, click **Axis Options** and set **Minimum** = 8 as shown).



10. However, this design permits the analysis of interactions, so we will now look at the **Pareto of ANOVA SS (Sum-of-Squares) % Contribution for SN Ratios** which includes Main Effects and Interactions:

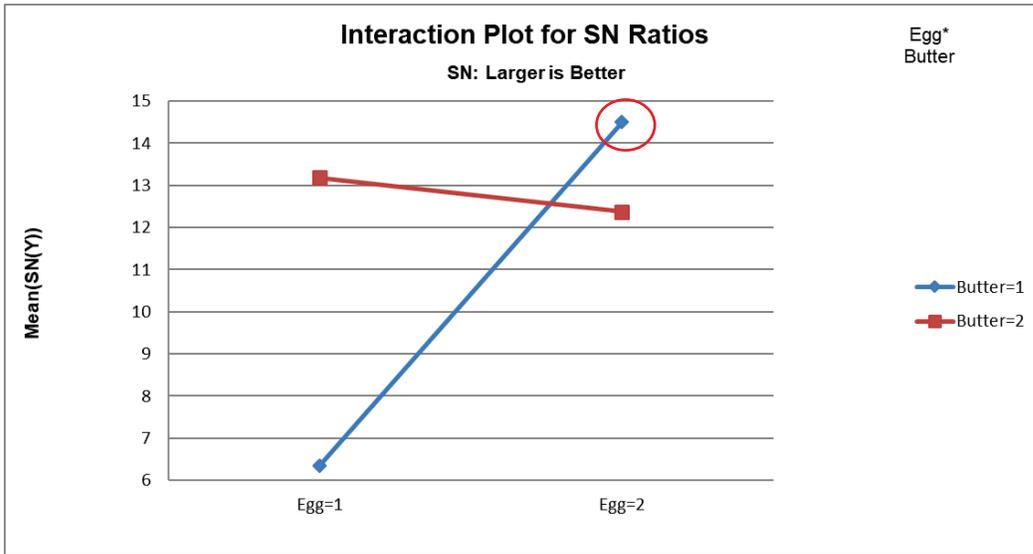


11. Now we see that the Egg*Butter interaction is clearly the dominant factor (in agreement with the analysis done earlier using the Three Factor Full-Factorial).

Note 1: The "Error" term in this model is the three-way interaction which is assumed to be negligible, so is not included in the model.

Note 2: SigmaXL does not compute a significance line in the Taguchi Templates, as done in the Three Factor Full-Factorial Template. Taguchi recommends pooling of effects to estimate experimental error which will be demonstrated later.

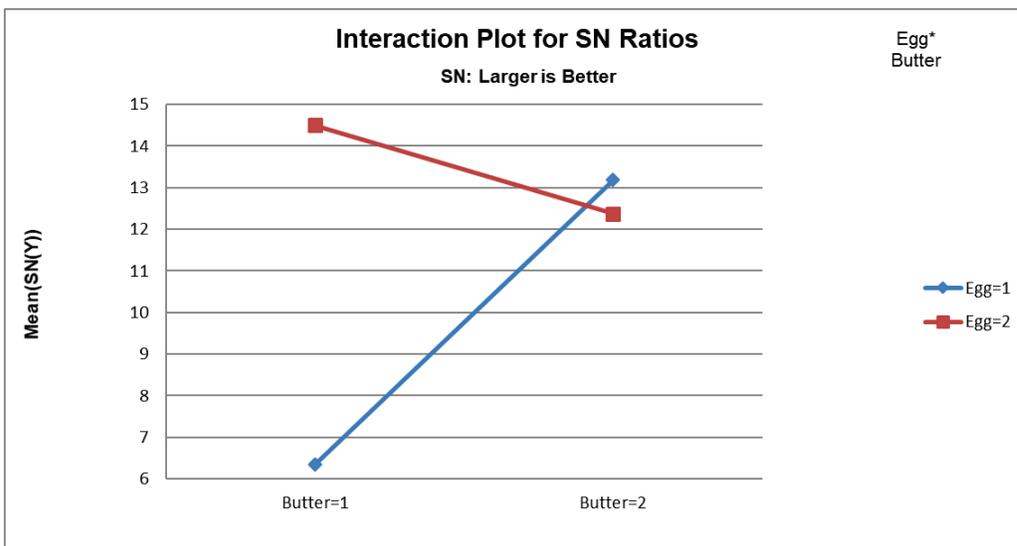
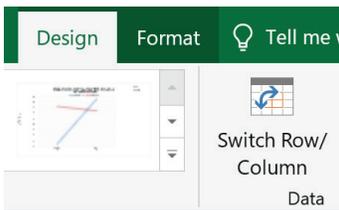
12. Scroll down to view the **Interaction Plot for SN Ratios** for the Egg*Butter Interaction:



Now we can clearly see the large effect due to the two-way interaction. The effect that Egg has on Taste depends on what the Butter setting is. The maximum SN Ratio for Taste Score is achieved with Egg = 2 and Butter = 1.

Note 1: The Y Axis has been adjusted to improve visibility (**Minimum = 6**).

Note 2: The Interaction Plot X-axis/Legend factors may be switched by clicking on the chart, then selecting Excel **Design > Switch Row/Column** as shown:



13. From the above plots, we conclude that the optimum setting to produce a maximum Taste Score SN Ratio is: Egg = 2, Butter = 1 and Flour = 2. (If we had only considered the Main Effects plot, we would have incorrectly set Butter = 2).
14. Scroll up and left to enter these values in the Predicted Output Calculator. Factor settings can be manually entered or selected from the drop-down list as shown:

Predicted Output for Y:

Factor Name	Enter Actual Factor Setting - Uncoded	Factor Setting Coded
Egg	2	2
Butter	1	1
Flour	2	2

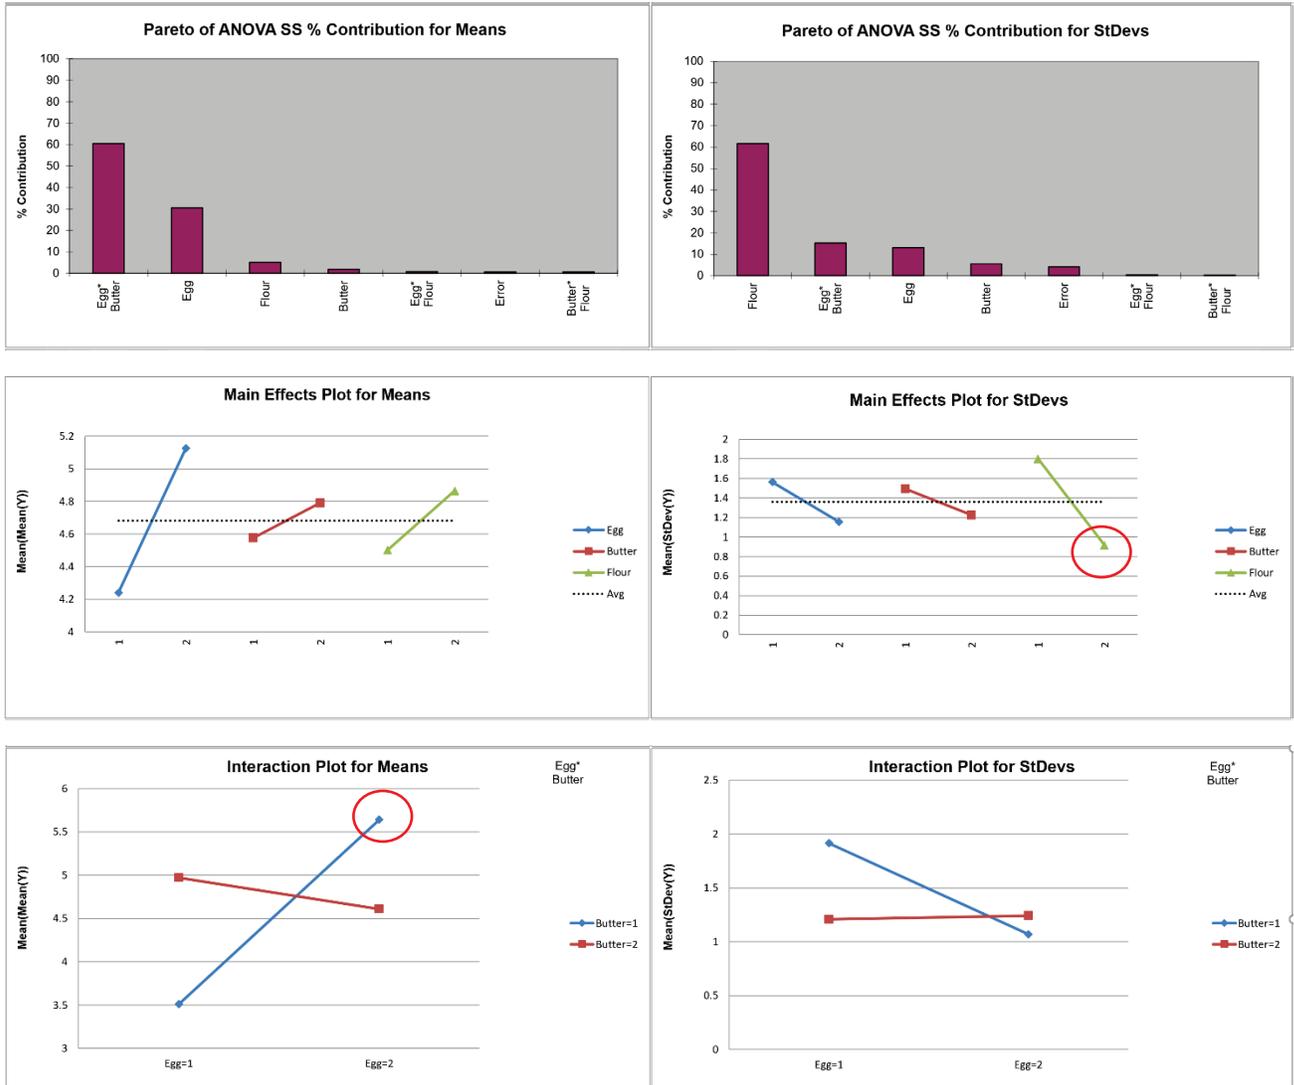
Predicted Output for Y:

Factor Name	Enter Actual Factor Setting - Uncoded	Factor Setting Coded	\hat{Y}	\hat{S}	SN: Larger is Better
Egg	2	2	5.8375	0.568826968	15.50283111
Butter	1	1			
Flour	2	2			

This gives the predicted Mean, Standard Deviation and SN Ratio for Taste Score and agrees with the results obtained in Part B, the analysis of the Three-Factor Full Factorial. (The predicted values are slightly different because the Full Factorial includes the three-way interaction term in the model, whereas this is considered as an error term in the Taguchi Template).

Note: This calculator treats the factors as categorical, so levels can be numeric or text. The drop-down list is a convenient way to select the factor levels, but unfortunately Excel Solver cannot be used with this calculator to perform optimization. We will demonstrate the use of Solver with Multiple Regression later.

15. As a validation of the SN Ratio analysis, we will review the Plots for Mean and Standard Deviation:



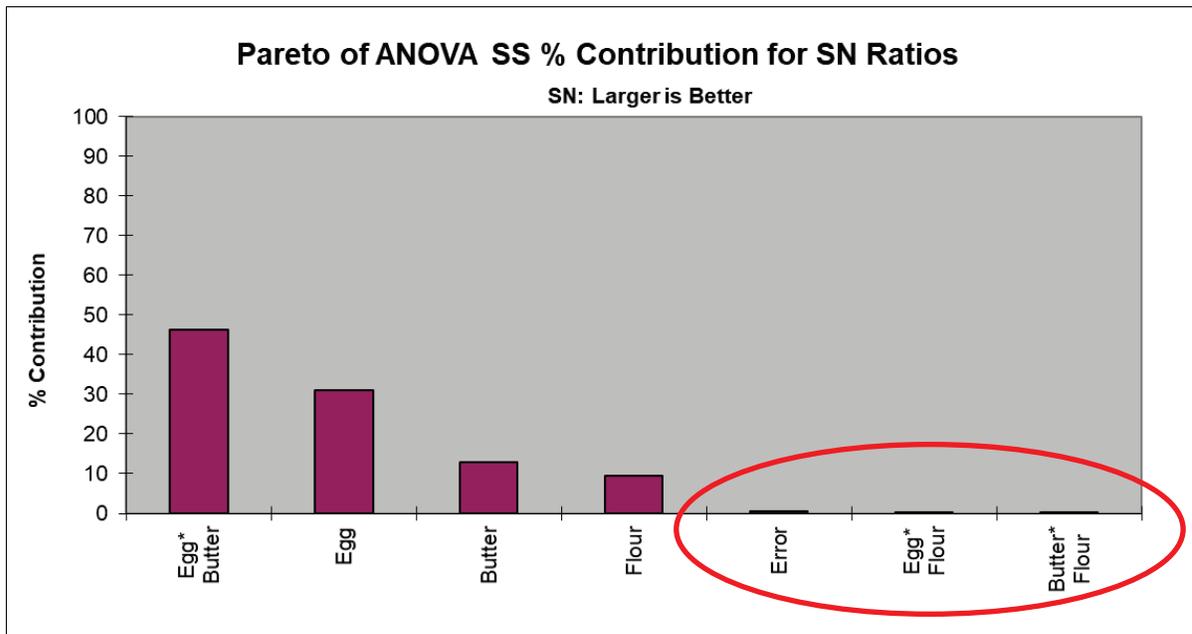
From the **Pareto of ANOVA SS % Contribution for Means**, we see that the Egg*Butter interaction is the dominant contributor to Mean Taste Score. Looking at the Egg*Butter **Interaction Plot for Means**, the optimum setting to maximize Mean Taste Score is Egg = 2 and Butter = 1. (Note: The Y Axis has been adjusted).

Looking at the **Pareto of ANOVA SS % Contribution for StDevs**, Flour is the dominant factor for Standard Deviation. The **Main Effects Plot for StDevs** shows that the optimum setting to minimize the StDev Taste Score is Flour = 2.

This is in agreement with the settings determined from the SN Ratio. Taguchi advocates the use of SN Ratio as it is simpler, looking at only one variable, but some practitioners prefer to analyze the Mean and Standard Deviation separately.

Multiple Regression and Excel Solver (Advanced Topics):

16. The Taguchi L8 Orthogonal Array (or L8 Dummy Coding) may be analyzed using Multiple Regression Analysis, adding and removing terms from the model as necessary. Terms that are removed will be used to estimate error for p-values. As noted earlier, SigmaXL does not compute a significance line in the Taguchi Template Pareto Charts. Taguchi recommends pooling of effects to estimate experimental error and this will be demonstrated here. Also, Solver cannot be used with the Predicted Output calculator in the Taguchi Template, but will be used in the Multiple Regression report.
17. We will be analyzing the SN: Larger is Better Ratio as the response in this example. The analysis may be repeated separately for the Mean and StDev.
18. From the **Pareto of ANOVA SS % Contribution for SN Ratios**, we see that the Error term (three-way interaction), Egg*Flour and Butter*Flour interactions are negligible compared to the other factors (each less than 1% contribution). As an initial attempt at pooling we will remove these terms from the regression model.



19. Highlight the L8 Orthogonal Array cells (**AL18:AV26**) as shown.

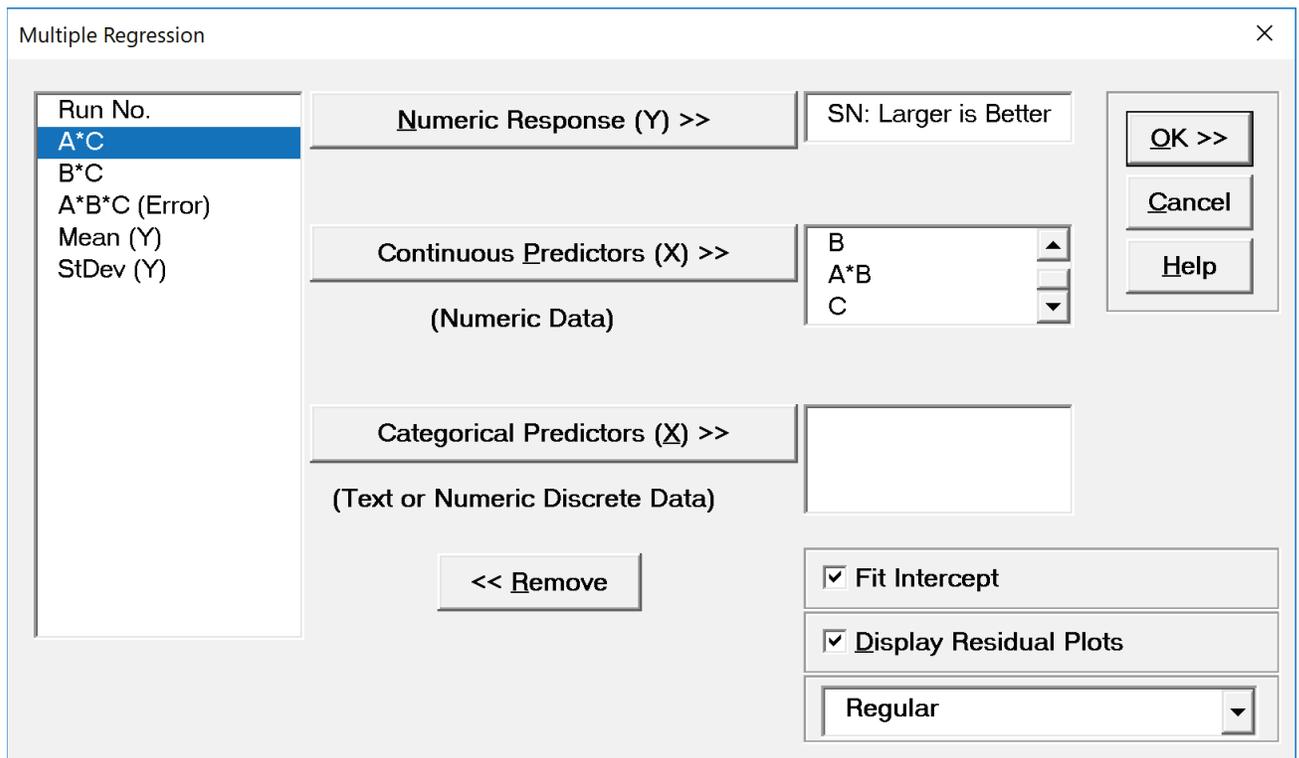
L8 Orthogonal Array:

Run No.	L8_1 A	L8_2 B	L8_3 A*B	L8_4 C	L8_5 A*C	L8_6 B*C	L8_7 A*B*C (Error)	Mean (Y)	StDev (Y)	SN: Larger is Better
6	1	1	1	1	1	1	1	3.52	2.449897957	5.010245078
2	1	1	1	2	2	2	2	3.5	1.380217374	7.673306706
7	1	2	2	1	1	2	2	4.74	1.479188967	12.39162469
1	1	2	2	2	2	1	1	5.2	0.935414347	13.97953162
5	2	1	2	1	2	1	2	5.38	1.454991409	13.71929378
8	2	1	2	2	1	2	1	5.9	0.681909085	15.27864445
3	2	2	1	1	2	2	1	4.36	1.817415748	11.23909439
4	2	2	1	2	1	1	2	4.86	0.665582452	13.51678365

The Orthogonal Array displays the factor letters, not the names, so recall that A is Egg, B is Butter and C is Flour.

20. Click **SigmaXL > Statistical Tools > Regression > Multiple Regression**. Click **Next**.

21. Select *SN: Larger is Better*, click **Numeric Response (Y) >>**; holding the CTRL key, select *A, B, A*B, and C*, click **Continuous Predictors (X) >>** as shown:



22. Click **OK**. The resulting regression report is shown:

Multiple Regression Model: SN: Larger is Better = (-7.210) + (3.675) * A + (2.361) * B + (4.482) * A*B + (2.022) * C

Model Summary:

R-Square	99.49%
R-Square Adjusted	98.82%
S (Root Mean Square Error)	0.382714

Parameter Estimates:

Predictor Term	Coefficient	SE Coefficient	T	P	VIF	Tolerance
Constant	-7.210	0.823057	-8.760	0.0031		
A	3.675	0.270619	13.579	0.0009	1	1.00000
B	2.361	0.270619	8.726	0.0032	1	1.00000
A*B	4.482	0.270619	16.564	0.0005	1	1.00000
C	2.022	0.270619	7.472	0.0050	1	1.00000

Analysis of Variance for Model:

Source	DF	SS	MS	F	P
Model	4	86.521	21.630	147.68	0.0009
Error	3	0.439409	0.146470		
Total (Model + Error)	7	86.961	12.423		

Durbin-Watson Test for Autocorrelation in Residuals:

DW Statistic	2.506
P-Value Positive Autocorrelation	0.5149
P-Value Negative Autocorrelation	0.0000

All terms are significant (using the pooled estimate of error) so we will use the model as is.

23. Now we will use the **Predicted Response Calculator** to determine the factor settings that maximize the SN Ratio:

Predicted Response Calculator:

Predictors	Enter Settings:	Predicted Response	Lower 95% CI	Upper 95% CI	Lower 95% PI	Upper 95% PI
A						
B						
A*B						
C						

24. All entries must be in coded units of 1 to 2 as given in the Orthogonal Array. While the calculator in the Taguchi Template uses strictly categorical factors, here we can use either categorical or continuous.

25. Since the Taguchi Orthogonal Array is coded 1, 2 rather than -1, +1, we cannot simply multiply Factor A * Factor B to compute the A*B interaction.

26. Turn on the Excel Row and Column headers in the Multiple Regression sheet: **File > Options > Advanced**. Scroll down to **Display options for the worksheet**. Check **Show row and column headers**. Click **OK**.

27. In cell A*B (**K14**), enter the formula:

$$=1.5 - 2*(K12-1.5)*(K13-1.5)$$

This formula centers and scales the A (**K12**) and B (**K13**) values so that the A*B (**K14**) value will match the values in the Taguchi array. Enter A = 1, B = 1, C = 1 (be careful to not overwrite the formula – you may wish to change the **K14** cell color as shown). The A*B value is 1.

Predicted Response Calculator:

Predictors	Enter Settings:	Predicted Response	Lower 95% CI	Upper 95% CI	Lower 95% PI	Upper 95% PI
A	1	5.330774831	4.367888571	6.29366109	3.778167389	6.883382272
B	1					
A*B	1					
C	1					

28. Optionally, test the formula entry, enter A=1, B=2, confirming that A*B = 2. Repeat for A=2, B=1 resulting in A*B = 2 and A=2, B=2 resulting in A*B = 1. Now reset back to A=1 and B=1.

29. Click Excel **Data > Solver** to open the Solver dialog (if the Solver menu option does not appear, click **File > Options > Add-Ins**. Adjacent to **Manage Excel Add-ins**, click **Go**. Check **Solver Add-in**. Click **OK**). Set the **Solver Parameters** as shown:

Solver Parameters

Set Objective:

To: Max Min Value Of:

By Changing Variable Cells:

Subject to the Constraints:

\$K\$12 <= 2
 \$K\$12 >= 1
 \$K\$13 <= 2
 \$K\$13 >= 1
 \$K\$15 <= 2
 \$K\$15 >= 1

Make Unconstrained Variables Non-Negative

Select a Solving Method:

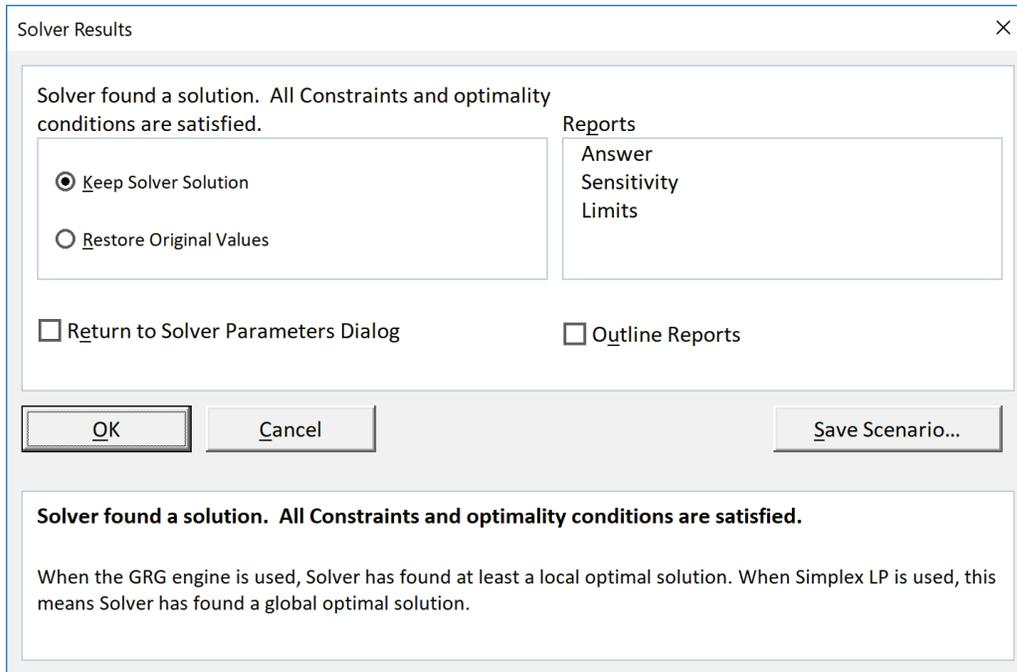
Solving Method
 Select the GRG Nonlinear engine for Solver Problems that are smooth nonlinear. Select the LP Simplex engine for linear Solver Problems, and select the Evolutionary engine for Solver problems that are non-smooth.

Cell **L12** is the predicted response (SN Ratio). Solver will try to maximize this value. Cells **K12**, **K13**, and **K15** are the A, B, and C Coded Factor Settings to be changed, with the constraints that the values are ≥ 1 and ≤ 2 . GRG Nonlinear is used here. (Note that if any of the factors are modeled as categorical, Evolutionary is recommended). Click **Solve**.

30. The Solver solution that maximizes SN Ratio is given in the **Predicted Response Calculator** as A (Egg) = 2, B (Butter) = 1 and C (Flour) = 2. These are in agreement with our earlier manual analysis. Confidence Intervals for the Mean SN Ratio and Prediction Intervals for Individual SN Ratio values are also reported.

Predicted Response Calculator:

Predictors	Enter Settings:	Predicted Response	Lower 95% CI	Upper 95% CI	Lower 95% PI	Upper 95% PI
A	2	15.50997018	14.54708392	16.47285643	13.95736273	17.06257762
B	1					
A*B	2					
C	2					



31. Click **OK** to keep the Solver solution.

32. This analysis may be repeated using the Mean and Standard Deviation as responses. A note of caution when using the Prediction Interval for the Mean, the individual values are averages across the Outer Array, not true individual values.

Taguchi L8 (2 Level) Four Factor – Catapult Example

This Catapult experiment is adapted from the article "THE CATAPULT PROBLEM: ENHANCED ENGINEERING MODELING USING EXPERIMENTAL DESIGN", by Schubert et al., Quality Engineering, 4(4), 463-473 (1992).

1. Click **SigmaXL > Help > Template Examples > Taguchi Examples > L8 Four Factor – Catapult** to open the example file. If you wish to start with a blank template and populate the values, click **SigmaXL > Design of Experiments > Basic Taguchi DOE Templates > Taguchi L8 (2 Level) > Four Factor (with Aliased Two-Way Interactions)**.
2. The response is Distance in Inches with Target = 50".
3. The three **Outer Array (Yrep1 to Yrep3)** values are repeat shots fired.
4. The goal is to hit a specific target of 50 inches with minimum variation, so we will maximize the Signal-to-Noise Ratio **SN: Nominal is Best (MSD with Target)**, selected from the list as shown:

Signal-to-Noise Ratio:	SN: Nominal is Best (MSD with Target)
Standard Deviation Response:	SN: Nominal is Best
	SN: Nominal is Best (Variance Only)
	SN: Nominal is Best (MSD with Target)
	SN: Larger is Better
	SN: Smaller is Better

Signal-to-Noise Ratio:	SN: Nominal is Best (MSD with Target)	SN Formula: $-10 \cdot \log_{10}(\text{Sum}((Y-T)^2)/n)$
Standard Deviation Response:	StDev (Y)	

5. Enter the Target value 50 in cell **H8**:

MSD Target:	50
-------------	----

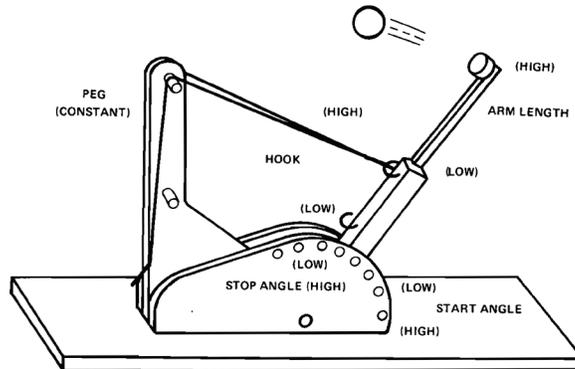
Note: If Target = 0, Nominal is Best (MSD with Target) is equivalent to Smaller Is Better. For use of MSD in SN Ratio, see Ranjit Roy (2010) *A Primer on the Taguchi Method*, Second Edition, Society of Manufacturing Engineers.

6. The **Inner Array** controllable factors used in this study are: Hook (Bottom/Top), Arm Length (Short/Long), Start Angle (150/177) and Stop Angle (Hole 4/Hole 5).

Factor	Factor Name	Level 1	Level 2
A	Hook	Bottom	Top
B	Arm Length	Short	Long
C	Start Angle	150	177
D	Stop Angle	Hole 4	Hole 5

Note that these are different factors and factor levels than those used earlier in the Full Factorial Catapult experiment (Part C), so the results cannot be compared. In that experiment

we used Pull Back Angle, Stop Pin and Pin Height. Here, Pin Height is kept fixed so not included, but Hook position and Arm Length are used. (In typical Catapult designs used in training today, Arm Length is adjusted by varying the Cup position).



7. Schubert used a Three Factor 8 Run Resolution IV design. We will use the Taguchi L8 Four Factor design. The column assignments in the SigmaXL Taguchi Templates are chosen to maximize the design resolution, so the L8 Four Factor Template is also Resolution IV, with main effects free and clear of two-way interactions and two-way interactions aliased with each other. The sort order used in this paper matches the Taguchi L8 standard sort order (left to right) so factor assignments remain the same. The paper uses the conventional -1 to +1 coding, whereas we will use the Taguchi 1, 2 coding.
8. Select Aliased Interactions using the drop-down list. Process knowledge, engineering or theory are used to make the selection and assume that the chosen interaction is dominant and the others are negligible. Aliased interactions are often associated with the largest main effects. Confirmation runs should always be used to validate the model. In this study, engineering knowledge was used to select all interactions that included the Stop Angle:

Aliased Interaction 1	Start Angle*Stop Angle
Aliased Interaction 2	Hook*Arm Length
Aliased Interaction 3	Start Angle*Stop Angle

Aliased Interaction 1	Start Angle*Stop Angle
Aliased Interaction 2	Arm Length*Stop Angle
Aliased Interaction 3	Hook*Stop Angle

Tip: For Taguchi L8 Five to Six Factors and L16 Nine to Fourteen Factors, not all possible interactions are available in the drop-down list (they are aliased with main effects). If an interaction of interest is not available in any of the drop-down lists, please modify the Factor Names so that they are assigned to columns shown in the list and then select that interaction.

9. A randomized run order was not given in the paper. The **Outer Array** Distance values are entered as shown:

Inner Array:

Std. Order	Run No.	Hook	Arm Length	Start Angle	Stop Angle
1	1	Bottom	Short	150	Hole 4
2	2	Bottom	Short	177	Hole 5
3	3	Bottom	Long	150	Hole 5
4	4	Bottom	Long	177	Hole 4
5	5	Top	Short	150	Hole 5
6	6	Top	Short	177	Hole 4
7	7	Top	Long	150	Hole 4
8	8	Top	Long	177	Hole 5

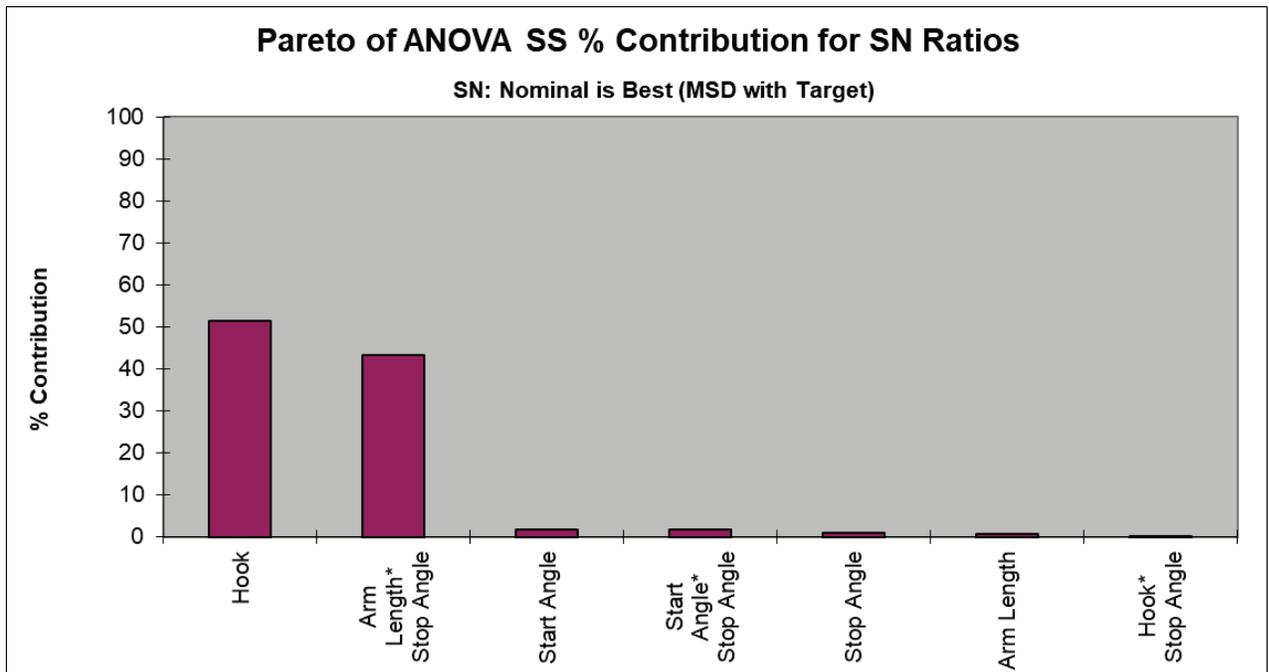
Outer Array:

Yrep1	Yrep2	Yrep3
28	27.1	26.2
46.3	43.5	46.5
21.9	21	20.1
52.9	53.7	52
75	73.1	74.3
127.7	126.9	128.7
86.2	86.5	87
195	195.9	195.7

Scroll to the right to view the calculated statistics:

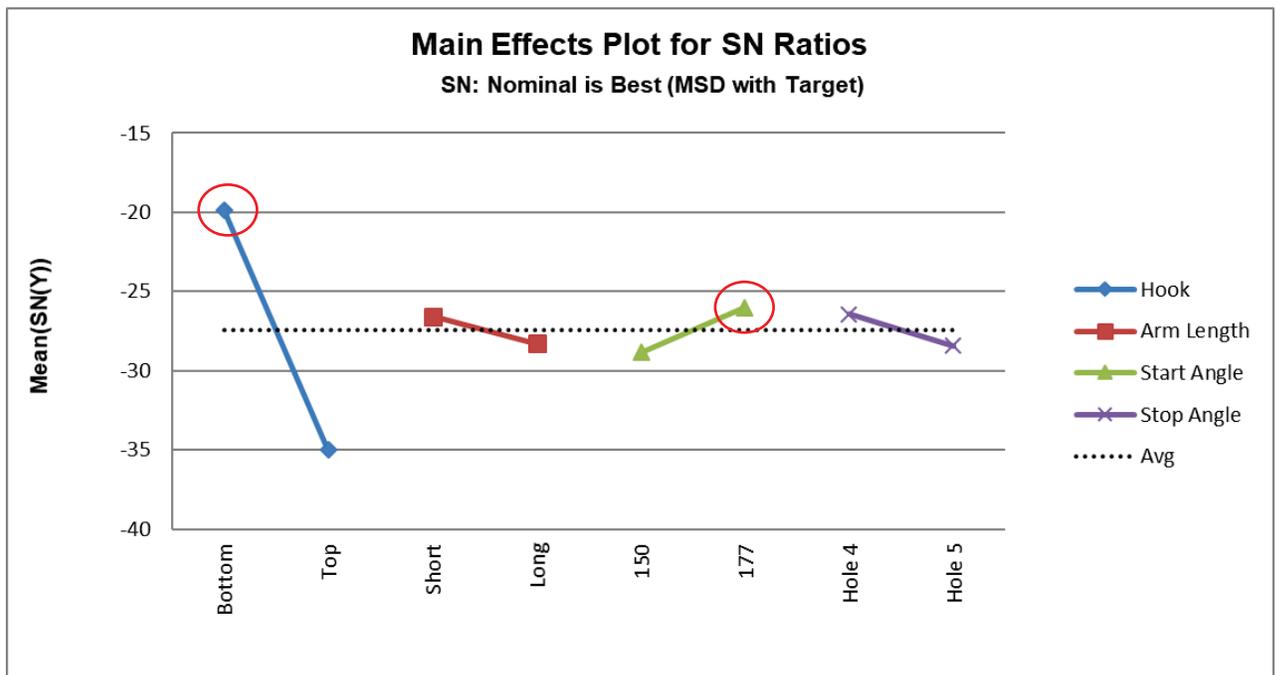
Mean (Y)	StDev (Y)	SN: Nominal is Best (MSD with Target)
27.1	0.9	-27.2011794
45.43333333	1.677299417	-13.56599436
21	0.9	-29.25074764
52.86666667	0.850490055	-9.395192526
74.13333333	0.960902354	-27.65693385
127.7666667	0.901849951	-37.81625904
86.56666667	0.404145188	-31.26206111
195.5333333	0.472581563	-43.25928006

10. We will focus our analysis on the Signal-To-Noise Ratio: Nominal is Best (MSD with Target). Since this design includes interactions, we will look at the **Pareto of ANOVA SS (Sum-of-Squares) % Contribution for SN Ratios**. Scroll right and down to view:



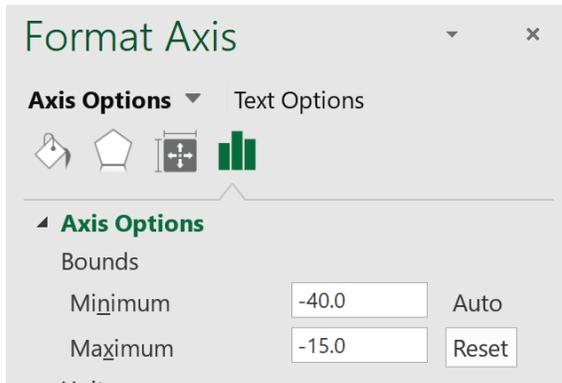
11. We see that the Hook position and Arm Length* Stop Angle interaction are the dominant factors.

12. Scroll down to view the **Main Effects Plot for SN Ratios**:

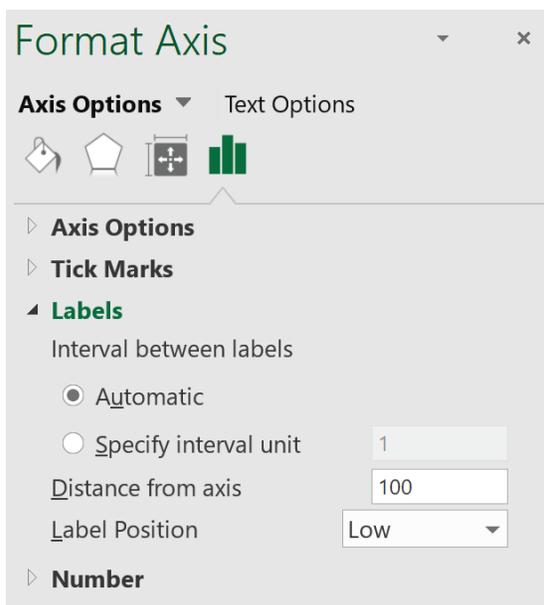


Setting Hook = Bottom and Start Angle = 177 maximizes the SN Ratio. (We will use the interaction plot to determine the optimum settings for Arm Length and Stop Angle).

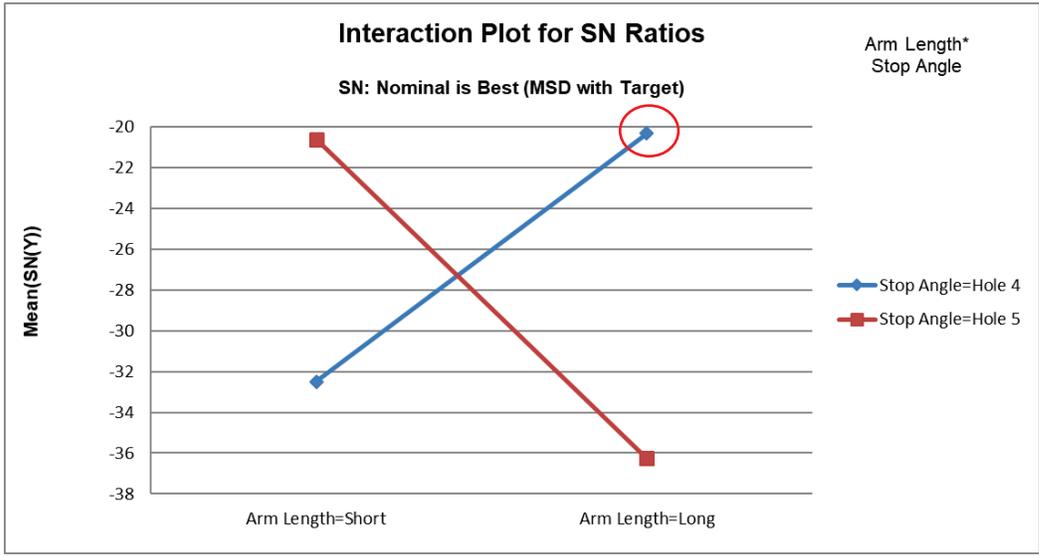
Note: The Y Axis has been adjusted to get a better scale (double click in the Y-Axis to open Excel Format Axis dialog, click **Axis Options** and set **Maximum** = -15 as shown).



With the negative SN values, the default X-axis labels appear as on top. To change the X-Axis location to bottom, double click on the X-Axis to open Excel Format Axis dialog, click **Axis Options, Labels** and set **Label Position** to Low as shown)



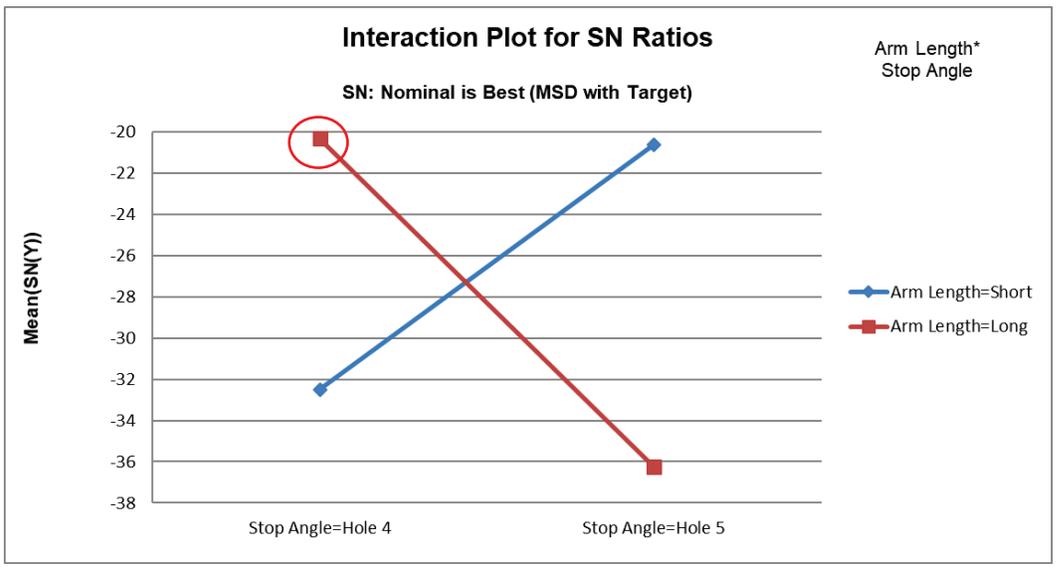
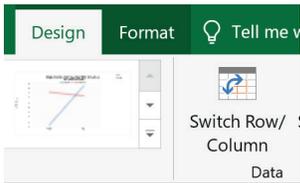
13. Scroll down to view the **Interaction Plot for SN Ratios** for the Arm Length * Stop Angle Interaction:



Here we can clearly see the large effect due to the two-way interaction. The effect that Arm Length has on Distance depends on what the Stop Angle is. The maximum SN Ratio is achieved with Arm Length = Long and Stop Angle = Hole 4.

Note 1: The Y-Axis has been adjusted to improve visibility (Maximum = -20).

Note 2: The Interaction Plot X-axis/Legend factors may be switched by clicking on the chart, then select Excel **Design > Switch Row/Column**:



14. From the above plots, we conclude that the optimum setting to produce a maximum SN Ratio for Distance = 50" is: Hook = Bottom, Arm Length = Long, Start Angle = 177 and Stop Angle = Hole 4.
15. Scroll up and left to enter these values in the **Predicted Output** Calculator. Factor settings can be selected from the drop-down list as shown:

Predicted Output for Y:

Factor Name	Enter Actual Factor Setting - Uncoded	Factor Setting Coded
Hook	Bottom	1
Arm Length	Bottom	2
Start Angle	Top	2
Stop Angle	Hole 4	1

Predicted Output for Y:

Factor Name	Enter Actual Factor Setting - Uncoded	Factor Setting Coded	\hat{Y}	\hat{S}	SN: Nominal is Best (MSD with Target)
Hook	Bottom	1	52.86666667	0.850490055	-9.395192526
Arm Length	Long	2			
Start Angle	177	2			
Stop Angle	Hole 4	1			

This gives the predicted Mean, Standard Deviation and SN Ratio for Distance and agrees with the results given in the paper. The predicted value of 52.9 inches is slightly off of the Target of 50 inches due to the discrete settings of the categorical factors.

Excel Solver cannot be used with this calculator to perform optimization, so we will use Multiple Regression and Solver, treating the Start Angle as continuous rather than categorical. This will allow interpolation in order to obtain a predicted Mean response of 50 inches.

Multiple Regression and Excel Solver (Advanced Topics):

16. The Taguchi L8 Orthogonal Array (or L8 Dummy Coding) may be analyzed using Multiple Regression Analysis, adding and removing terms from the model as necessary. However, for this example, we will not remove any terms from the model, following the analysis as performed in the paper.
17. We will be analyzing the Mean Distance to refine the settings determined above from **SN: Nominal is Best (MSD with Target)**.

18. Highlight the **L8 Orthogonal Array** cells (**AM23:AW31**) as shown.

L8 Orthogonal Array:

Run No.	L8_1 A	L8_2 B	L8_3 C*D	L8_4 C	L8_5 B*D	L8_6 A*D	L8_7 D	Mean (Y)	StDev (Y)	SN: Nominal is Best (MSD with Target)
1	1	1	1	1	1	1	1	27.1	0.9	-27.2011794
2	1	1	1	2	2	2	2	45.43333333	1.677299417	-13.56599436
3	1	2	2	1	1	2	2	21	0.9	-29.25074764
4	1	2	2	2	2	1	1	52.86666667	0.850490055	-9.395192526
5	2	1	2	1	2	1	2	74.13333333	0.960902354	-27.65693385
6	2	1	2	2	1	2	1	127.7666667	0.901849951	-37.81625904
7	2	2	1	1	2	2	1	86.56666667	0.404145188	-31.26206111
8	2	2	1	2	1	1	2	195.5333333	0.472581563	-43.25928006

The Orthogonal Array displays the factor letters, not the names, so recall that A is Hook, B is Arm Length, C is Start Angle and D is Stop Angle.

19. Click **SigmaXL > Statistical Tools > Regression > Multiple Regression**. Click **Next**.

20. Select *Mean (Y)*, click **Numeric Response (Y) >>**; now, one at a time carefully select the factors in the sequence: *A, B, C, D, A*D, B*D and C*D*, clicking **Continuous Predictors (X) >>** each time. Using this order will make subsequent analysis easier. Uncheck **Display Residual Plots**.

Multiple Regression

Run No.
StDev (Y)
SN: Nominal is Best (MSD with Target)

Numeric Response (Y) >> Mean (Y)

Continuous Predictors (X) >>
(Numeric Data)
A*D
B*D
C*D

Categorical Predictors (X) >>
(Text or Numeric Discrete Data)

<< Remove

Fit Intercept

Display Residual Plots

Regular

OK >>
Cancel
Help

21. Click **OK**. The resulting regression report is shown:

Multiple Regression Model: Mean (Y) = (-76.300) + (84.400) * A + (20.383) * B + (53.200) * C + (10.450) * D + (-17.

Model Summary:	
R-Square	100.00%
R-Square Adjusted	0.00%
S (Root Mean Square Error)	

Parameter Estimates:

Predictor Term	Coefficient	SE Coefficient	T	P	VIF	Tolerance
Constant	-76.300					
A	84.400				1	1.00000
B	20.383				1	1.00000
C	53.200				1	1.00000
D	10.450				1	1.00000
A*D	-17.217				1	1.00000
B*D	-28.100				1	1.00000
C*D	-19.717				1	1.00000

Analysis of Variance for Model:

Source	DF	SS	MS	F	P
Model	7	23906	3415.2		
Error	0	2.38026E-23			
Total (Model + Error)	7	23906	3415.2		

ANOVA and P-Values are not computed since there are no degrees of freedom for the error term.

22. Now we will use the **Predicted Response Calculator** to determine the factor settings that achieve the Target value of 50 inches.

Predicted Response Calculator:

Predictors	Enter Settings:	Predicted Response
A		
B		
C		
D		
A*D		
B*D		
C*D		

23. All entries must be in coded units of 1 to 2 as given in the Orthogonal Array. While the calculator in the Taguchi Template uses strictly categorical factors, here we can use either categorical or continuous. We will use the categorical settings obtained from the SN Ratio analysis for: A = 1 (Hook = Bottom), B = 2 (Arm Length = Long), and D = 1 (Stop Angle = Hole 4). Factor C (Start Angle) is continuous and will be determined using Solver in order to achieve the Target Distance = 50 inches. The initial setting will be C = 2 (Start Angle = 177).

24. Since the Taguchi Orthogonal Array is coded 1, 2 rather than -1, +1, we cannot simply multiply Factor A * Factor D to compute the A*D interaction.

25. Turn on the Excel Row and Column headers in the Multiple Regression sheet: **File > Options > Advanced**. Scroll down to **Display options for the worksheet**. Check **Show row and column headers**. Click **OK**.

26. In cell A*D (**K16**), enter the formula:

$$=1.5 - 2*(K12-1.5)*(K15-1.5)$$

This formula centers and scales the A (**K12**) and D (**K15**) values so that the A*D (**K16**) value will match the values in the Taguchi array.

27. In cell B*D (**K17**), enter the formula:

$$=1.5 - 2*(K13-1.5)*(K15-1.5)$$

28. In cell C*D (**K18**), enter the formula:

$$=1.5 - 2*(K14-1.5)*(K15-1.5)$$

29. Enter A = 1, B = 1, C=1, D=1 (be careful to not overwrite the formulas – you may wish to change the **K16** to **K18** cell color as shown). The cells should show A*D = 1, B*D = 1 and C*D = 1.

	I	J	K	L
10	Predicted Response Calculator:			
11		Predictors	Enter Settings:	Predicted Response
12		A	1	27.1
13		B	1	
14		C	1	
15		D	1	
16		A*D	1	
17		B*D	1	
18		C*D	1	

Confidence and Prediction Intervals are not given because the model does not include an error term.

30. As mentioned above, we will use the categorical settings obtained from the SN Ratio analysis, but Factor C (Start Angle) is continuous and will be determined using Solver in order to achieve the Target Distance = 50 inches. Enter the values as shown:

Predicted Response Calculator:			
	Predictors	Enter Settings:	Predicted Response
	A	1	52.86666667
	B	2	
	C	2	
	D	1	
	A*D	1	
	B*D	2	
	C*D	2	

The initial Predicted Response matches the Mean in the Predicted Output given in the Taguchi Template.

31. Click Excel **Data > Solver** to open the Solver dialog (if the Solver menu option does not appear, click **File > Options > Add-Ins**. Adjacent to **Manage Excel Add-ins**, click **Go**. Check **Solver Add-in**. Click **OK**). Set the **Solver Parameters** as shown:

Solver Parameters

Set Objective:

To: Max Min Value Of:

By Changing Variable Cells:

Subject to the Constraints:

\$K\$14 <= 2
\$K\$14 >= 1

Make Unconstrained Variables Non-Negative

Select a Solving Method:

Solving Method
Select the GRG Nonlinear engine for Solver Problems that are smooth nonlinear. Select the LP Simplex engine for linear Solver Problems, and select the Evolutionary engine for Solver problems that are non-smooth.

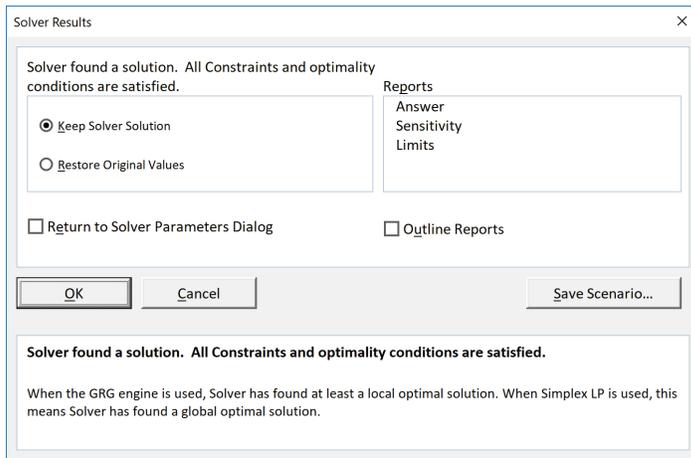
Buttons: Add, Change, Delete, Reset All, Load/Save, Options, Help, Solve, Close

Cell **L12** is the predicted response (Mean Distance). Solver will try to set this to the Target value of 50 inches. Cell **K14** is the C Coded Factor Setting to be changed, with the constraint that the value is ≥ 1 and ≤ 2 . GRG Nonlinear is used here. (Note that if any of the Solver factors are modeled as categorical, Evolutionary is recommended).

32. Click **Solve**.

Predicted Response Calculator:

Predictors	Enter Settings:	Predicted Response
A	1	50
B	2	
C	1.914385266	
D	1	
A*D	1	
B*D	2	
C*D	1.914385266	



33. The Predicted Response is 50 inches. Click **OK** to keep the Solver solution.

34. Now we have to convert the coded C value to uncoded units. In cell **L14**, enter the formula:

$$=27*K14+123$$

This converts the coded 1 to 2 back to the uncoded 150 to 177 for the Start Angle.

Predicted Response Calculator:

Predictors	Enter Settings:	Predicted Response
A	1	50
B	2	
C	1.914385266	174.6884022
D	1	
A*D	1	
B*D	2	
C*D	1.914385266	

35. The factor setting of 174.7 degrees to obtain a Mean Distance of 50 inches matches the value given in the paper. Confirmation runs were performed and validated the model prediction, resulting in greater precision than a theoretical model requiring 200 person-hours to develop!

Taguchi L9 (3 Level) Four Factor – Paper Airplane Example

This Paper Airplane Experiment is adapted from the article "Teaching Taguchi's Approach to Parameter Design", by Sanjiv Sarin, Quality Progress, May 1997.

1. Click **SigmaXL > Help > Template Examples > Taguchi Examples > L9 Four Factor – Paper Airplane** to open the example file. If you wish to start with a blank template and populate the values, click **SigmaXL > Design of Experiments > Basic Taguchi DOE Templates > Taguchi L9 (3 Level) > Four Factor**.
2. The response is Distance of Flight in Inches. The goal is to maximize Distance.
3. We will maximize the Signal-to-Noise Ratio **SN: Larger is Better**, selected from the list as shown:

Signal-to-Noise Ratio:	SN: Larger is Better
Standard Deviation Response:	SN: Nominal is Best SN: Nominal is Best (Variance Only) SN: Nominal is Best (MSD with Target)
Factor	SN: Larger is Better SN: Smaller is Better

Signal-to-Noise Ratio:	SN: Larger is Better
Standard Deviation Response:	StDev (Y)

SN Formula: $-10 \cdot \log_{10}(\text{Sum}(1/Y^2)/n)$

4. The **Inner Array** Controllable Input factors used in this study are: Paper Weight (1 Sheet, 2 Sheets, 3 Sheets), Design (Design 1, Design 2, Design 3), Paper Width (4 inches, 6 inches, 8 inches) and Paper Length (6 inches, 8 inches, 10 inches).

Factor	Factor Name	Level 1	Level 2	Level 3
A	Paper Weight	1	2	3
B	Design	1	2	3
C	Paper Width	4	6	8
D	Paper Length	6	8	10

5. The Uncontrollable Factors are Launch Height, Launch Angle and Ground Surface. The **Outer Array** was created using an L4 Three Factor design. Scroll right to view the screen capture:

Outer Array created using Taguchi L4 Three Factor Template:

Factor	Factor Name	Level 1	Level 2
A	Launch Height	Ground	Chair
B	Launch Angle	Horizontal	45 Degrees
C	Ground Surface	Concrete	Wood

Run No.	Launch Height	Launch Angle	Ground Surface
1	Ground	Horizontal	Concrete
2	Ground	45 Degrees	Wood
3	Chair	Horizontal	Wood
4	Chair	45 Degrees	Concrete

6. The distance values are entered into the **Outer Array** of the L9 as **Yrep1** to **Yrep4**:

Inner Array:

Std. Order	Run No.	Paper Weight	Design	Paper Width	Paper Length
1	1	1	1	4	6
2	2	1	2	6	8
3	3	1	3	8	10
4	4	2	1	6	10
5	5	2	2	8	6
6	6	2	3	4	8
7	7	3	1	8	8
8	8	3	2	4	10
9	9	3	3	6	6

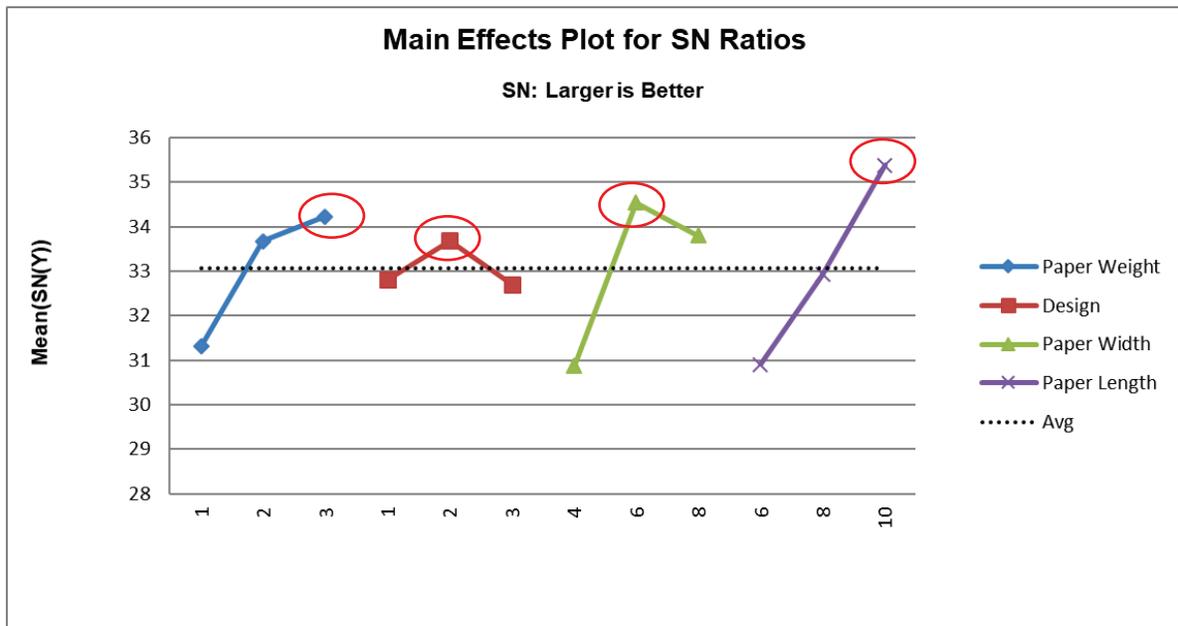
Outer Array:

Yrep1	Yrep2	Yrep3	Yrep4
49	44	12	38
91	42	44	38
59	48	39	67
116	89	48	88
32	38	56	108
24	55	39	46
42	76	122	41
50	73	47	65
76	34	80	37

Scroll to the right to view the calculated statistics:

Mean (Y)	StDev (Y)	SN: Larger is Better
35.75	16.45954637	26.69079814
53.75	24.95829855	33.24061006
53.25	12.28481447	33.97524775
85.25	28.01636427	37.1912413
58.5	34.53983208	32.85315172
41	13.08943594	30.97352457
70.25	38.14336989	34.55282905
58.75	12.33896268	34.95858081
56.75	24.6221445	33.17653559

7. Since there are no interactions, we will focus our analysis on the Main Effects Plots for Signal-To-Noise Ratio and Mean Distance as done in the paper. Scroll right and down to view the **Main Effects Plot for SN Ratios**:



The optimum settings to maximize SN Ratio are: Paper Weight = 3, Design = 2, Paper Width = 6 and Paper Length = 10.

8. Scroll left to look at the **Main Effects Plot for Means**:



The optimum settings to maximize the Mean are: Paper Weight = 3, Design = 1, Paper Width = 6 and Paper Length = 10. (Note, The Y-Axis minimum has been adjusted to 40 to improve the scale).

9. Commenting on this difference, Sarin states:

Factor B. For the factor of design, a plane built using design 1 was superior based on the $\bar{\bar{X}}$ value. Based on the average S/N ratio, however, the difference between design 1 and design 2 was marginal. When there is disagreement between the level that optimizes the $\bar{\bar{X}}$ value and the one that maximizes the average S/N ratio, it is usually recommended that the average S/N ratio be given priority. This is due to the fact that the Taguchi approach focuses first on selecting levels of controllable factors that reduce variability; bringing the process on target is the secondary objective. Thus, according to the Taguchi approach, design 2 should be recommended. But as a matter of practicality, design 1 was chosen since it results in significant improvement in the $\bar{\bar{X}}$ value while sacrificing the S/N value to a relatively lesser degree.

10. We will compare the two settings using the Predicted Output Calculator. Scroll up and enter (or select from the drop-down list) as shown:

Predicted Output for Y:

Factor Name	Enter Actual Factor Setting - Uncoded	Factor Setting Coded	\hat{Y}	\hat{S}	SN: Larger is Better
Paper Weight	3	3	78.83333333	24.24191679	38.62040862
Design	2	2			
Paper Width	6	2			
Paper Length	10	3			

Now change the Design to 1 as shown:

Predicted Output for Y:

Factor Name	Enter Actual Factor Setting - Uncoded	Factor Setting Coded	\hat{Y}	\hat{S}	SN: Larger is Better
Paper Weight	3	3	85.58333333	27.83597919	37.74791725
Design	1	1			
Paper Width	6	2			
Paper Length	10	3			

11. As noted by the author, depending on the priority given to Mean or Signal-To-Noise Ratio, one would select Design 1 or 2. In either case, confirmation runs need to be performed to validate the experimental results.

SigmaXL: Control Phase Tools: Statistical Process Control (SPC) Charts and Time Series Forecasting

Part A - Individuals Charts

Tip: See [Part E – Control Chart Selection Tool](#). The Control Chart Selection Tool makes it easy for you to select the correct statistical process control chart depending on data type and subgroup/sample type and size.

Individuals Chart Template

Click **SigmaXL > Templates and Calculators > Control Chart Templates > Basic > Individuals.**

This template is also located at **SigmaXL > Control Charts > Control Chart Templates > Basic > Individuals.**

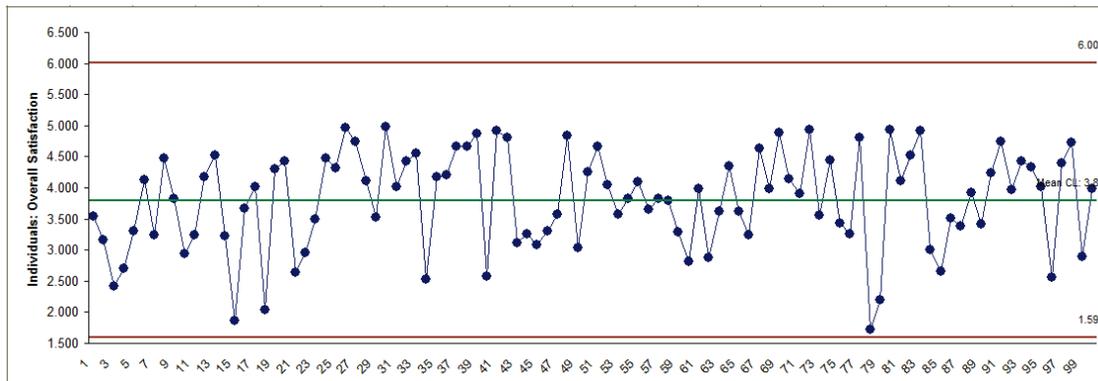
See **Measure Phase Part B – Templates and Calculators** for Individuals Chart template example:

[Basic Control Chart Templates – Individuals Chart](#)

Individuals Chart

1. Open **Customer Data.xlsx**, click on **Sheet 1**. Click **SigmaXL>Control Charts>Individuals**. Ensure that entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
2. Select *Overall Satisfaction*, click **Numeric Data Variable (Y) >>**. Ensure that **Calculate Limits** is selected.

3. Click **OK**. Resulting Individuals control chart:



4. We have seen this data earlier as a run chart. The Control Chart adds calculated control limits. Note that the Upper Control Limit exceeds the survey upper limit of 5. Here it would be appropriate to change the UCL to 5.0. Click **Recall SigmaXL Dialog** menu or press **F3** to recall last dialog.
5. Select *Overall Satisfaction*, click **Numeric Data Variable (Y) >>**, change **UCL** to 5.

Individuals Chart

Customer Record No
Order Date
Customer Type
Avg No. of orders per mo
Avg days Order to deliver
Loyalty - Likely to Recom
Overall Satisfaction
Responsive to Calls
Ease of Communications
Staff Knowledge
Size of Customer
Major-Complaint
Product Type
Sat-Discrete

Numeric Data Variable (Y) >> Overall Satisfaction

Optional X-Axis Labels >>

<< Remove

Calculate Limits
 Enter Limits

5 UCL
3.8013 CL
1.596271 LCL

Tests for Special Causes

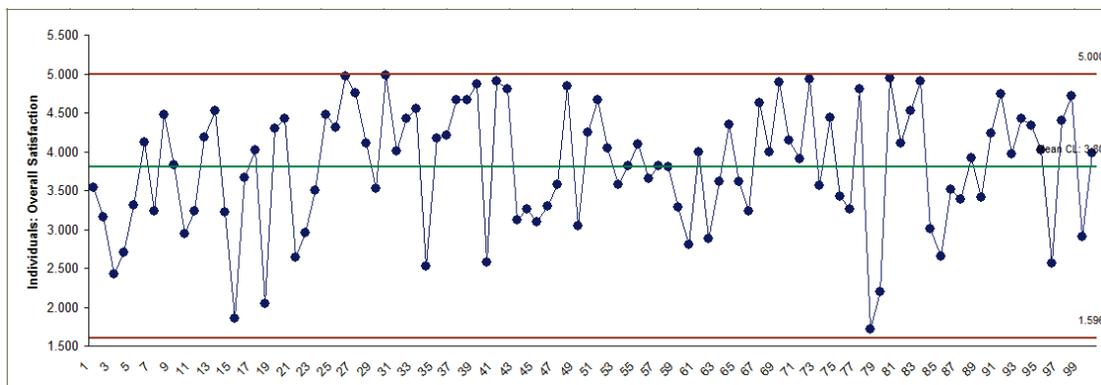
Sigma Zone Lines

Advanced Options

Add Title

OK >>
Cancel
Help

6. Click **OK**. Resulting Individuals chart with modified UCL:



Tip: You should not change the calculated control limits unless you have a legitimate reason such as a boundary condition. Do not replace the control limits with specification limits – the control chart will lose its statistical ability to detect assignable causes. We will redo the Individuals chart for Overall Satisfaction later using the Individuals Nonnormal tool.

7. Open **Delivery Times.xlsx**. Click **Sheet 1** Tab. This data set contains room service delivery time deviations in minutes. The Critical Customer Requirement is target time +/- 10 minutes.
8. Click **SigmaXL > Control Charts > Individuals**. Ensure that entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
9. Select *Delivery Time Deviation*, click **Numeric Data Variable (Y) >>**, check **Tests for Special Causes**, check **Sigma Zone Lines**. Ensure that **Calculate Limits** is selected.
10. Select **Advanced Options**, specify **LSL = -10, Target = 0, USL = 10** as shown:

Individuals Chart
X

Delivery Time Deviation
 Defects
 Floor

Numeric Data Variable (Y) >>

Optional X-Axis Labels >>

Calculate Limits
 Historical Limits

UCL

Tests for Special Causes

CL

Sigma Zone Lines

LCL

Add Title

Process Capability Report
 LSL Target USL

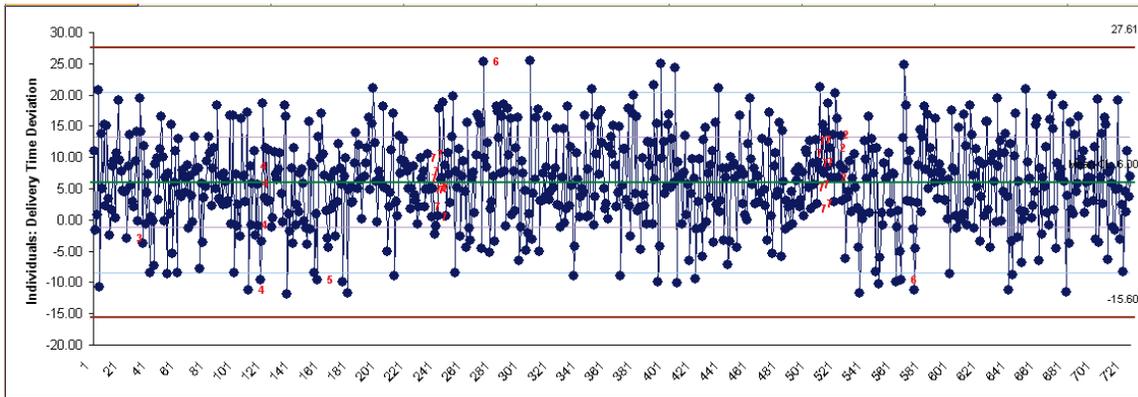
Calculate Control Limits:
 Specify Subgroup Number for Calculation of Control Limits
 Specify Historical Group Column (Split Limits)

Estimate:
 Average MR
 Median MR

Delivery Time Deviation
 Defects
 Floor

Individual CL
 Mean
 Median

11. Click **OK**. The resulting chart is shown:

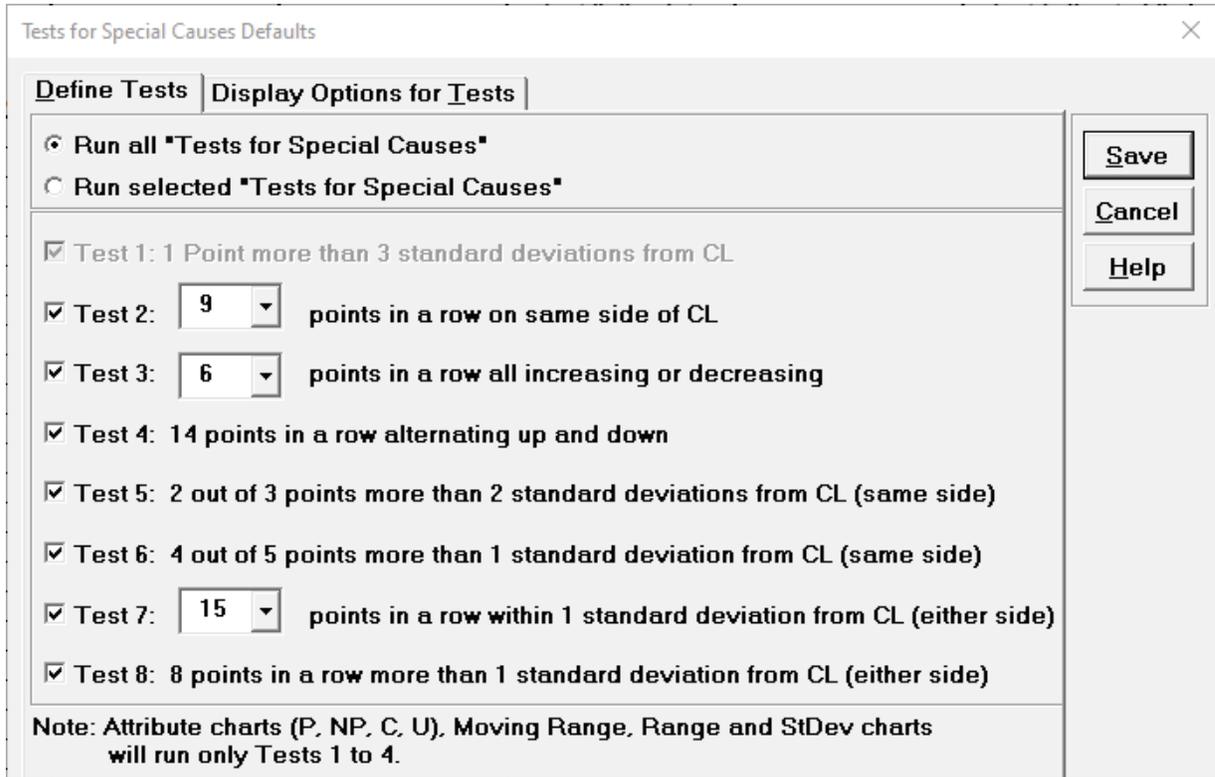


Some tests for special causes are indicated on the chart. If more than one test fails, the number corresponds to the first failed test.

12. There are no points that exceed the +/- 3 sigma limits on this chart, but we see some indication of instability with tests for special causes. The **Tests for Special Causes** report below the chart provides detailed information about each observation identified as a special cause. Note that the control chart also shows the +/- 1 sigma and +/- 2 sigma lines to aid in viewing these tests.

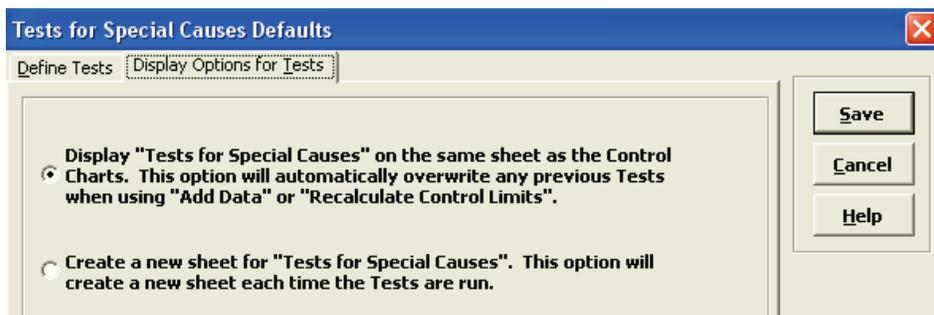
Tests for Special Causes - Individuals: Delivery Time Deviation								
Number of Data Points Failing Tests = 30								
Observation No.	Test 1: 1 point more than 3 Stdev from CL	Test 2: 9 points in a row on same side of CL	Test 3: 6 points in a row all increasing or all decreasing	Test 4: 14 points in a row alternating up and down	Test 5: 2 out of 3 points more than 2 Stdev from CL (same side)	Test 6: 4 out of 5 points more than 1 Stdev from CL (same side)	Test 7: 15 points in a row within 1 Stdev from CL (either side)	Test 8: 8 points in a row more than 1 Stdev from CL (either side)
24			x					
109				x				
110				x				
111				x				
112				x				
157					x			
229							x	
230							x	
231							x	
232							x	
233							x	
234							x	
235							x	
236							x	
237							x	
238							x	
273						x		
499							x	
500							x	
501							x	
502							x	
503							x	
504							x	
505							x	
506							x	
507							x	
515		x						
516		x						
517		x						
565						x		

13. These tests for special causes can have defaults set to apply any or all of Tests 1-8. Test 2 can be set to 7, 8, or 9 points in a row on same side of CL. Test 3 can be set to 6 or 7 points in a row all increasing or decreasing. Test 7 can be set to 14 or 15 points in a row within 1 standard deviation from CL. Click **Sheet 1 Tab**. Click **SigmaXL > Control Charts > "Tests for Special Causes" Defaults** to run selected tests for special causes:



Note that these defaults will apply to Individuals and X-bar charts. Test 1 to 4 settings will be applied to Attribute Charts and (if checked) Moving Range, Range, StDev and Attribute Charts.

14. Click on the Tab **Display Options for Tests**.



15. If you prefer to create a separate worksheet for each **Tests for Special Causes** report, choose **Create a new sheet** option. The default is to display **Tests for Special Causes** on the same sheet as the Control Chart. Note that this report will be overwritten when you add data or recalculate control limits.

16. Click **Save**.

17. Click on **Indiv Proc Cap** Tab to view the Process Capability Report, which includes potential (short term) capability indices Cp and Cpk:

Process Capability Report: Delivery Time Deviation		
Count		725
Mean		6.004
StDev (Overall, Long Term)		7.162
StDev (Within, Short Term)		7.202
USL		10
Target		0
LSL		-10

Capability Indices using Overall StDev		
Pp		0.47
Ppu		0.19
Ppl		0.74
Ppk		0.19
Cpm		0.36

Potential Capability Indices using Within StDev		
Cp		0.46
Cpu		0.18
Cpl		0.74
Cpk		0.18

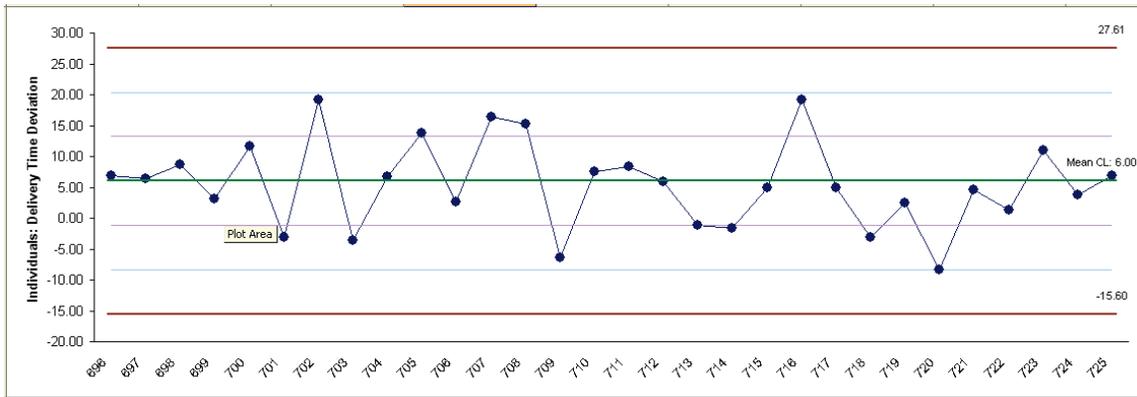
Expected Overall Performance		
ppm > USL		288409
ppm < LSL		12720.5
ppm Total		301129.8
% > USL		28.84%
% < LSL		1.27%
% Total		30.11%

Actual (Empirical) Performance		
% > USL		26.90%
% < LSL		1.38%
% Total		28.28%

Anderson-Darling Normality Test		
A-Squared		0.708616
P-Value		0.064143

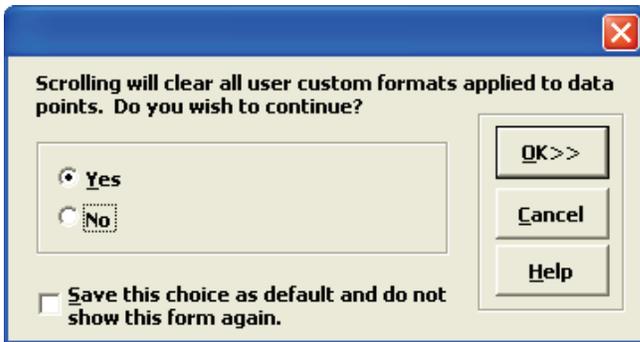
18. While this process demonstrated some slight instability on the control charts, the bigger issue was being late 6 minutes on average and having a Standard Deviation of 7.2 minutes! One improvement implemented was rescheduling the service elevators so that Room Service and Maintenance were not both trying to use them during peak times.

19. Click on the **Indiv** sheet. With 725 data points, you may want to have a closer look at the most recent data. To do this, click **SigmaXL Chart Tools > Show Last 30 Points**. (If this menu item does not appear, click on any cell adjacent to the chart.) The resulting chart is shown:



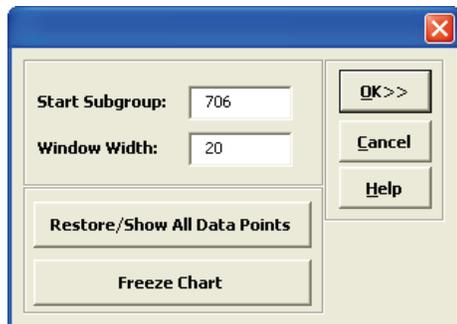
20. To reset the chart, click **SigmaXL Chart Tools > Show All Data Points**.

21. To enable scrolling, click **SigmaXL Chart Tools > Enable Scrolling**. A warning message is given:



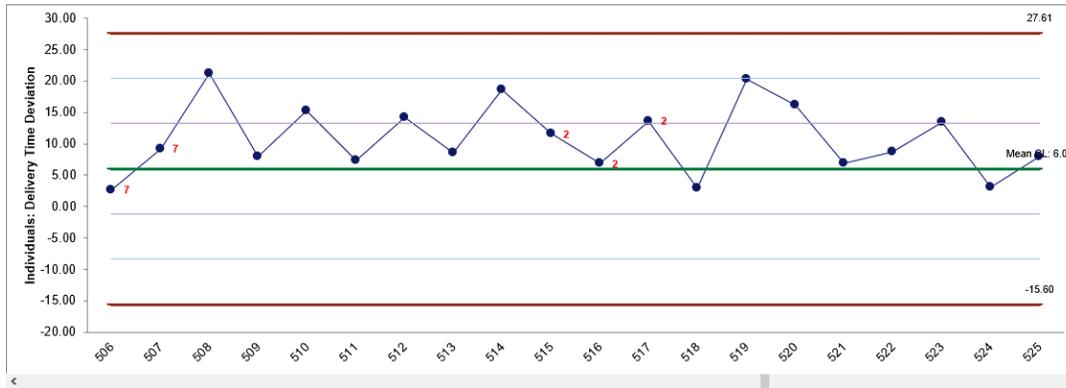
Scrolling will clear user custom formats, but does not affect Tests for Special Causes.

22. Click **OK**. A scroll dialog appears allowing you to specify the **Start Subgroup** and **Window Width**:



23. At any point, you can click **Restore/Show All Data Points** or **Freeze Chart**. Freezing the chart will remove the scroll and unload the dialog. The scroll dialog will also unload if you change worksheets. To restore the dialog, click **SigmaXL Chart Tools > Enable Scrolling**.

24. Click **OK**. The control chart appears with the scroll bar beneath it. You can also change the **Start Subgroup** and **Window Width** and **Update**.



X

Start Subgroup: <input style="width: 40px;" type="text" value="506"/>	<input type="button" value="Update"/>
Window Width: <input style="width: 40px;" type="text" value="20"/>	<input type="button" value="Cancel"/>
<input type="button" value="Help"/>	
<input type="button" value="Restore/Show All Data Points"/>	
<input type="button" value="Freeze Chart"/>	

Individuals Charts: Advanced Limit Options – Historical Groups

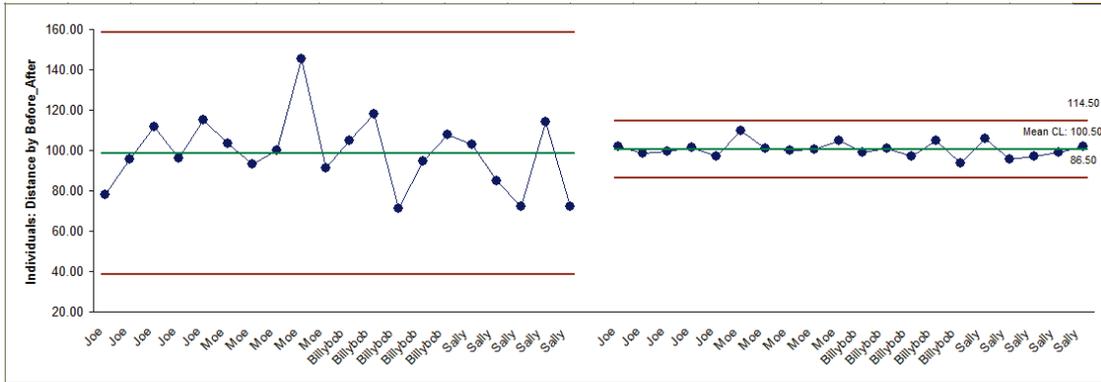
1. Open **Catapult Data - Before After Improvement.xlsx**. Click **Sheet 1 Tab**. This data set contains Catapult firing distances. Before_After denotes before improvement and after improvement. The target distance is 100 inches, with the goal being to hit the target and minimize variation about the target. We would like to use an individuals control chart with historical groups to split the limits demonstrating the before versus after improvement.
2. Click **SigmaXL > Control Charts > Individuals**. Ensure that entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
3. Select *Distance*, click **Numeric Data Variable (Y) >>**; select *Operator*, click **Optional X-Axis Labels >>**.
4. Click **Advanced Options**. Select **Specify Historical Groups**. Select *Before_After*, click **Historical Groups >>**.

The screenshot shows the 'Individuals Chart' dialog box with the following settings:

- Numeric Data Variable (Y) >>**: Distance
- Optional X-Axis Labels >>**: Operator
- Calculate Limits**: Calculate Limits, Historical Limits
- Tests for Special Causes**: Tests for Special Causes
- Sigma Zone Lines**: Sigma Zone Lines
- Advanced Options**: Add Title
- Process Capability Report**: LSL, Target, USL (all empty)
- Estimate**: Average MR, Median MR
- Individual CL**: Mean, Median
- Calculate Control Limits**: Specify Subgroup Number for Calculation of Control Limits, Specify Historical Group Column (Split Limits)
- Historical Groups >>**: Before_After

Note: Process Capability analysis is not permitted when Historical Groups are used.

- Click **OK**. The resulting Individuals Control Chart with split limits based on historical groups is shown, demonstrating a clear process improvement:

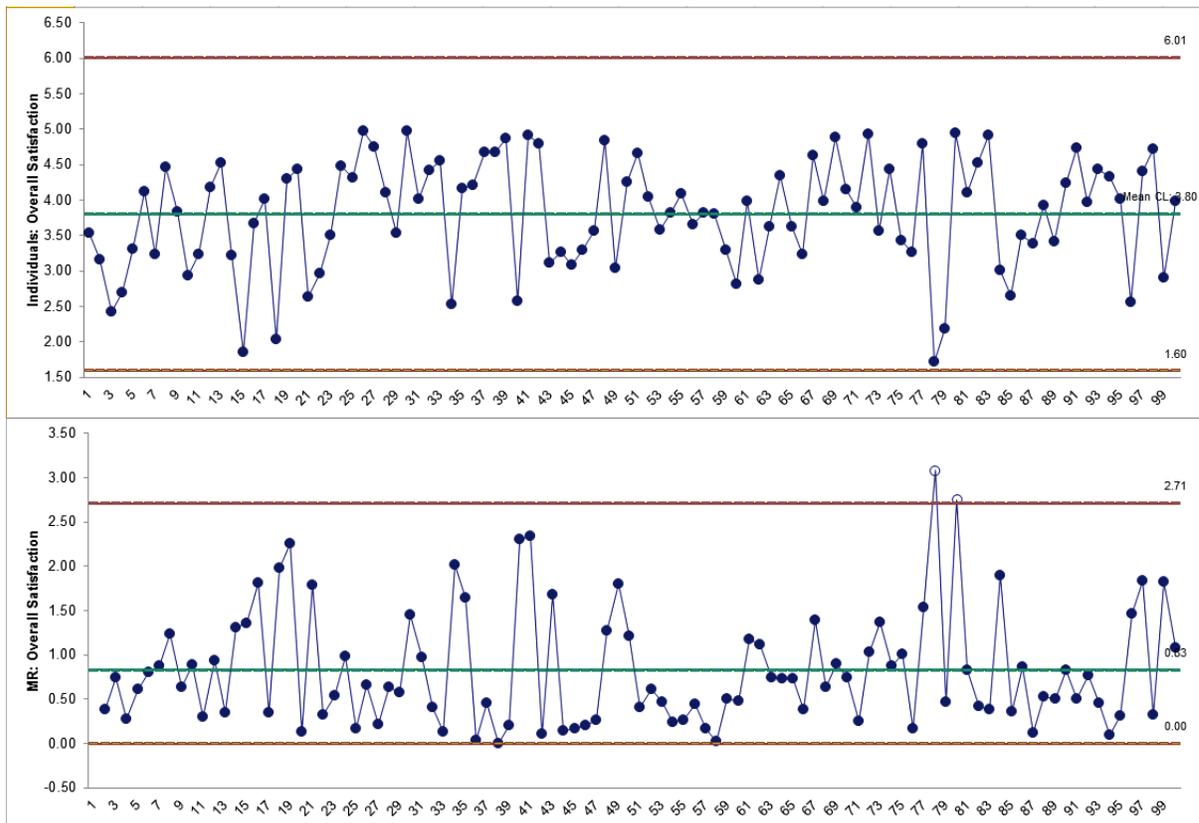


Individuals & Moving Range Charts

1. Open **Customer Data.xlsx**, click on **Sheet 1**. Click **SigmaXL>Control Charts>Individuals & Moving Range**. Ensure that entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
2. Select **Overall Satisfaction**, click **Numeric Data Variable (Y) >>**. Ensure that **Calculate Limits** is selected.

The screenshot shows the 'Individuals and Moving Range Chart' dialog box. On the left, a list of variables includes 'Overall Satisfaction', which is selected. The 'Numeric Data Variable (Y) >>' field contains 'Overall Satisfaction'. The 'Optional X-Axis Labels >>' field is empty. The 'Calculate Limits' radio button is selected, and 'Tests for Special Causes' and 'Sigma Zone Lines' are unchecked. The 'Advanced Options' button is visible. The 'Individuals' and 'Moving Range' sub-dialogs are also shown, with 'UCL', 'CL', and 'LCL' fields for both.

3. Click **OK**. Resulting Individuals & Moving Range control chart:



Individuals Charts for Nonnormal Data (Box-Cox Transformation)

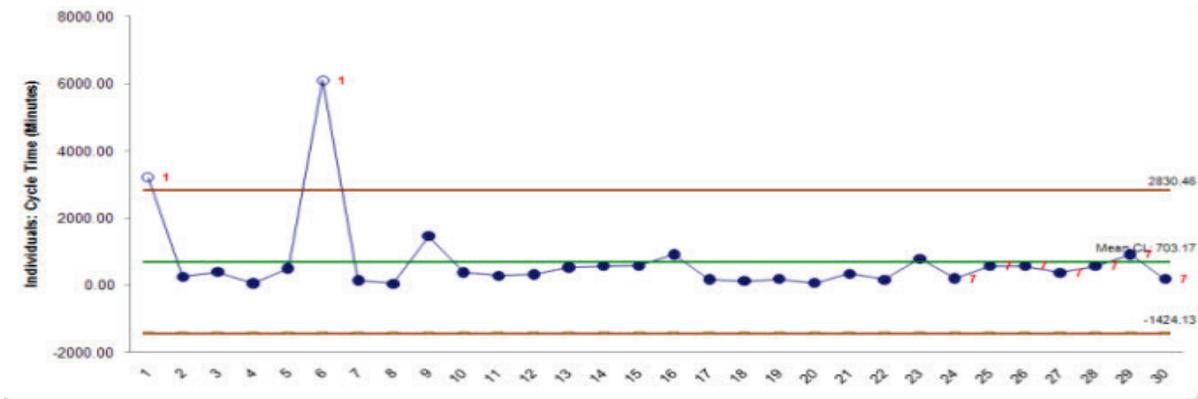
An important assumption for Individuals Charts is that the data be normally distributed (unlike the X-Bar Chart which is robust to nonnormality due to the central limit theorem). If the data is nonnormal, the Box-Cox Transformation tool can be used to convert nonnormal data to normal by applying a power transformation. The Johnson transformation and other distributions may also be used with Automatic Best Fit (see **Measure Phase Tools, Part J – Process Capability, Capability Combination Report Individuals Nonnormal**).

1. Open the file **Nonnormal Cycle Time2.xlsx**. This contains continuous nonnormal data of process cycle times. We performed a Process Capability study with this data earlier in the Measure Phase, Part H.
2. Initially, we will ignore the nonnormality in the data and construct an Individuals Chart. Click **SigmaXL > Control Charts > Individuals**. Ensure that entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
3. Select *Cycle Time (Minutes)*, click **Numeric Data Variable (Y) >>**. Select **Calculate Limits**. Check **Tests for Special Causes**. Click **OK**.

The screenshot shows the 'Individuals Chart' dialog box with the following settings:

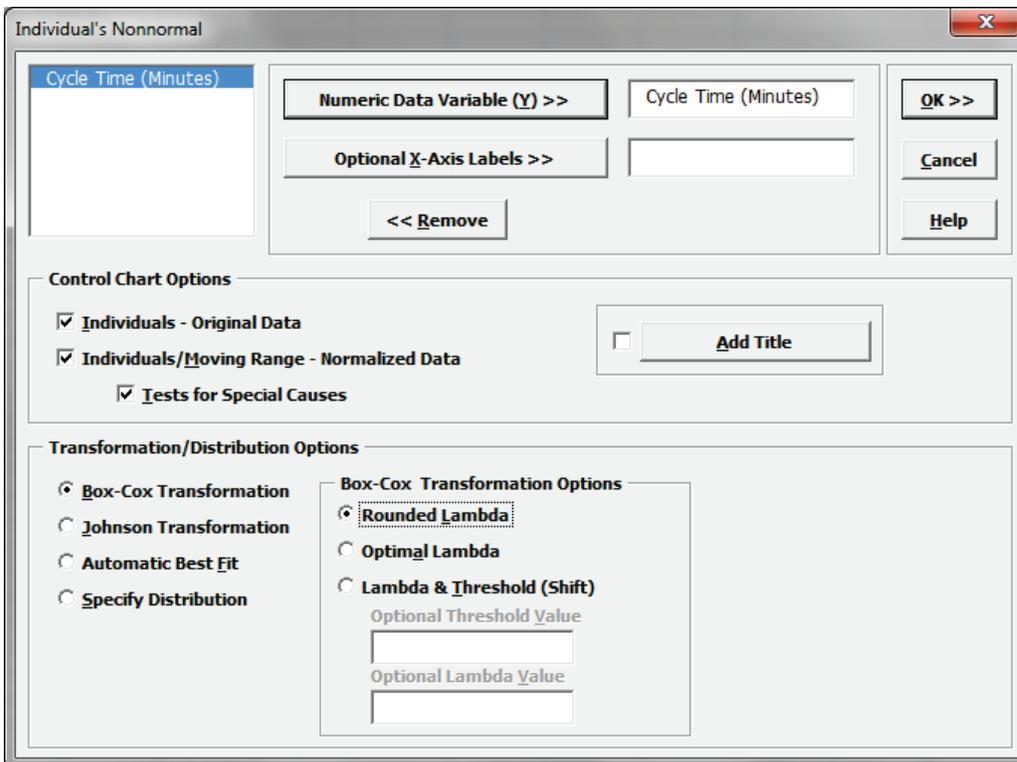
- Numeric Data Variable (Y) >>**: Cycle Time (Minutes)
- Optional X-Axis Labels >>**: (Empty)
- << Remove**: (Button)
- Calculate Limits**: (Selected)
- Historical Limits**: (Unselected)
- Tests for Special Causes**: (Checked)
- Sigma Zone Lines**: (Unselected)
- Advanced Options**: (Button)
- Add Title**: (Unselected)
- UCL**: (Empty text box)
- CL**: (Empty text box)
- LCL**: (Empty text box)
- OK >>**: (Button)
- Cancel**: (Button)
- Help**: (Button)

4. The resulting Individuals Chart is shown:

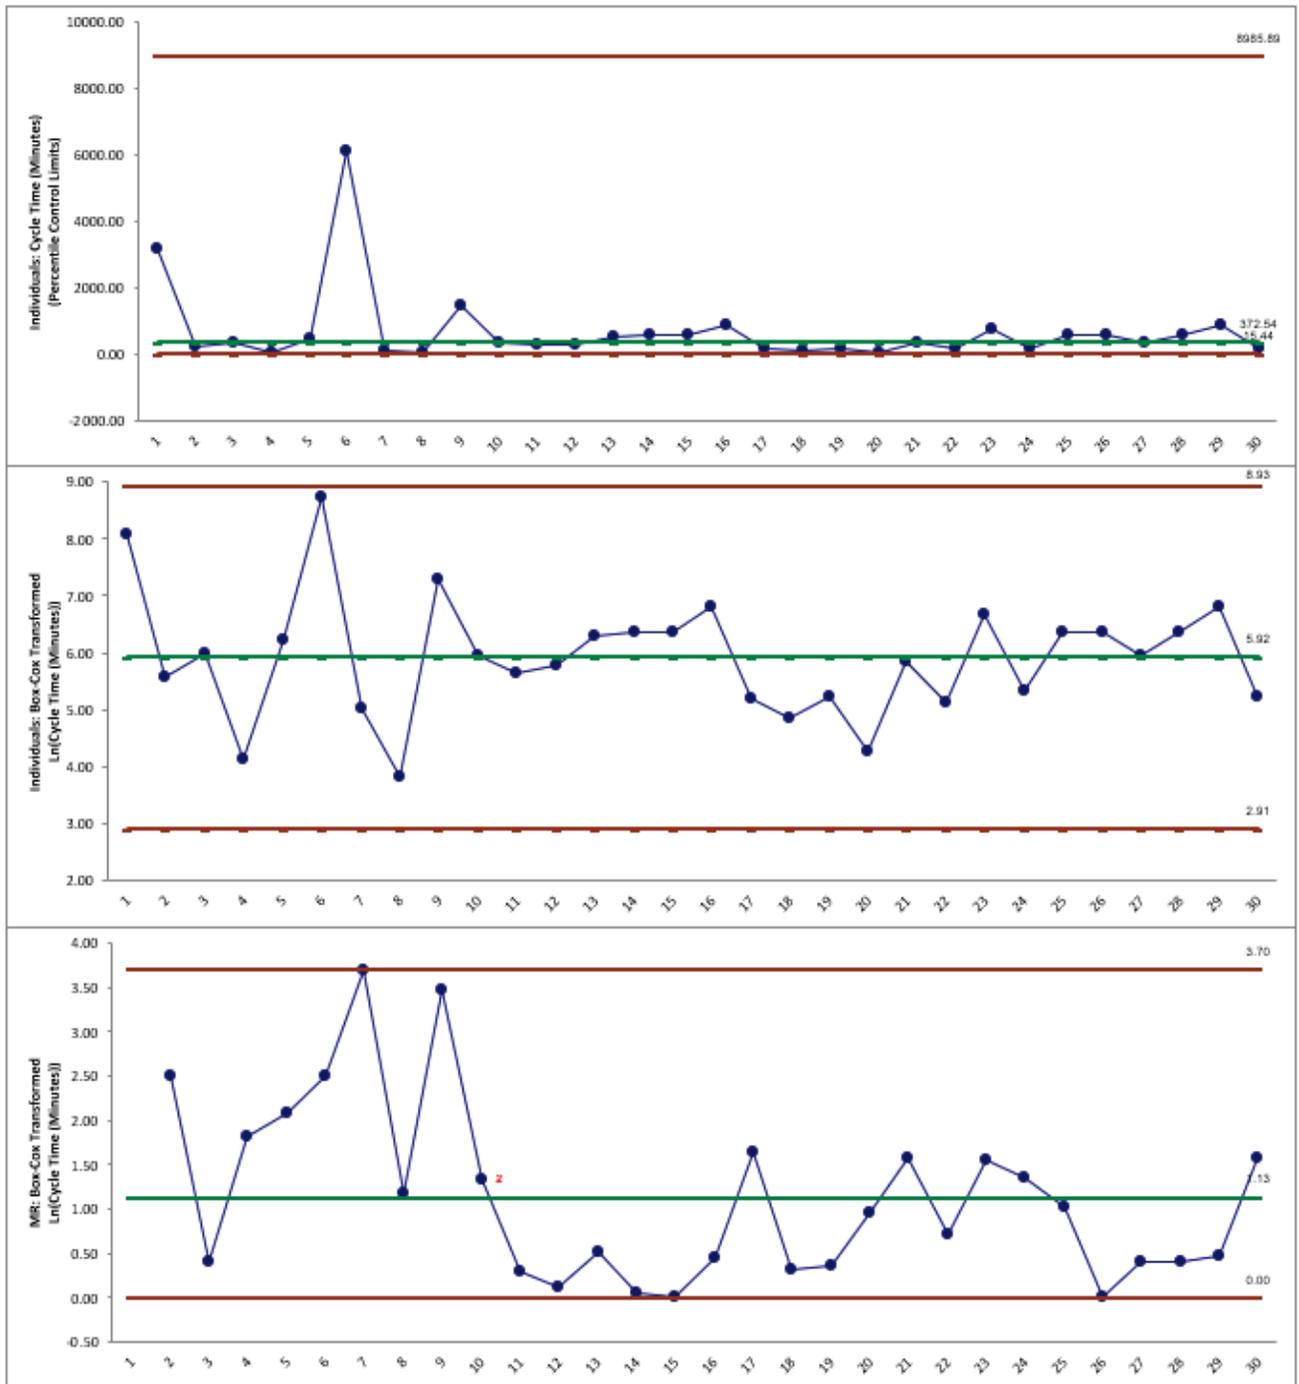


This chart clearly shows that the process is “out-of-control”. But is it really? Nonnormality can cause serious errors in the calculation of Individuals Chart control limits, triggering false alarms (Type I errors) or misses (Type II errors).

5. We will now construct Individuals Control Charts for nonnormal data. Select **Sheet 1** Tab (or press **F4**). Click **SigmaXL > Control Charts > Nonnormal > Individuals Nonnormal**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
6. Select *Cycle Time (Minutes)*, click **Numeric Data Variable (Y) >>**. We will use the default selection for **Transformation/Distribution Options: Box-Cox Transformation with Rounded Lambda**. Check **Tests for Special Causes** as shown:



7. Click **OK**. The resulting control charts are shown below:



Tests for Special Causes - Individuals: Cycle Time (Minutes)
 Number of Data Points Failing Tests = 0

Note that there are no out-of-control signals on the control charts, so the signals observed earlier when normality was assumed, were false alarms.

The **Individuals – Original Data** chart displays the untransformed data with control limits calculated as:

UCL = 99.865 percentile

CL = 50th percentile

LCL = 0.135 percentile

The benefit of displaying this chart is that one can observe the original untransformed data. Since the control limits are based on percentiles, this represents the overall, long term variation rather than the typical short term variation. The limits will likely be nonsymmetrical.

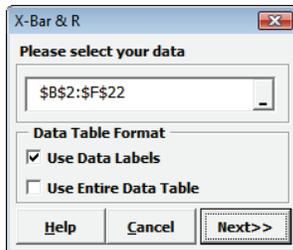
The **Individuals/Moving Range – Normalized Data** chart displays the transformed z-values with control limits calculated using the standard Shewhart formulas for Individuals and Moving Range charts. The benefit of using this chart is that tests for special causes can be applied and the control limits are based on short term variation. The disadvantage is that one is observing transformed data on the chart rather than the original data.

Part B - X-Bar & Range/StDev Charts

Tip: See [Part E – Control Chart Selection Tool](#). The Control Chart Selection Tool makes it easy for you to select the correct statistical process control chart depending on data type and subgroup/sample type and size.

X-Bar & R Charts

1. Open the file **Catapult Data – Xbar Control Charts.xlsx**. Each operator fires the ball 3 times. The target distance is 100 inches. The Upper Specification Limit (USL) is 108 inches. The Lower Specification Limit (LSL) is 92 inches.
2. Select **B2:F22** (if not already selected); here, we will only use the first 20 subgroups to determine the control limits.
3. Select **SigmaXL > Control Charts > X-Bar & R**. Do not check **Use Entire Data Table!**



4. Click **Next**. Select **Subgroups across Rows**, select *Shot 1, Shot 2, Shot 3*, click **Numeric Data Variables (Y) >>**; select *Operator*, click **Optional X-Axis Labels >>**. Check **Tests for Special Causes** as shown:

X-Bar and Range Chart

Subgroup No
Operator
Shot 1
Shot 2
Shot 3

Stacked Column Format (1 Numeric Data Column & Subgroup Size or Column)
 Subgroups across Rows (2 or More Numeric Data Columns)

Numeric Data Variables (Y) >> Shot 1
Shot 2
Shot 3

Optional X-Axis Labels >> Operator

<< Remove

Calculate Limits
 Historical Limits

Tests for Special Causes

Sigma Zone Lines

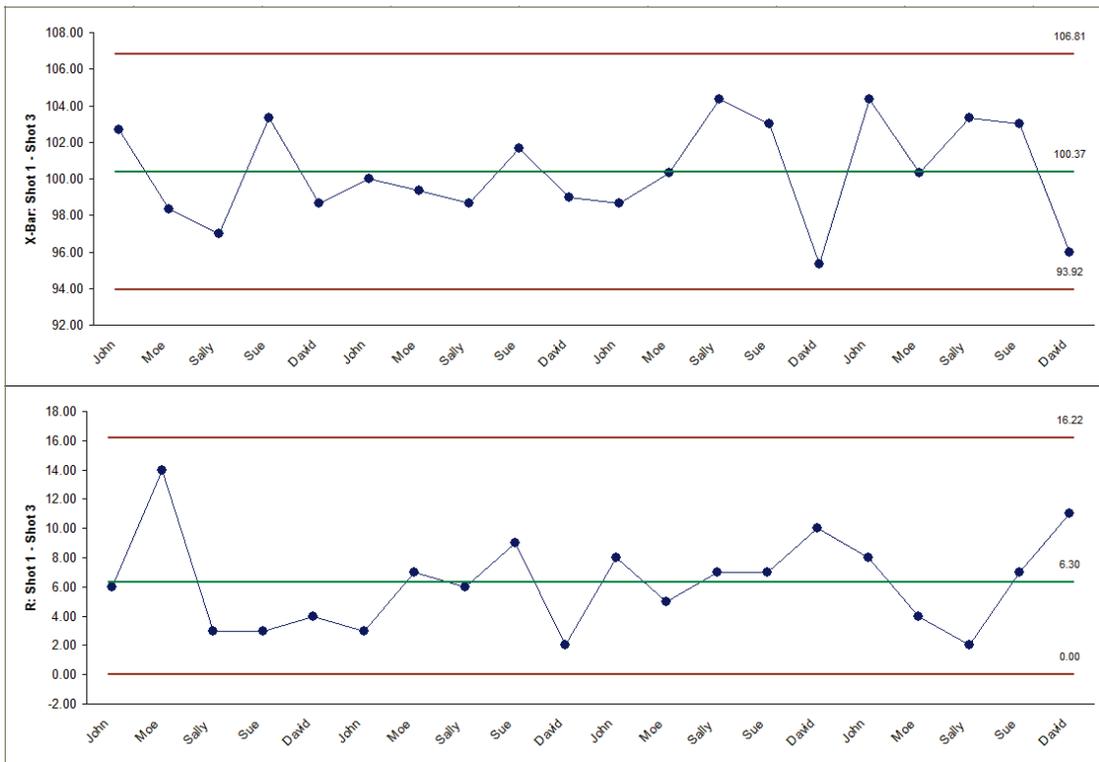
Advanced Options

Add Title

X-Bar		R
	UCL	
	CL	
	LCL	

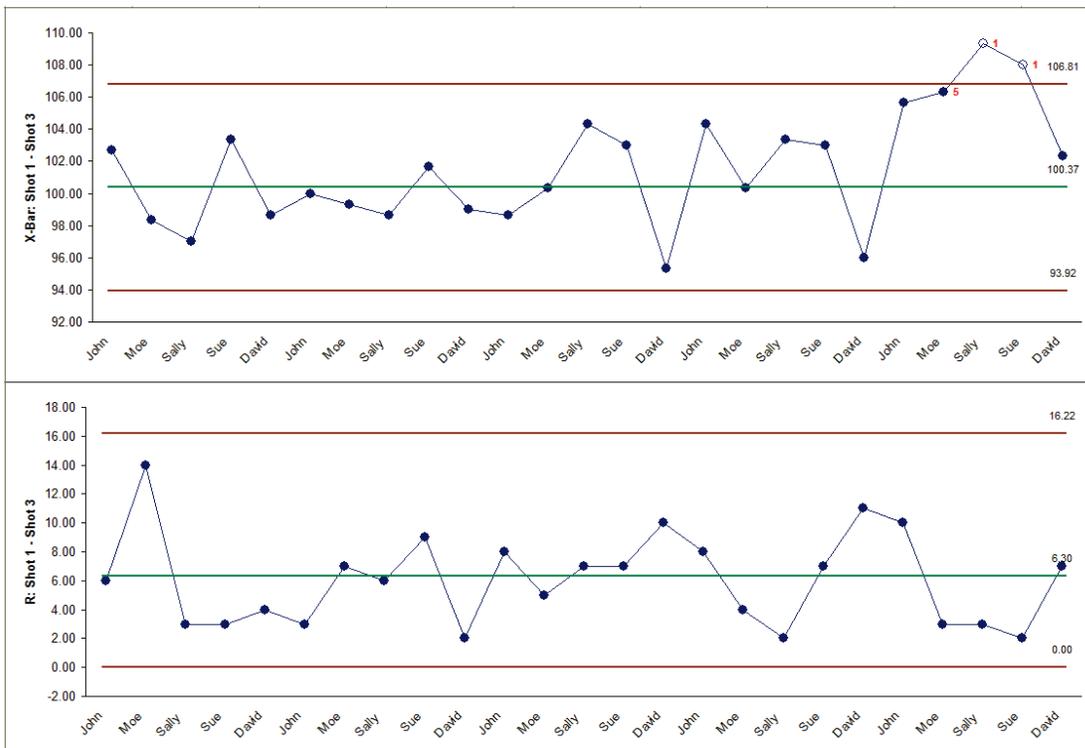
OK >>
Cancel
Help

5. Click **OK**. Resulting X-bar & R charts:



6. This is currently a stable catapult process. Subgroups 21 to 25 were added afterwards. To add the additional data to this chart, click **SigmaXL Chart Tools > Add Data to this Control Chart**

 [Add Data to this Control Chart](#) as shown:



- Note that the Add Data button does NOT recalculate the control limits. Once control limits are established, they should only be recalculated when a deliberate process change or improvement is introduced.
- The **Tests for Special Causes** report gives us more detail on the recent instability:

Tests for Special Causes - X-Bar: Shot 1 - Shot 3
Number of Data Points Failing Tests = 3

Observation No.	Test 1: 1 point more than 3 StDev from CL	Test 2: 9 points in a row on same side of CL	Test 3: 6 points in a row all increasing or all decreasing	Test 4: 14 points in a row alternating up and down	Test 5: 2 out of 3 points more than 2 StDev from CL (same side)	Test 6: 4 out of 5 points more than 1 StDev from CL (same side)	Test 7: 15 points in a row within 1 StDev from CL (either side)	Test 8: 8 points in a row more than 1 StDev from CL (either side)
22					x	x		
23	x				x	x		
24	x				x	x		

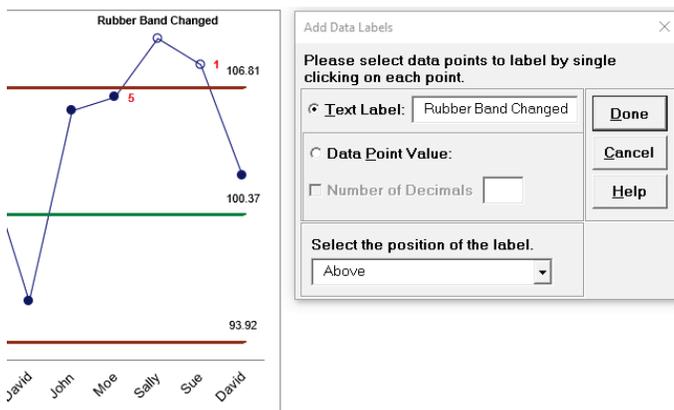
The X-bar chart and **Tests for Special Causes** report clearly shows that this process is now out of control with an unstable mean. The process must be stopped, and the Out-of-Control Action Plan must be followed to determine and fix the root cause. In this case, the assignable cause was a change of rubber band requiring a reset of the pull back angle. The use of tests for special causes gave us an early warning of this at observation number 22.

Note that the Range chart is in-control even though the X-Bar chart is out-of-control:

Tests for Special Causes - Range: Shot 1 - Shot 3
Number of Data Points Failing Tests = 0

Observation No.	Test 1: 1 point more than 3 StDev from CL	Test 2: 9 points in a row on same side of CL	Test 3: 6 points in a row all increasing or all decreasing	Test 4: 14 points in a row alternating up and down

- The tests for special causes can have defaults set to apply any or all of Tests 1-8. Test 2 can be set to 7, 8, or 9 points in a row on same side of CL. Click **SigmaXL > Control Charts > “Tests for Special Causes” Defaults** to run selected tests for special causes. Note that these defaults will apply to Individuals and X-bar charts. Test 1 to 4 settings will be applied to Attribute Charts and (if checked) Moving Range, Range, StDev. and Attribute Charts.
- To add a comment to a data point, select **SigmaXL Chart Tools > Add Data Label**. Select **Text Label**. Enter a comment as shown. Click on the data point to add the comment. Click **Done**.



11. Now we will look at Process Capability Indices for this process. Click on **Sheet 1** (or press **F4** to activate last worksheet). Click **SigmaXL > Control Charts > X-Bar & R**. Check **Use Entire Data Table**. Click **Next**. (Alternatively select **B2:F27**, press **F3**.)
12. Select *Shots 1-3*, click **Numeric Data Variables (Y) >>**.
13. Click **Advanced Options**. Enter **LSL = 92, Target = 100, USL = 108**.
14. The resulting dialog box settings are shown:

The screenshot shows the 'X-Bar and Range Chart' dialog box with the following settings:

- Format:** Subgroups across Rows (2 or More Numeric Data Columns)
- Numeric Data Variables (Y) >>:** Shot 1, Shot 2, Shot 3
- Optional X-Axis Labels >>:** (Empty)
- Calculate Limits:** Calculate Limits
- Historical Limits:** Historical Limits
- Tests for Special Causes:** Tests for Special Causes
- Sigma Zone Lines:** Sigma Zone Lines
- Advanced Options:** Add Title
- Process Capability Report:** LSL: 92, Target: 100, USL: 108
- Calculate Control Limits:** Specify Historical Group Column (Split Limits)
- Historical Groups >>:** (Empty)

On the right side of the dialog, there are input fields for X-Bar and R control limits:

	X-Bar		R
UCL			
CL			
LCL			

15. Click **OK**. Click **X-Bar & R – Proc Cap** sheet for the Process Capability report:

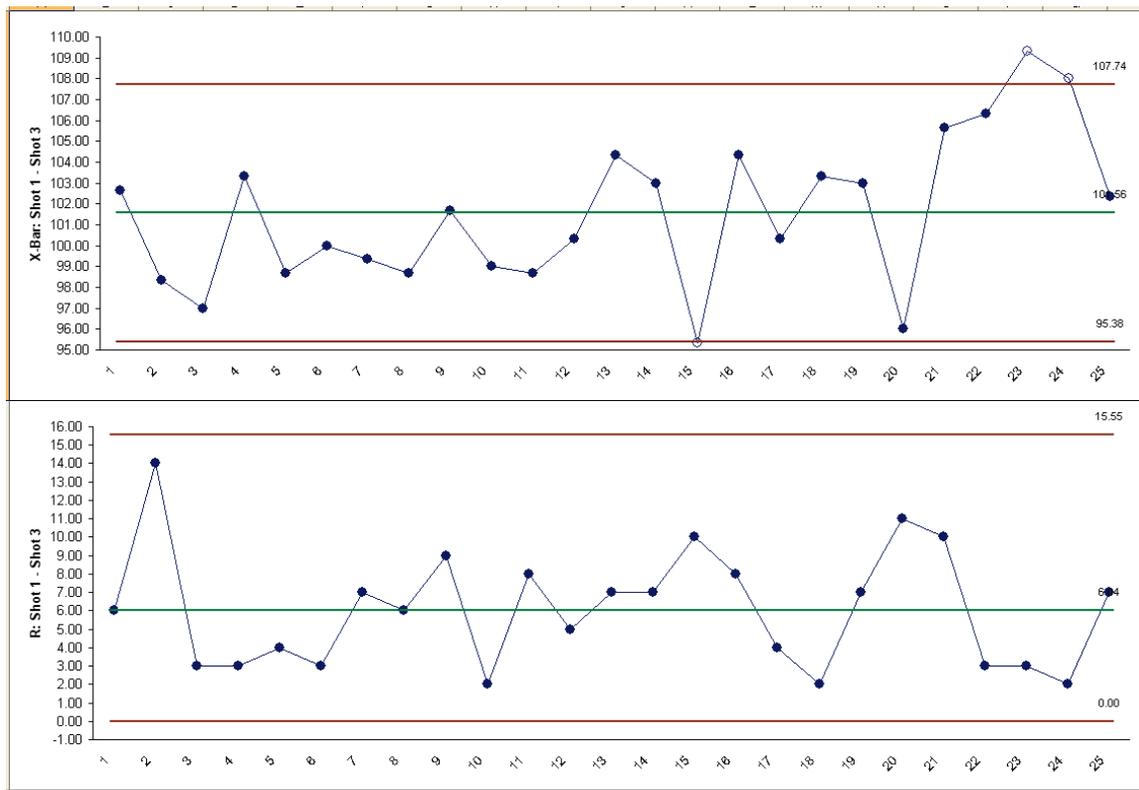
Process Capability Report: X-Bar: Shot 1 - Shot 3	
Count	75
Mean	101.56
StDev (Overall, Long Term)	4.616
StDev (Within, Short Term)	3.568
USL	108
Target	100
LSL	92
Capability Indices using Overall StDev	
Pp	0.58
Ppu	0.47
Ppl	0.69
Ppk	0.47
Cpm	0.55
Potential Capability Indices using Within StDev	
Cp	0.75
Cpu	0.60
Cpl	0.89
Cpk	0.60
Expected Overall Performance	
ppm > USL	81468
ppm < LSL	19168.4
ppm Total	100636.0
% > USL	8.15%
% < LSL	1.92%
% Total	10.06%
Actual (Empirical) Performance	
% > USL	5.33%
% < LSL	4.00%
% Total	9.33%
Anderson-Darling Normality Test	
A-Squared	0.486247
P-Value	0.2192

Note the difference between Pp and Cp; Ppk and Cpk. This is due to the process instability. If the process was stable, the actual performance indices Pp and Ppk would be closer to the Cp and Cpk values.

X-Bar & R Charts – Exclude Subgroups

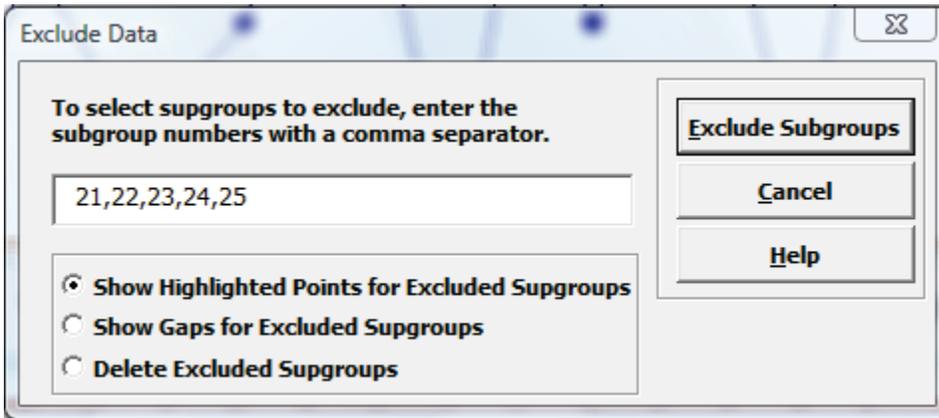
After creating a control chart, you can specify subgroups (or rows) to exclude by using the **Exclude Data** tool.

16. Click on **Sheet 1** (or press **F4** to activate last worksheet). Click **SigmaXL > Control Charts > X-Bar & R**. Check **Use Entire Data Table**. Click **Next**.
17. Select *Shots 1-3*, click **Numeric Data Variables (Y) >>**. Ensure that **Calculate Limits** is selected. Click **OK**.
18. The resulting X-bar & R charts are displayed:

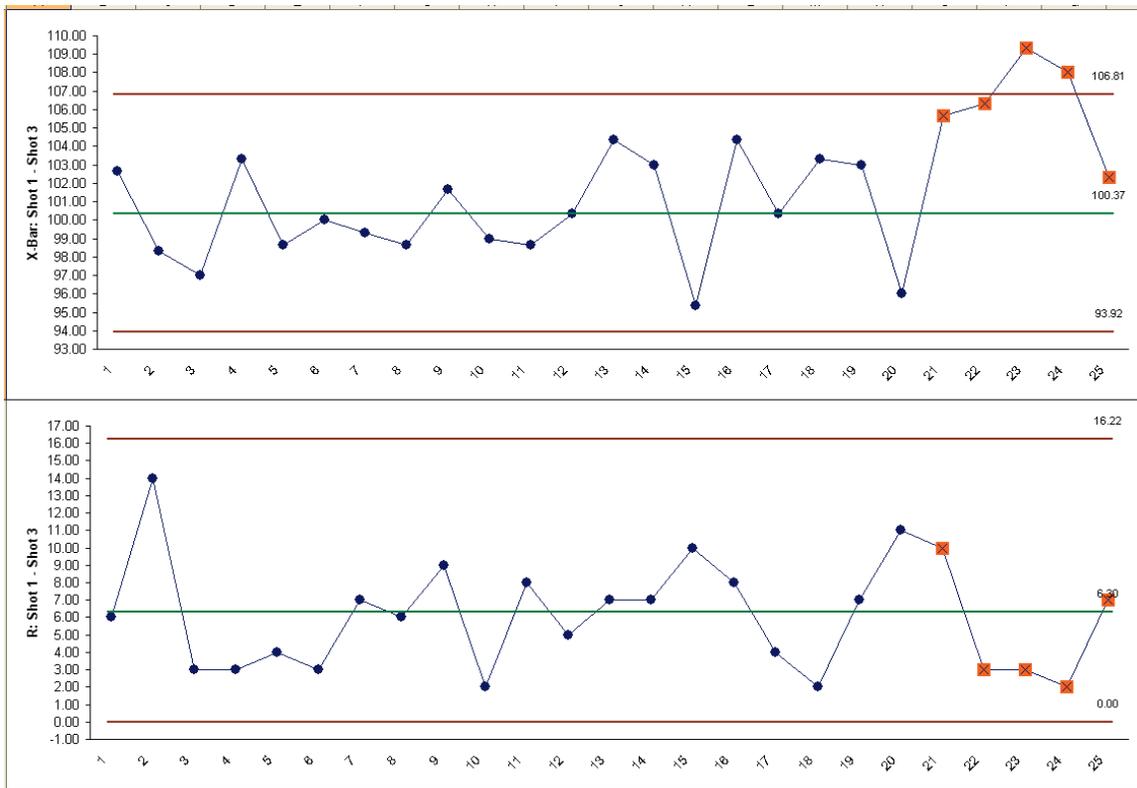


The control limits here were calculated including subgroups 21 to 25 which have a known assignable cause.

19. To calculate the control limits excluding subgroups 21 to 25, click **SigmaXL Chart Tools > Exclude Subgroups**. Select **Show Highlighted Points for Excluded Subgroups**. Enter 21,22,23,24,25 as shown:



20. Click **Exclude Subgroups**. The control chart limits are recalculated and the excluded points are highlighted:



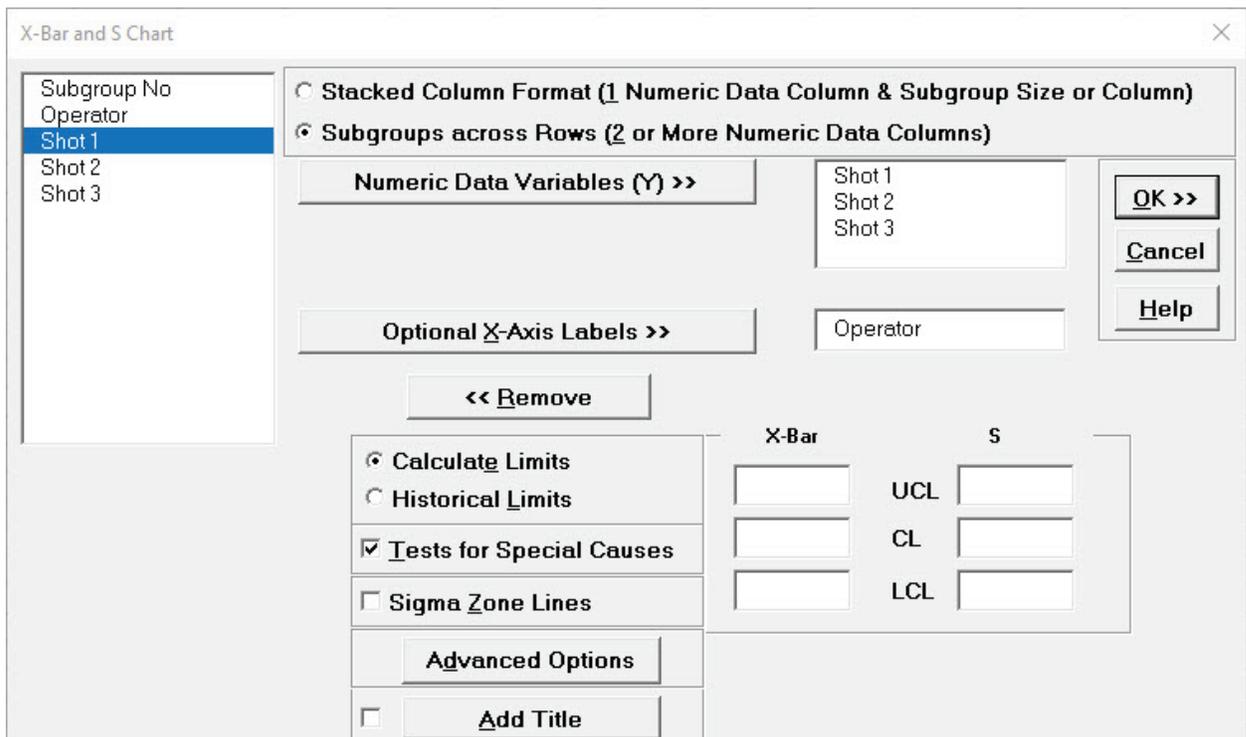
Tip: You can also choose to show gaps for excluded subgroups or delete excluded subgroups from the charts.

X-Bar & S Charts

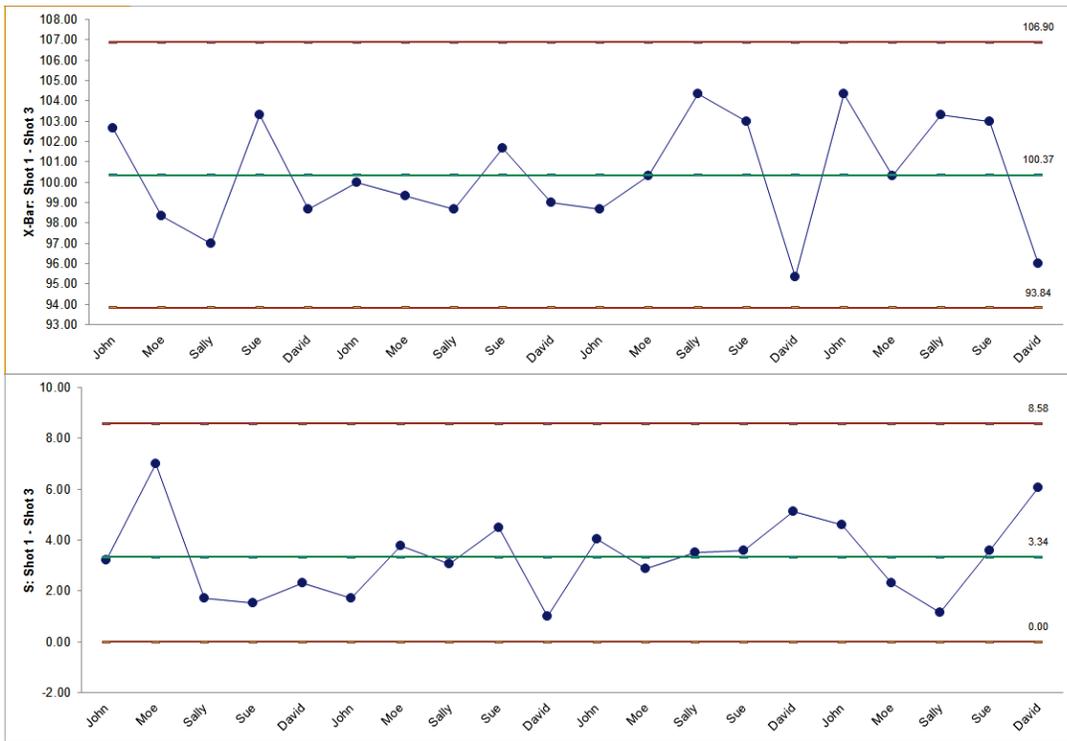
1. Open the file **Catapult Data – Xbar Control Charts.xlsx**. Each operator fires the ball 3 times. The target distance is 100 inches. The Upper Specification Limit (USL) is 108 inches. The Lower Specification Limit (LSL) is 92 inches.
2. Select **B2:F22**; here, we will only use the first 20 subgroups to determine the control limits.
3. Select **SigmaXL > Control Charts > X-Bar & S**.
4. Do not check **Use Entire Data Table**!



5. Click **Next**. Select **Subgroups across Rows**, select *Shot 1, Shot 2, Shot 3*, click **Numeric Data Variables (Y) >>**; select *Operator*, click **Optional X-Axis Labels >>**. Check **Tests for Special Causes** as shown:

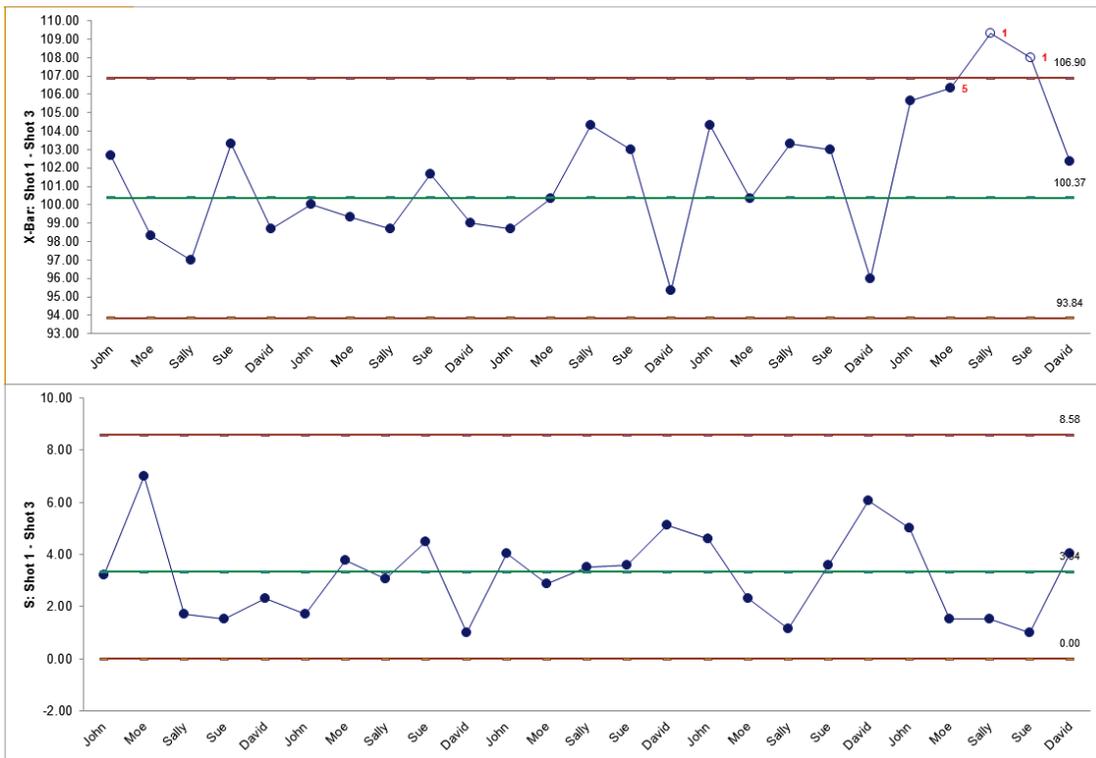


6. Click **OK**. Resulting X-bar & S charts:



7. This is currently a stable catapult process. Subgroups 21 to 25 were added afterwards. To add the additional data to this chart, click **SigmaXL Chart Tools > Add Data to this Control Chart**

 **Add Data to this Control Chart** as shown:



Part C – Attribute Charts (P, NP, C, U)

Tip: See [Part E – Control Chart Selection Tool](#). The Control Chart Selection Tool makes it easy for you to select the correct statistical process control chart depending on data type and subgroup/sample type and size.

P-Charts

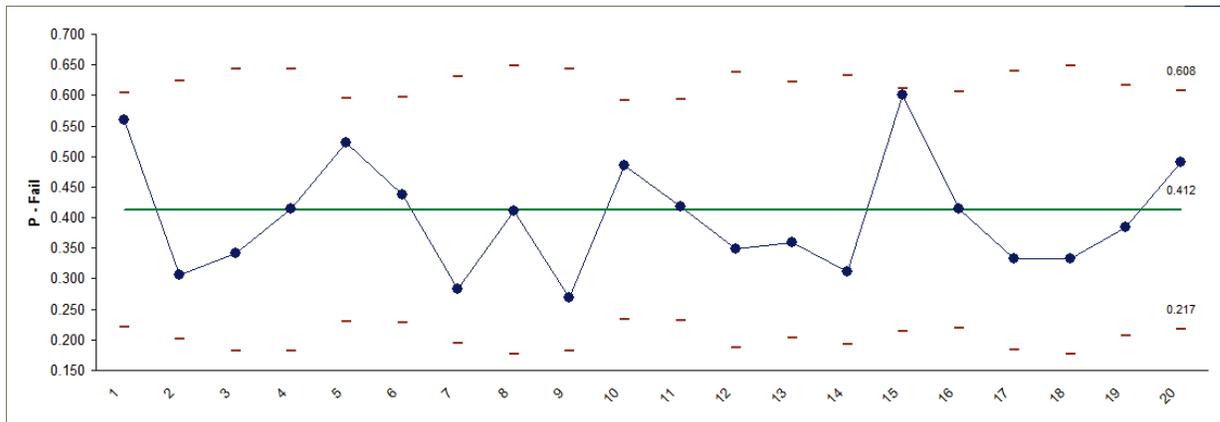
1. Open **New York Daily Cycle Time – Discrete.xlsx**. This is data from the Sigma Savings and Loans Company, New York location. Each day, the cycle time (in days) for completed loans and leases was recorded. *N* indicates the number of loans counted. A *Fail* was recorded if the cycle time exceeded the critical customer requirement of 8 days. Note that we are not recommending that continuous data be converted to discrete data in this manner, but rather using this data to illustrate the use of P charts for Discrete or Attribute data. P Charts (for Defectives) can have fixed or varying subgroup sizes.
2. Select **SigmaXL > Control Charts > Attribute Charts > P**. Ensure that **B3:E23** are selected, click **Next**.
3. Select *Fail*, click **Numeric Data Variable (Y) >>**; select *N*, click **Subgroup Column or Size >>**. If we had a fixed subgroup size, the numerical value of the subgroup size could be entered instead of Column N. Check **Tests for Special Causes**.

The screenshot shows the 'P-Chart' dialog box. On the left, a list of variables includes 'Day', 'Fail', 'Pass', and 'N', with 'N' selected. The main area contains the following fields and options:

- Numeric Data Variable (Y) >>**: Fail
- Subgroup Column or Size >>**: N
- Optional X-Axis Labels >>**: (empty)
- << Remove**: (button)
- Calculate Limits**: Calculate Limits, Historical Limits
- Tests for Special Causes**:
- Sigma Zone Lines**:
- Advanced Options**: (button)
- Add Title**:
- UCL**: (input field)
- CL**: (input field)
- LCL**: (input field)

On the right side, there are buttons for **OK >>**, **Cancel**, and **Help**.

4. Click **OK**. The resulting P-Chart is shown:



The moving limits are due to the varying subgroup sizes. While this P-chart shows stability, a much bigger concern is the average 41% failure rate to deliver the loans/leases in 8 days or less!

P-Charts: Advanced Limit Options – Historical Groups

1. Open **New York Daily Cycle Time Discrete – Before After Improvement.xlsx**. *Before_After* denotes before improvement and after improvement. We would like to use a P-Chart with historical groups to split the limits demonstrating the before versus after improvement.
2. Click **SigmaXL > Control Charts > Attribute Charts > P**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
3. Select *Fail*, click **Numeric Data Variable (Y) >>**; select *N*, click **Subgroup Column or Size >>**.
4. Click **Advanced Options**. Select **Specify Historical Group Column**. Select *Before_After*, click **Historical Groups >>**.

P-Chart [Close]

Day
Before_After
Fail
Pass
N

Numeric Data Variable (Y) >> Fail

Subgroup Column or Size >> N

Optional X-Axis Labels >>

<< Remove

Calculate Limits
 Historical Limits
 Tests for Special Causes
 Sigma Zone Lines

UCL
 CL
 LCL

Advanced Options

Add Title

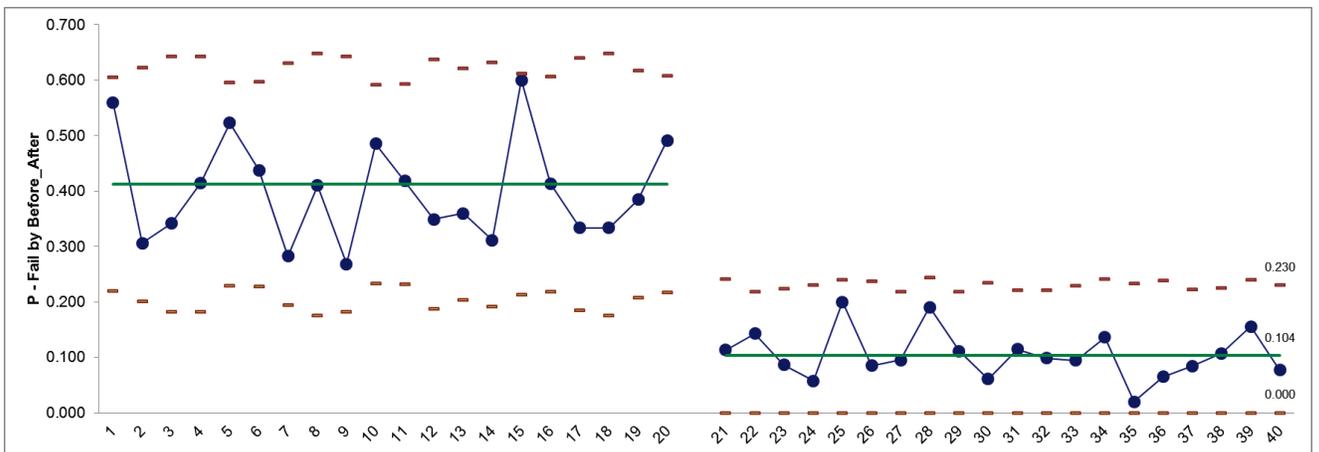
Calculate Control Limits:

Specify Subgroup Number for Calculation of Control Limits
 Specify Historical Group Column (Split Limits)

Day
Before_After
 Fail
 Pass

Historical Groups >> Before_After

- Click **OK**. The resulting P-Chart with split limits based on historical groups is shown, demonstrating a clear process improvement:

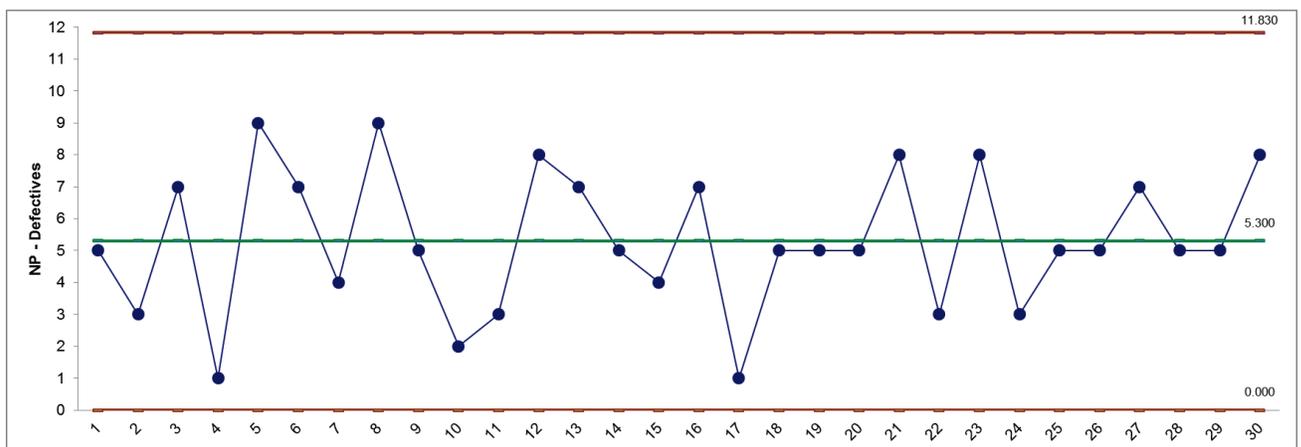


NP-Charts

1. Open **Attribute Data NP Chart Defectives.xlsx**. The subgroup size is constant at 50. NP Charts (for Defectives) require a fixed subgroup size.
2. Select **SigmaXL > Control Charts > Attribute Charts > NP**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
3. Select *Defectives*, click **Numeric Data Variable (Y) >>**; enter 50 for **Enter Subgroup Size**. Check **Tests for Special Causes**.

The screenshot shows the 'NP-Chart' dialog box. On the left, a list contains 'Defectives' and 'N', with 'Defectives' selected. In the center, there are three main buttons: 'Numeric Data Variable (Y) >>', 'Enter Subgroup Size', and 'Optional X-Axis Labels >>'. The 'Enter Subgroup Size' field is set to '50'. Below these are 'Advanced Options' and 'Add Title' buttons. On the right, there are three input fields for 'UCL', 'CL', and 'LCL', all of which are empty. At the bottom right, there are 'OK >>', 'Cancel', and 'Help' buttons. The 'Tests for Special Causes' checkbox is checked, while 'Calculate Limits', 'Historical Limits', and 'Sigma Zone Lines' are unchecked.

4. Click **OK**. The resulting NP-Chart is shown:



This is a stable process, with no tests for special causes flagged.

C Chart Template

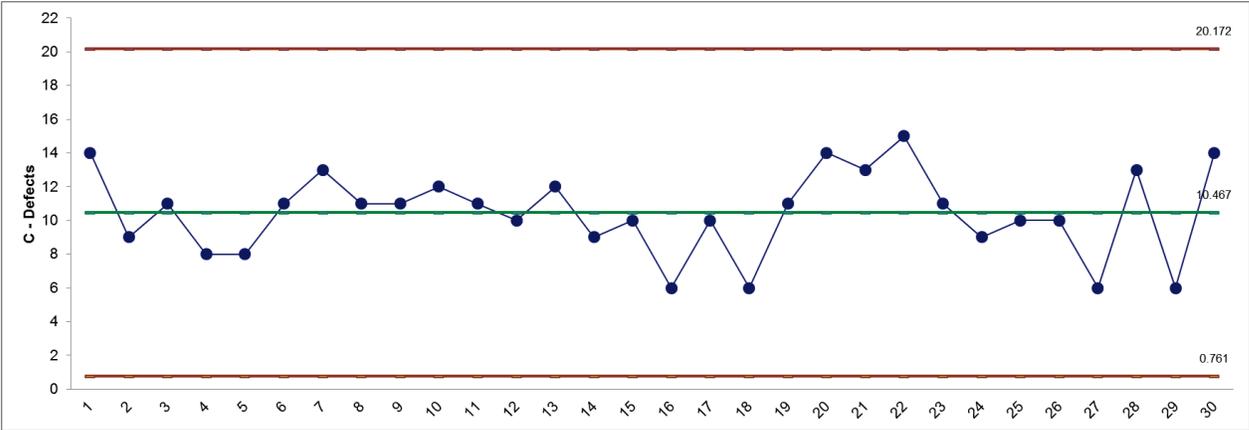
Click **SigmaXL > Templates and Calculators > Control Chart Templates > Basic > C (Count)**.

This template is also located at **SigmaXL > Control Charts > Control Chart Templates > Basic > C (Count)**.

C-Charts

1. Open **Attribute Data C Chart Defects.xlsx**. C Charts (for Defects) assume a fixed subgroup size.
2. Select **SigmaXL > Control Charts > Attribute Charts > C**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
3. Select *Defects*, click **Numeric Data Variable (Y) >>**. Check **Tests for Special Causes**.

4. Click **OK**. The resulting C-Chart is shown:



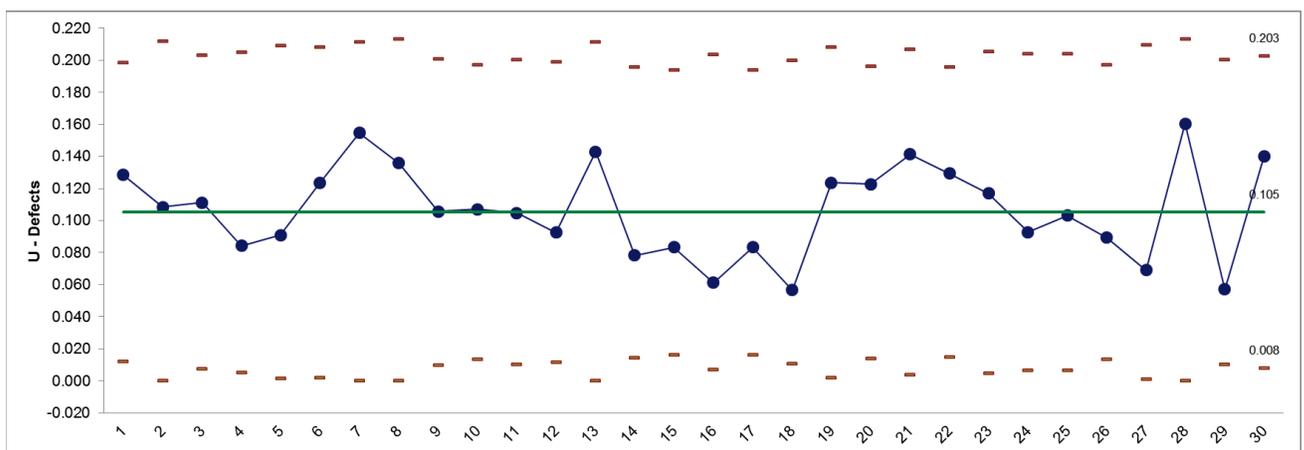
This is a stable process, with no tests for special causes flagged.

U-Charts

1. Open **Attribute Data U Chart Defects.xlsx**. U Charts (for Defects) can have fixed or varying subgroup sizes.
2. Select **SigmaXL > Control Charts > Attribute Charts > U**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
3. Select *Defects*, click **Numeric Data Variable (Y) >>**. Select *N*, click **Subgroup Column or Size >>**. Check **Tests for Special Causes**.

The screenshot shows the 'U-Chart' dialog box. On the left, a list contains 'Defects' and 'N', with 'N' selected. In the center, there are three buttons: 'Numeric Data Variable (Y) >>', 'Subgroup Column or Size >>', and 'Optional X-Axis Labels >>'. Below these is a '<< Remove' button. To the right, three input fields contain 'Defects', 'N', and are empty. Further right are 'OK >>', 'Cancel', and 'Help' buttons. Below the input fields, there are radio buttons for 'Calculate Limits' (selected) and 'Historical Limits'. A checked checkbox is labeled 'Tests for Special Causes'. Below that is an unchecked checkbox for 'Sigma Zone Lines'. To the right of these are three empty input fields labeled 'UCL', 'CL', and 'LCL'. At the bottom, there is an 'Advanced Options' button and an unchecked 'Add Title' checkbox.

4. Click **OK**. The resulting U-Chart is shown:



This is a stable process, with no tests for special causes flagged.

Part D – P' & U' Charts (Laney)

P' and U' (Laney) Control Charts are attribute control charts that should be used when the subgroup/sample size is large and assumptions are not met. Typically, you will see that the control limits do not “look right,” being very tight with many data points appearing to be out-of-control. This problem is also referred to as *overdispersion*. This occurs when the assumption of a Binomial distribution for defectives or Poisson distribution for defects is not valid. Individuals charts are often recommended in these cases, but Laney’s P' and U' charts are a preferred alternative.

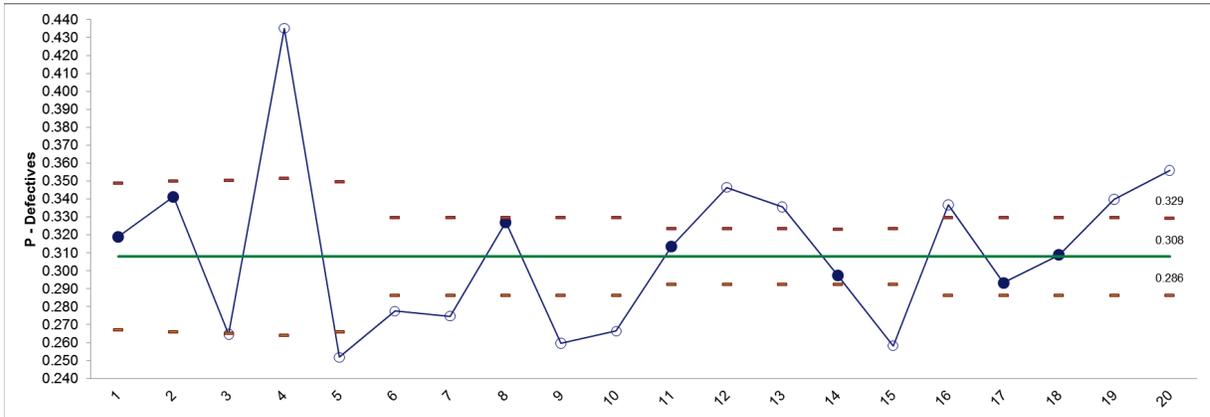
Tip: See [Part E – Control Chart Selection Tool](#). The Control Chart Selection Tool makes it easy for you to select the correct statistical process control chart depending on data type and subgroup/sample type and size.

References:

- [1] Crossley, Mark L.(2007), *The Desk Reference of Statistical Quality Methods*, Second Edition, Milwaukee, WI, ASQ Quality Press, pp. 352 – 355.
- [2] Laney, David B., “P-Charts and U-Charts Work (But Only Sometimes),” *Quality Digest*, http://www.qualitydigest.com/sept07/departments/what_works.shtml.
- [3] Laney, David B. (2002), “Improved Control Charts for Attribute Data,” *Quality Engineering* 14:531–7.
- [4] M. A. Mohammed and D. Laney (2006), “Overdispersion in health care performance data: Laney’s approach,” *Qual. Saf. Health Care* 15; 383-384.

P'-Charts

1. Open **Laney – Quality Digest – Defectives.xlsx**. This data is used with permission from David Laney and was used in the example given in reference [2].
2. We will begin with the creation of a regular P-Chart for this data. Select **SigmaXL > Control Charts > Attribute Charts > P**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
3. Select *Defectives* as the **Numeric Data Variable (Y)**, *N* as the **Subgroup Column or (Size)**. Click **OK**. The resulting P-Chart is shown:



At first glance, this chart suggests that the process is out of control. The problem here is actually due to the Binomial assumption not being valid (also known as *overdispersion*). This problem becomes evident when the sample sizes are large.

4. Select Sheet **Defectives** (or press **F4**). Click **SigmaXL > Control Charts > Attribute Charts > P' (Laney)**. Click **Next**.
5. Select *Defectives*, click **Numeric Data Variable (Y) >>**; select *N*, click **Subgroup Column or Size >>**. (If we had a fixed subgroup size, the numerical value of the subgroup size could be entered instead of Column N.) Check **Tests for Special Causes**.

P' (Laney)

Defectives

N

Numeric Data Variable (Y) >> Defectives

Subgroup Column or Size >> N

Optional X-Axis Labels >>

<< Remove

Calculate Limits
 Historical Limits

Tests for Special Causes

Sigma Zone Lines

Add Title

UCL

CL

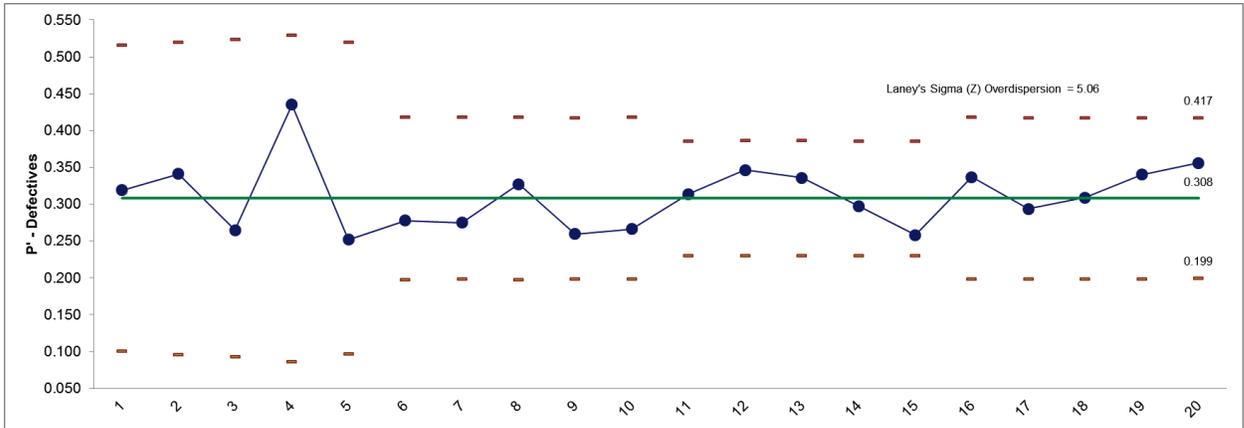
LCL

OK >>

Cancel

Help

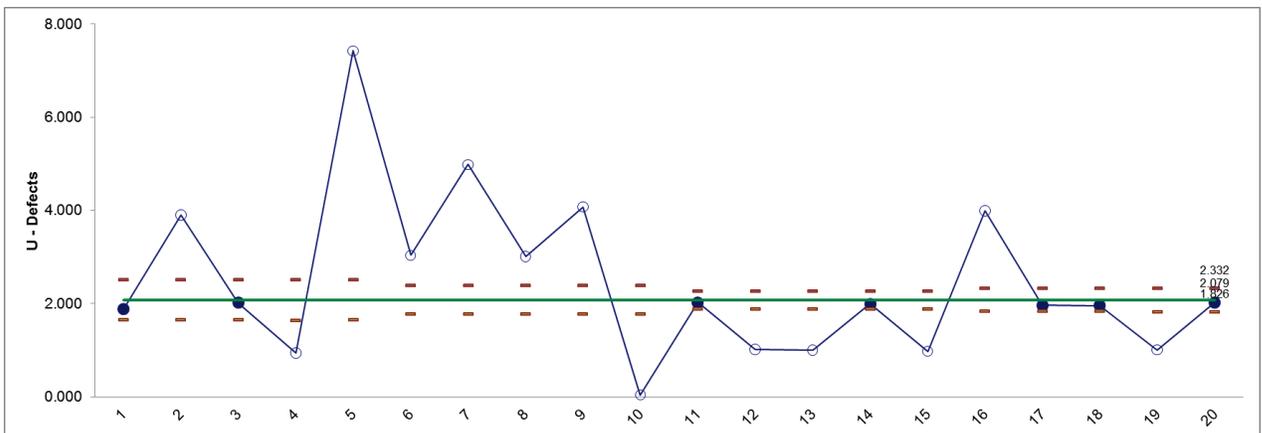
6. Click **OK**. The resulting P'-Chart is shown:



7. Now we see that the process is actually “in-control.” Laney’s Sigma (Z) is a measure of the *overdispersion*. See referenced articles for further details.

U'-Charts

1. Open **Attribute Data – U' Defects - Crossley.xlsx**. This data is used with permission from Mark Crossley and was used in the example given in reference [1]. Note that there are multiple defects per unit.
2. We will begin with the creation of a regular U-Chart for this data. Select **SigmaXL > Control Charts > Attribute Charts > U**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
3. Select *Defects* as the **Numeric Data Variable (Y)**, *N* as the **Subgroup Column (Size)**. Click **OK**. The resulting U-Chart is shown:



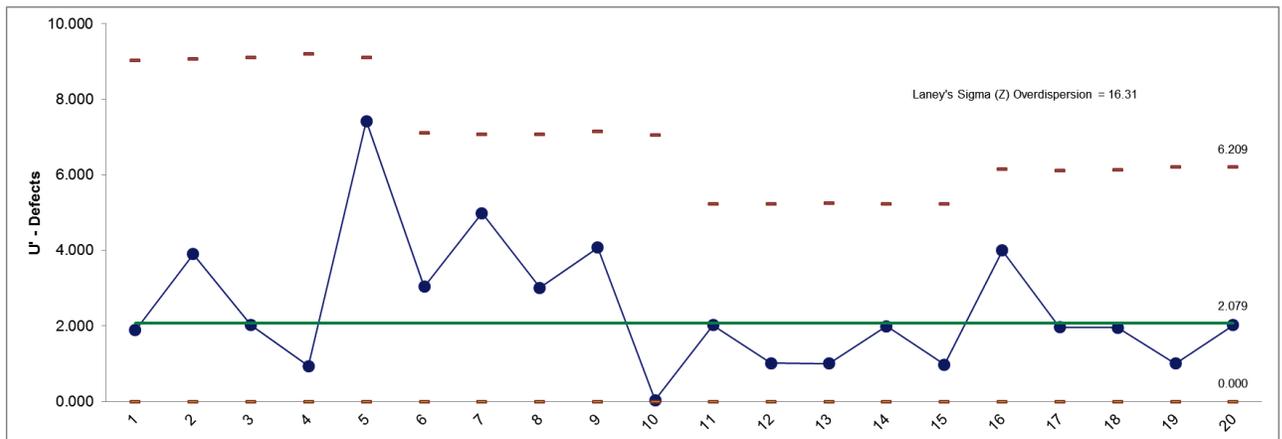
This chart suggests that the process is out of control, but the problem here is actually due to the Poisson assumption not being valid (also known as *overdispersion*). This problem becomes evident when the sample sizes are large.

4. Select Sheet **U' Chart Defects - Crossley** (or press **F4**). Click **SigmaXL > Control Charts > Attribute Charts > U' (Laney)**. Click **Next**.

- Select *Defects*, click **Numeric Data Variable (Y) >>**; select *N*, click **Subgroup Column or Size >>**. (If we had a fixed subgroup size, the numerical value of the subgroup size could be entered instead of Column N.) Check **Tests for Special Causes**.

The screenshot shows the 'U' (Laney) dialog box. On the left, a list contains 'Defects' and 'N', with 'N' selected. In the center, there are three buttons: 'Numeric Data Variable (Y) >>', 'Subgroup Column or Size >>', and 'Optional X-Axis Labels >>'. Below these is a '<< Remove' button. On the right, there are three input fields: the first contains 'Defects', the second contains 'N', and the third is empty. Further right are 'OK >>', 'Cancel', and 'Help' buttons. Below the main controls, there are radio buttons for 'Calculate Limits' (selected) and 'Historical Limits'. A checked checkbox is labeled 'Tests for Special Causes'. There are also checkboxes for 'Sigma Zone Lines' and 'Add Title'. On the far right, there are three input fields for 'UCL', 'CL', and 'LCL', all of which are currently empty.

- Click **OK**. The resulting U'-Chart is shown:



- Now we see that the process is actually “in-control.” Laney’s Sigma (Z) is a measure of the *overdispersion*. See reference [1] for further details.

Part E – Control Chart Selection Tool

The Control Chart Selection Tool makes it easy for you to select the correct statistical process control chart depending on data type and subgroup/sample type and size.

Data Types and Definitions

Continuous/Variable: Data that is measured on a continuous scale where a mid-point (or other subdivision) has meaning. For example, when measuring cycle time, 2.5 days has meaning. Other examples include distance, weight, thickness, length and cost. Customer Satisfaction on a 1 to 5 scale can be considered as continuous in that a satisfaction score of 3.5 has meaning. Continuous data is always in numeric format.

Discrete/Attribute: Data that is categorical in nature. If we have defect types 1, 2, and 3, defect type 1.5 has no meaning. Other examples of discrete data would be customer complaints and reasons for product return. Discrete data can be text or integer numeric format.

Defective: An entire unit that is nonconforming to customer requirements. A unit may be defective because of one or more defects. For example, an application form is good only if all critical entry fields are correct. Any error in a critical field is a defect, resulting in a defective form. A single form can have more than one defect.

Defect: Any specific nonconformity to customer requirements. There can be more than one defect per unit or area of opportunity, such as the entry errors described above.

Subgroup/Sample: Data for a subgroup are usually collected within a short period of time to ensure homogeneous conditions within the subgroup (common cause variation), in order to detect differences between subgroups (special cause variation).

Subgroup/Sample Size: The number of observations within your sample, not the number of samples. Subgroup sizes of 3 to 5 are common for continuous measures in parts manufacturing, while individual measurements are common in chemical processes (temperature, pH) and transactional areas (financial). Subgroup size for discrete data should be a minimum of 50.

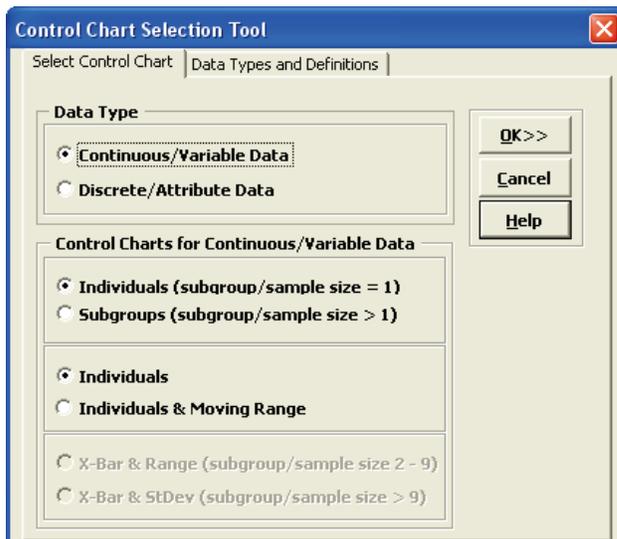
Subgroup/Sample Size is constant: The number of observations within your sample remains fixed over time.

Subgroup/Sample Size varies: The number of observations within your sample varies over time.

Subgroup/Sample Size is very large and assumptions not met: This applies to discrete data when the subgroup sizes are approximately 5,000 or higher and the control limits do not “look right,” being very tight with many data points appearing to be out-of-control. This problem is also referred to as “overdispersion.” This occurs when the assumption of a Binomial distribution for defectives or Poisson distribution for defects is not valid. (Note: If the problem of overdispersion is apparent with your continuous data, use **SigmaXL > Control Charts > Advanced Charts > I-MR-R or I-MR-S**).

Control Chart Selection Tool – Individuals Chart

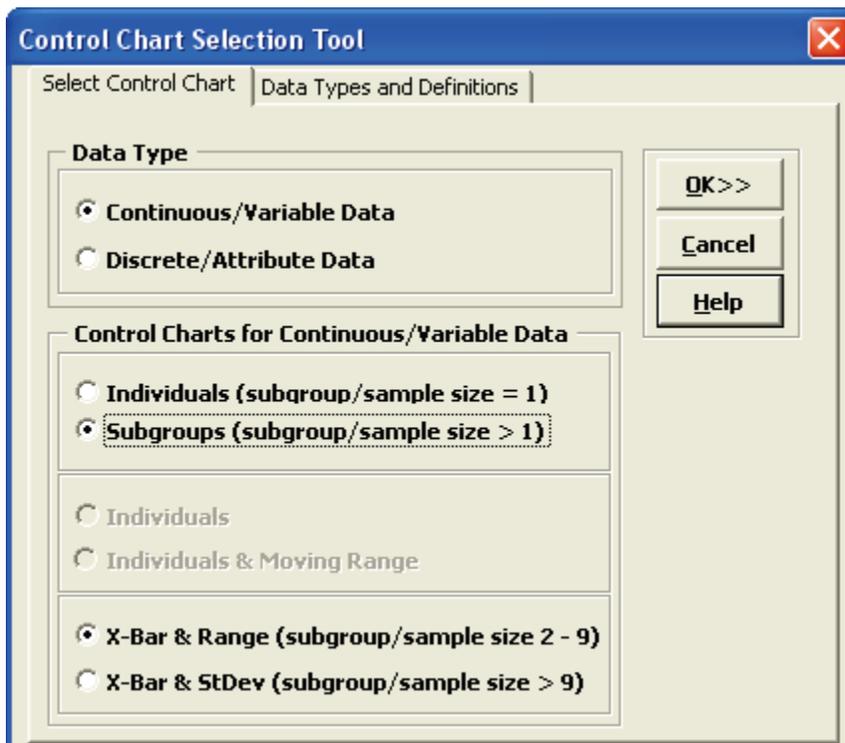
1. Open **Customer Data.xlsx**, click on **Sheet 1**. Click **SigmaXL > Control Charts > Control Chart Selection Tool**.
2. We would like to create a control chart of the Overall Satisfaction data. Since this can be considered as continuous data, the data type is **Continuous/Variable Data**. The subgroup/sample size is 1 (i.e. there is no subgrouping), so select **Individuals (subgroup/sample size = 1)**. At this point, we can choose **Individuals** or **Individuals and Moving Range**. We will keep the simpler **Individuals** selection as shown. (Note that the above data types and definitions can be viewed by clicking the **Data Types and Definitions** tab):



3. Click **OK**. This starts up the Individuals Chart dialog (see Part A – Individuals Charts for continuation).

Control Chart Selection Tool – X-Bar & R Chart

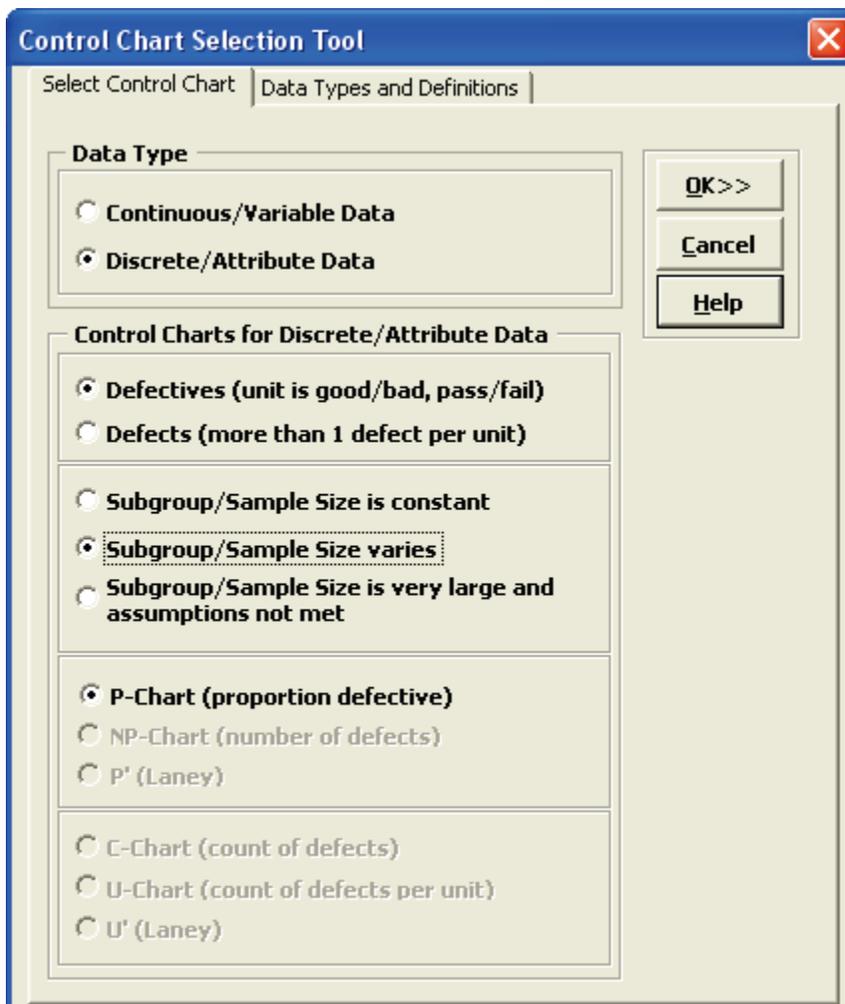
1. Open the file **Catapult Data – Xbar Control Charts.xlsx**. Each operator fires the ball 3 times. The target distance is 100 inches. Select **B2:F22**; here we will only use the first 20 subgroups to determine the control limits.
2. Click **SigmaXL > Control Charts > Control Chart Selection Tool**.
3. Since catapult shot distance measurement is continuous, we keep the default selection **Continuous/Variable Data**. The catapult shot data are in subgroups, so select **Subgroups (subgroup/sample size > 1)**. The subgroup/sample size is small (3), so we will use the **X-Bar & Range Chart** as shown:



4. Click **OK**. This starts up the **X-Bar & Range** dialog (see Part B – X-Bar & Range Charts for continuation).

Control Chart Selection Tool – P-Chart

1. Open **New York Daily Cycle Time – Discrete.xlsx**. This is data from the Sigma Savings and Loans Company, New York location. Each day, the cycle time (in days) for completed loans and leases was recorded. *N* indicates the number of loans counted. A *Fail* was recorded if the cycle time exceeded the critical customer requirement of 8 days.
2. Click **SigmaXL > Control Charts > Control Chart Selection Tool**.
3. Since this data is discrete, select **Discrete/Attribute Data**. We are looking at Defectives data since each loan is a pass or fail, so select **Defectives (unit is good/bad, pass/fail)**. The subgroup/sample size varies day to day so **Subgroup/Sample Size varies** is selected as shown. The recommended chart is the **P-Chart (proportion defective)**:



4. Click **OK**. This starts up the P-Chart dialog (see Part C – P-Charts for continuation).

Part F – I-MR-R/S Charts

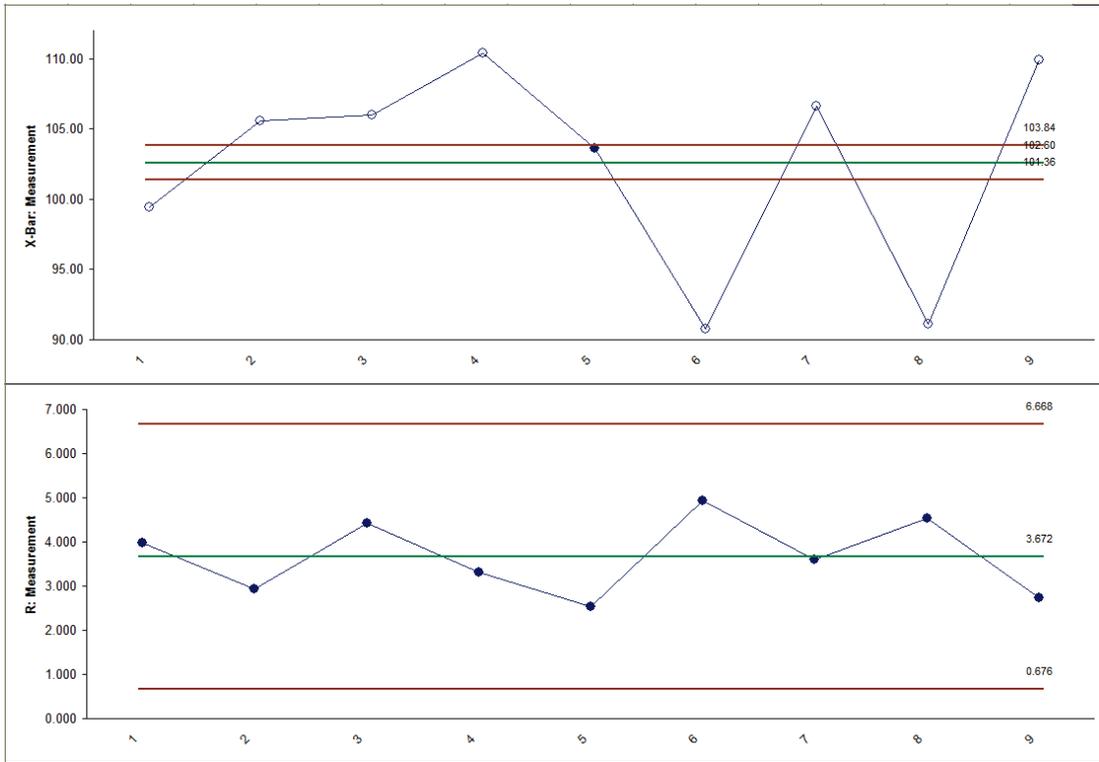
I-MR-R Chart

If the within-subgroup variability is much smaller than between subgroup, the classical X-bar & R (or S) chart will not work, producing numerous (false) alarms. The correct chart to use, in this case, is the I-MR-R (or S) chart. The subgroup averages are treated as individual values (I-MR) and the within subgroup ranges are plotted on the Range chart.

1. Open **Multi-Vari Data.xlsx**. Select Sheet **Between**. We saw this data previously using Multi-Vari charts. First, we will incorrectly use the X-bar & R chart, and then apply the correct I-MR-R chart.
2. Click **SigmaXL > Control Charts > X-bar & R**. Check **Use Entire Data Table**.
3. Click **Next**. Select **Stacked Column Format**. Select *Measurement*, click **Numeric Data Variable (Y) >>**; select *unit*, click **Subgroup Column or Size >>**.

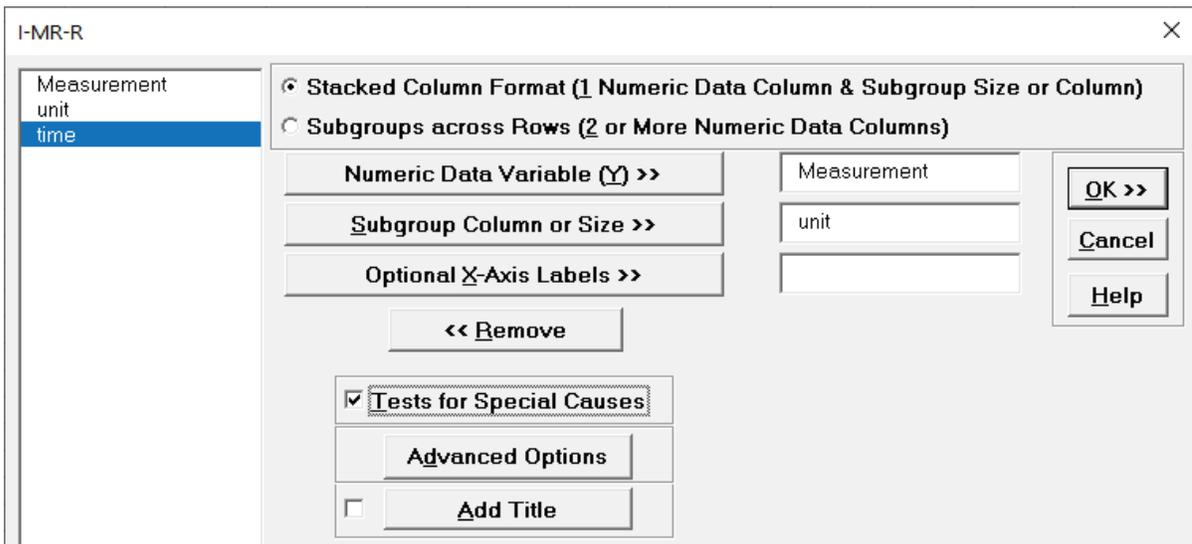
The screenshot shows the 'X-Bar and Range Chart' dialog box. On the left, a list of variables includes 'Measurement', 'unit', and 'time', with 'unit' selected. The main area has two radio buttons: 'Stacked Column Format (1 Numeric Data Column & Subgroup Size or Column)' (selected) and 'Subgroups across Rows (2 or More Numeric Data Columns)'. Below these are three input fields: 'Numeric Data Variable (Y) >>' containing 'Measurement', 'Subgroup Column or Size >>' containing 'unit', and 'Optional X-Axis Labels >>' which is empty. To the right of these fields are 'OK >>', 'Cancel', and 'Help' buttons. Below the input fields is a '<< Remove' button. At the bottom left, there are three checkboxes: 'Calculate Limits' (checked), 'Historical Limits' (unchecked), and 'Tests for Special Causes' (unchecked). Below these is a 'Sigma Zone Lines' checkbox (unchecked) and an 'Advanced Options' button. At the very bottom is an 'Add Title' checkbox (unchecked). On the right side of the dialog, there are two columns of input fields: 'X-Bar' and 'R'. Under 'X-Bar' are three empty fields. Under 'R' are three empty fields. Between the columns are labels 'UCL', 'CL', and 'LCL' corresponding to the rows.

4. Click **OK**. Resulting X-bar & R chart:

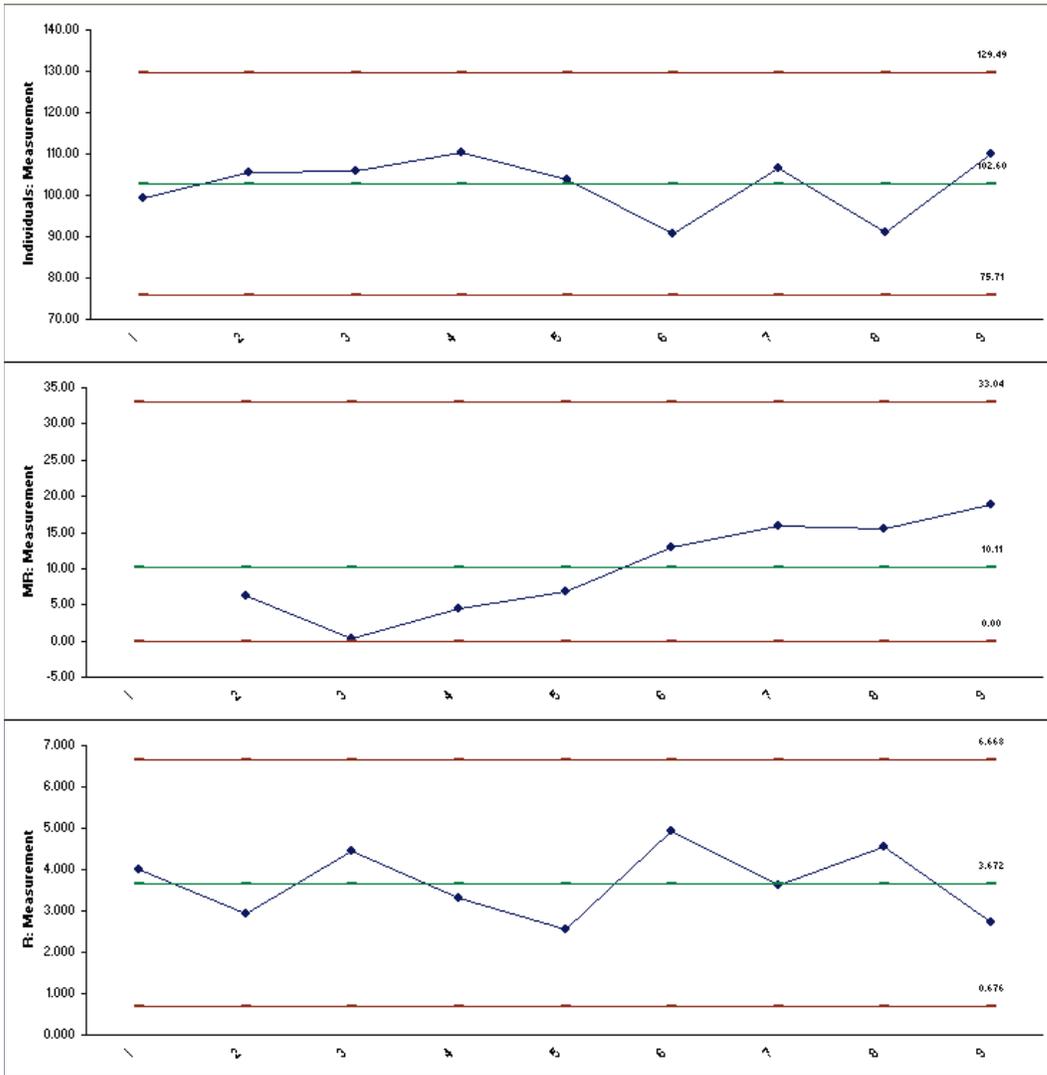


Almost all of the data points in the X-bar chart are out-of-control! This is due to the small within-subgroup variability (the control limits are calculated from the within-subgroup variability).

5. Click Sheet **Between**. Select **SigmaXL > Control Charts > I-MR-R**.
6. Click **Next**. Select **Stacked Column Format**. Select *Measurement*, click **Numeric Data Variable (Y) >>**; select *unit*, click **Subgroup Column or Size >>**. Check **Tests for Special Causes**.



7. Click **OK**. Resulting I-MR-R chart:



This chart is much cleaner, showing a stable Individuals and Range chart. The MR chart may be trending up, but we would want to collect more data before making this conclusion. Typically, you want at least 20 (30 preferred) subgroups before calculating final control limits.

Part G – Control Chart Templates: Rare Events

Rare Events T Chart - Introduction

The Rare Events T Chart (or time-between chart) is an alternative to a standard attribute chart when the adverse event of interest is relatively rare and a measurement of time between each occurrence can be obtained. If the rate of occurrence follows a Poisson distribution (the usual assumption for a C or U chart), then the times between occurrences will have an exponential distribution. The exponential distribution can be transformed to a symmetric Weibull distribution by raising the time measure to the (1/3.6) power. A Shewhart Individuals chart is calculated on the transformed data and then an inverse transformation is applied to the control limits in order to get back to the original time between units. This is an approximate model to the exponential and will result in asymmetric control limits. Note that an "out-of-control" signal above the UCL is desirable, indicating a significant increase in time between adverse rare events. See Provost and Murray, 2011 and Nelson, 1994. The Provost and Murray book is popular in health care, so SigmaXL uses the control limit calculations as given in the book:

t = time between incidents

z = transformed time [$z = t^{(1/3.6)}$]

Construct a Shewhart individuals control chart of z values:

\overline{MR}_z = average moving range of z 's

Remove moving range outliers (i.e., exceed UCL for moving range) for added robustness:

\overline{MR}'_z = average moving range of z 's with any MRs $> (3.27 * \overline{MR}_z)$ removed

$CL_z = \bar{Z}$ (average of transformed time)

$$UCL_z = \bar{Z} + 3 * \left(\frac{\overline{MR}'_z}{1.128} \right)$$

$$LCL_z = \bar{Z} - 3 * \left(\frac{\overline{MR}'_z}{1.128} \right)$$

Transform the center line and the limits back to time scale by raising them to the 3.6 power:

$$CL_t = CL_z^{3.6}, UCL_t = UCL_z^{3.6}, LCL_t = LCL_z^{3.6}$$

Tip: Alternatively, a Rare Events T Chart may be created using an Exponential or Weibull Control Chart on the Days Between data. Click **SigmaXL > Control Charts > Nonnormal > Individuals Nonnormal**. Select **Specify Distribution, Exponential (1 Parameter)** or **Weibull (2 Parameter)**. To select which distribution is the best fit, use distribution fitting: **SigmaXL > Control Charts > Nonnormal > Distribution Fitting**. Select **All Transformations & Distributions** and compare.

Rare Events T Chart - Example

1. Click **SigmaXL > Templates & Calculators > Control Chart Templates > Rare Events > Rare Events T**. This template is also located at **SigmaXL > Control Charts > Control Chart Templates > Rare Events > Rare Events T**.
2. Open **Days Between Surgical Site Infections.xlsx (Sheet 1 tab)**. Days between surgical site infections are an important health care metric. Since they are rare events, days between (or time between) events are tracked rather than events per day. The event date/times 1/1/2019 to 6/13/2019 are “Before Improvement” and the days between will be used to calculate the control limits. Date/times 6/14/2019 to 2/26/2020 are “After Improvement” and the days between will be added to the control chart. The goal is to decrease the occurrence of Surgical Site Infections which will increase the days between, with an “out-of-control” signal above the UCL being desirable and confirmation of the improvement effort.

	A	B
1	Date/Time SSI	Before/After Improvement
2	1/1/2019	Before
3	1/4/2019	Before
4	1/17/2019	Before
5	2/2/2019	Before
6	2/14/2019	Before
7	2/20/2019	Before
8	2/24/2019	Before
9	3/2/2019	Before
10	3/11/2019	Before
11	3/17/2019	Before
12	3/18/2019	Before
13	4/6/2019	Before
14	4/6/2019	Before
15	4/29/2019	Before
16	5/3/2019	Before
17	5/4/2019	Before
18	5/18/2019	Before
19	5/26/2019	Before
20	5/30/2019	Before
21	6/11/2019	Before
22	6/13/2019	Before
23	6/14/2019	After
24	7/26/2019	After
25	8/3/2019	After
26	8/6/2019	After
27	8/25/2019	After
28	9/8/2019	After
29	11/13/2019	After
30	2/15/2020	After
31	2/16/2020	After
32	2/26/2020	After

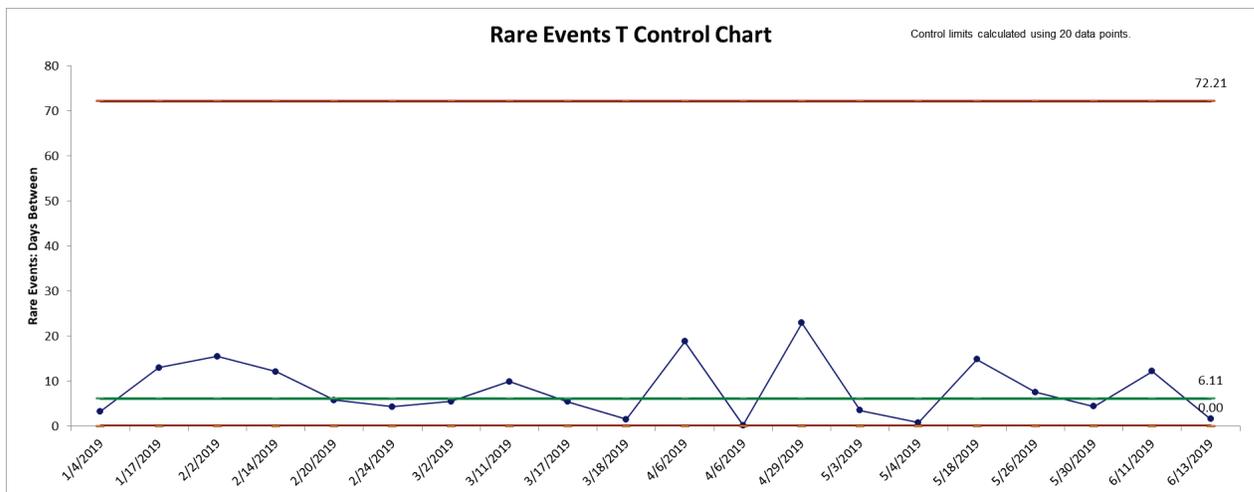
- Copy the “Before Improvement” date/times in cells **A1:A22** and Paste Values to the template at **A1**. **Days Between** are automatically calculated using the Excel formula =A3-A2, =A4-A3, etc. Cell **B2** is not used.

	A	B	C	D	E
1	Date/Time SSI	Days Between			
2	1/1/2019				
3	1/4/2019	3.3			
4	1/17/2019	13			
5	2/2/2019	15.5			
6	2/14/2019	12.1			
7	2/20/2019	5.8			
8	2/24/2019	4.3			
9	3/2/2019	5.5			
10	3/11/2019	9.9			
11	3/17/2019	5.4			
12	3/18/2019	1.5			
13	4/6/2019	18.8			
14	4/6/2019	0.2			
15	4/29/2019	22.9			
16	5/3/2019	3.5			
17	5/4/2019	0.8			
18	5/18/2019	14.8			
19	5/26/2019	7.5			
20	5/30/2019	4.4			
21	6/11/2019	12.2			
22	6/13/2019	1.6			
23					

Rare Events T Control Chart

Add Data

- Click the **Rare Events T Control Chart** button to create the Rare Events T Control Chart:



- This confirms that the process is “in-control”. As discussed above the control limits are asymmetrical.

6. Although the process is “in-control”, as a critical health care metric, efforts were made to improve the process and this data will be now added to the chart. Switch back to **Days Between Surgical Site Infections.xlsx (Sheet 1 tab)**. Select and copy the “After Improvement” date/times in cells **A23: A32** as shown.

	A	B
1	Date/Time SSI	Before/After Improvement
2	1/1/2019	Before
3	1/4/2019	Before
4	1/17/2019	Before
5	2/2/2019	Before
6	2/14/2019	Before
7	2/20/2019	Before
8	2/24/2019	Before
9	3/2/2019	Before
10	3/11/2019	Before
11	3/17/2019	Before
12	3/18/2019	Before
13	4/6/2019	Before
14	4/6/2019	Before
15	4/29/2019	Before
16	5/3/2019	Before
17	5/4/2019	Before
18	5/18/2019	Before
19	5/26/2019	Before
20	5/30/2019	Before
21	6/11/2019	Before
22	6/13/2019	Before
23	6/14/2019	After
24	7/26/2019	After
25	8/3/2019	After
26	8/6/2019	After
27	8/25/2019	After
28	9/8/2019	After
29	11/13/2019	After
30	2/15/2020	After
31	2/16/2020	After
32	2/26/2020	After

7. Switch to the Rare Events T Control Chart template and Paste Values to cell **A23** as shown.

	A	B	C	D	E
1	Date/Time SSI	Days Between			
2	1/1/2019				
3	1/4/2019	3.3			
4	1/17/2019	13			
5	2/2/2019	15.5			
6	2/14/2019	12.1			
7	2/20/2019	5.8			
8	2/24/2019	4.3			
9	3/2/2019	5.5			
10	3/11/2019	9.9			
11	3/17/2019	5.4			
12	3/18/2019	1.5			
13	4/6/2019	18.8			
14	4/6/2019	0.2			
15	4/29/2019	22.9			
16	5/3/2019	3.5			
17	5/4/2019	0.8			
18	5/18/2019	14.8			
19	5/26/2019	7.5			
20	5/30/2019	4.4			
21	6/11/2019	12.2			
22	6/13/2019	1.6			
23	6/14/2019	0.9			
24	7/26/2019	42.1			
25	8/3/2019	7.5			
26	8/6/2019	3.3			
27	8/25/2019	19.1			
28	9/8/2019	14.1			
29	11/13/2019	66.3			
30	2/15/2020	93.6			
31	2/16/2020	1			
32	2/26/2020	10.2			
33					

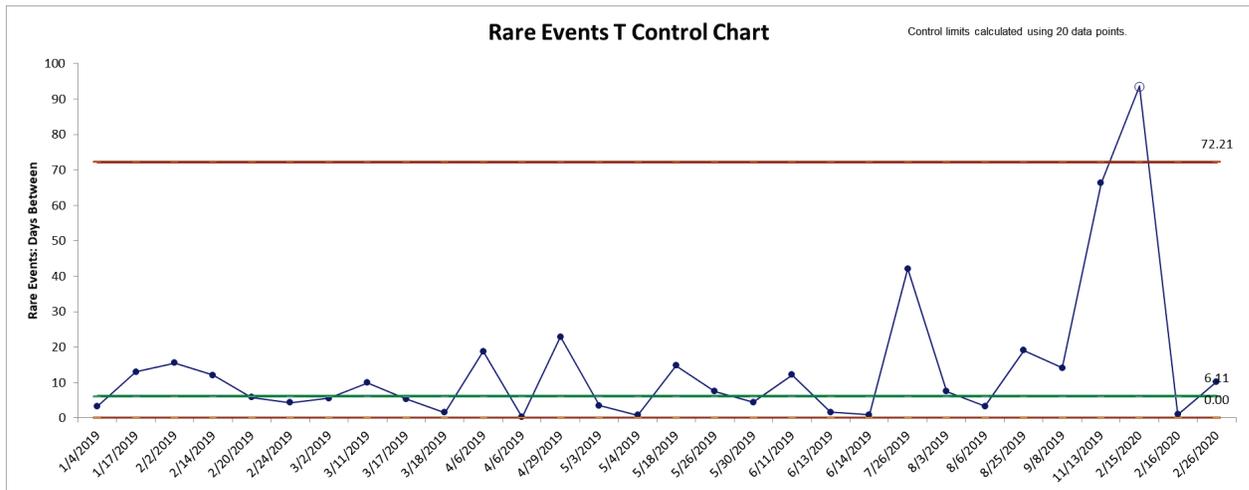
Rare Events T Control Chart

Add Data

Notes:

1. This Rare Events T Con
2. You can replace the Da
3. Enter date or date/time
4. Alternatively, you may r
5. Click the **Rare Events**
6. After the control chart h
7. An "out-of-control" sign:

8. Click the **Add Data** button to add the new **Days Between** data to the Rare Events T Control Chart:



9. This confirms that the process is now “out-of-control”, so the improvement efforts to decrease the occurrence of Surgical Site Infections and increase the Days Between have been successful. Note that the control limits should be recalculated with just the “After Improvement” data when 20 data points are available.

Template Notes and References:

1. This Rare Events T Control Chart template should be used with days or time between (typically adverse) rare events.
2. You can replace the **Date/Time** and **Days Between** column headings with any headings that you wish.
3. Enter date or date/time in the **Date/Time** column. Days between are automatically calculated and entered into **Days Between** column. Cell **B2** is not used in this case.
4. Alternatively, you may manually enter data in **Days Between**. Note, this will overwrite the cell formula.
5. Click the **Rare Events T Control Chart** button to create a control chart. This will overwrite any existing control chart.
6. After the control chart has been created and additional new **Date/Time** or **Days Between** data entered, click the **Add Data** button to add the data to the existing chart.
7. An "out-of-control" signal above the UCL is desirable, indicating a significant increase in time between adverse rare events.
8. Be careful to not have any zeros in the data.
9. This chart uses the control limit formulas given in Provost and Murray, 2011.
10. The data are transformed with $Y^{(1/3.6)}$, which transforms an exponential distribution to a symmetric Weibull. See Nelson, 1994.
11. A Shewhart Individuals Chart is created on the transformed data and the final control limits are then calculated as $UCL^{3.6}$, $CL^{3.6}$, $LCL^{3.6}$.

12. Alternatively, a Rare Events T Chart may be created using a Weibull or Exponential Control Chart. Click **SigmaXL > Control Charts > Nonnormal > Individuals Nonnormal**. Select **Specify Distribution, Exponential (1 Parameter)** or **Weibull (2 Parameter)**.

13. References:

Nelson, L.S. (1994), "A Control Chart for Parts-Per-Million Nonconforming Items", *Journal of Quality Technology*, 26:3, pp. 239-240.

Provost L, Murray S. (2011), *The Health Care Data Guide: Learning from Data for Improvement*. San Francisco: Jossey-Bass, pp. 230-231.

Rare Events G Chart - Introduction

The G chart (or Geometric chart) is an alternative to a standard attribute chart when the adverse event of interest is rare and discrete opportunities between events are counted (e.g., number of units or days between).

The calculation of control limits is an approximation based on the geometric distribution. An "out-of-control" signal above the UCL is desirable, indicating a significant increase in units/opportunities or days between adverse rare events. See Provost and Murray, 2011. The Provost and Murray book is popular in health care, so SigmaXL uses the control limit calculations as given in the book:

g = number of opportunities or units between incidents

\bar{g} = average of g 's

$CL = 0.693 * \bar{g}$ (CL is the theoretical median for a geometric distribution)

$UCL = \bar{g} + 3 * \sqrt{\bar{g}(1 + \bar{g})}$

$LCL = 0$

For more accurate probability-based limits, see the Rare Events Prob G Control Chart.

Rare Events G Chart - Example

1. Click **SigmaXL > Templates & Calculators > Control Chart Templates > Rare Events > Rare Events G**. This template is also located at **SigmaXL > Control Charts > Control Chart Templates > Rare Events > Rare Events G**.
2. Open **Doses Dispensed Between Adverse Drug Events.xlsx (Sheet 1 tab)**. Adverse Drug Events (ADE) are rare, so Doses Dispensed Between ADEs is used as a health care metric. The data is discrete and distributed geometrically, so the G chart is appropriate to use here. The event dates 7/2/2019 to 9/24/2019 are "Before Improvement" and the Doses Between ADEs will be used to calculate the control limits. Dates 9/28/2019 to 1/17/2020 are "After Improvement" and the Doses Between ADEs will be added to the control chart. The goal is to decrease the occurrence of ADEs which will increase the doses between, with an "out-of-control" signal above the UCL being desirable and confirmation of the improvement effort.

SigmaXL: Control Phase Tools: SPC

	A	B	C
1	Date ADE	Doses Between Events	Before/After Improvement
2	7/2/2019	758	Before
3	7/14/2019	11911	Before
4	7/21/2019	7593	Before
5	7/22/2019	398	Before
6	7/24/2019	2516	Before
7	7/26/2019	1147	Before
8	7/29/2019	3322	Before
9	8/13/2019	15414	Before
10	8/14/2019	626	Before
11	8/15/2019	752	Before
12	8/23/2019	7819	Before
13	8/25/2019	2445	Before
14	8/30/2019	4937	Before
15	9/4/2019	4663	Before
16	9/4/2019	175	Before
17	9/6/2019	2187	Before
18	9/7/2019	592	Before
19	9/11/2019	3721	Before
20	9/20/2019	9317	Before
21	9/24/2019	4269	Before
22	9/28/2019	3058	After
23	10/22/2019	24638	After
24	10/27/2019	4579	After
25	10/28/2019	1497	After
26	11/1/2019	3761	After
27	11/19/2019	17947	After
28	12/1/2019	11790	After
29	12/3/2019	1725	After
30	1/4/2020	32255	After
31	1/17/2020	12793	After

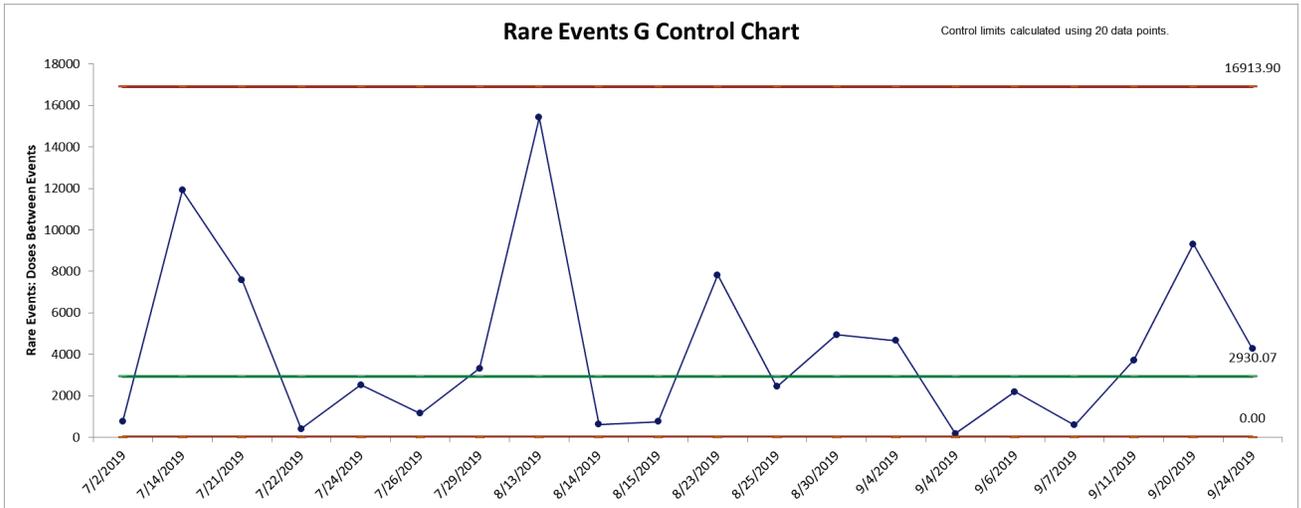
- Copy the “Before Improvement” data in cells **A1:B21** and Paste Values to the template at **A1**. This will overwrite the formulas in Column B. Adjust the template column widths as necessary.

	A	B	C	D	E
1	Date ADE	Doses Between Events			
2	7/2/2019	758			
3	7/14/2019	11911			
4	7/21/2019	7593			
5	7/22/2019	398			
6	7/24/2019	2516			
7	7/26/2019	1147			
8	7/29/2019	3322			
9	8/13/2019	15414			
10	8/14/2019	626			
11	8/15/2019	752			
12	8/23/2019	7819			
13	8/25/2019	2445			
14	8/30/2019	4937			
15	9/4/2019	4663			
16	9/4/2019	175			
17	9/6/2019	2187			
18	9/7/2019	592			
19	9/11/2019	3721			
20	9/20/2019	9317			
21	9/24/2019	4269			
22					

Rare Events G Control Chart

Add Data

4. Click the **Rare Events G Control Chart** button to create the Rare Events G Control Chart:



5. This confirms that the process is “in-control”.

6. Although the process is “in-control”, as a critical health care metric, efforts were made to improve the process and this data will be now added to the chart. Switch back to **Doses Dispensed Between Adverse Drug Events.xlsx (Sheet 1 tab)**. Select and copy the “After Improvement” date/times in cells **A22: B31** as shown.

	A	B	C
	Date ADE	Doses Between Events	Before/After Improvement
1			
2	7/2/2019	758	Before
3	7/14/2019	11911	Before
4	7/21/2019	7593	Before
5	7/22/2019	398	Before
6	7/24/2019	2516	Before
7	7/26/2019	1147	Before
8	7/29/2019	3322	Before
9	8/13/2019	15414	Before
10	8/14/2019	626	Before
11	8/15/2019	752	Before
12	8/23/2019	7819	Before
13	8/25/2019	2445	Before
14	8/30/2019	4937	Before
15	9/4/2019	4663	Before
16	9/4/2019	175	Before
17	9/6/2019	2187	Before
18	9/7/2019	592	Before
19	9/11/2019	3721	Before
20	9/20/2019	9317	Before
21	9/24/2019	4269	Before
22	9/28/2019	3058	After
23	10/22/2019	24638	After
24	10/27/2019	4579	After
25	10/28/2019	1497	After
26	11/1/2019	3761	After
27	11/19/2019	17947	After
28	12/1/2019	11790	After
29	12/3/2019	1725	After
30	1/4/2020	32255	After
31	1/17/2020	12793	After

- Switch to the Rare Events G Control Chart template and Paste Values to cell **A22** as shown. This will overwrite the formulas in Column B. Adjust the template column widths as necessary.

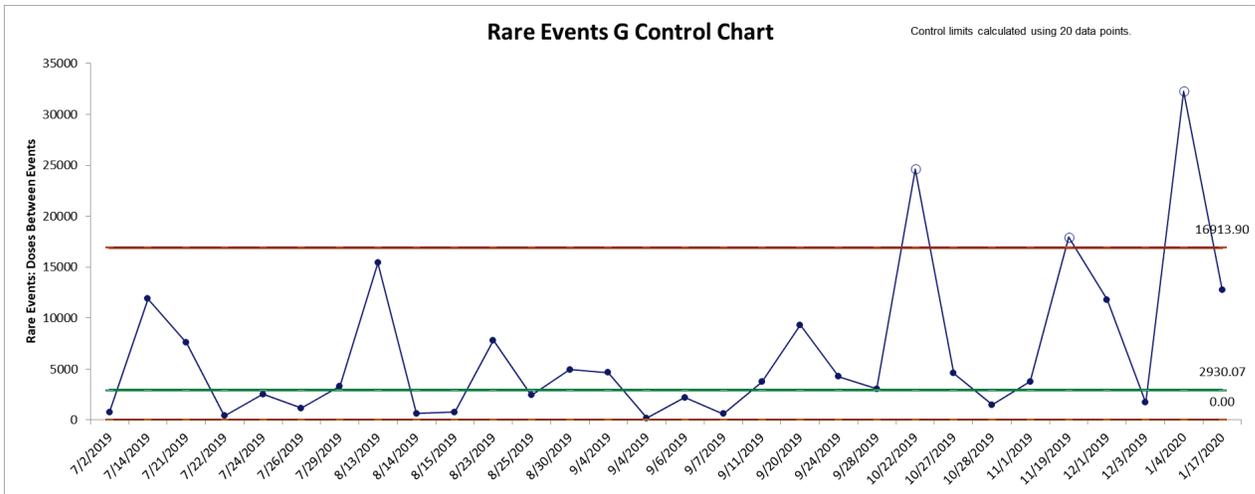
	A	B	C	D	E
1	Date ADE	Doses Between Events			
2	7/2/2019	758			
3	7/14/2019	11911			
4	7/21/2019	7593			
5	7/22/2019	398			
6	7/24/2019	2516			
7	7/26/2019	1147			
8	7/29/2019	3322			
9	8/13/2019	15414			
10	8/14/2019	626			
11	8/15/2019	752			
12	8/23/2019	7819			
13	8/25/2019	2445			
14	8/30/2019	4937			
15	9/4/2019	4663			
16	9/4/2019	175			
17	9/6/2019	2187			
18	9/7/2019	592			
19	9/11/2019	3721			
20	9/20/2019	9317			
21	9/24/2019	4269			
22	9/28/2019	3058			
23	10/22/2019	24638			
24	10/27/2019	4579			
25	10/28/2019	1497			
26	11/1/2019	3761			
27	11/19/2019	17947			
28	12/1/2019	11790			
29	12/3/2019	1725			
30	1/4/2020	32255			
31	1/17/2020	12793			

Rare Events G Control Chart

Add Data

- Notes:
- This Rare Events G Control Chart displays the number of doses between events.
 - You can replace the data in the chart by clicking the Add Data button.
 - Enter date in the Date ADE column.
 - Alternatively, you can enter the date in the Date ADE column.
 - Click the Rare Events G Control Chart button.
 - After the control chart is created, you can click the Add Data button to add new data.

- Click the **Add Data** button to add the new **Doses Between Events** data to the Rare Events G Control Chart:



- This confirms that the process is now “out-of-control”, so the improvement efforts to decrease the occurrence of ADEs and increase the Doses Between have been successful. Note that the control limits should be recalculated with just the “After Improvement” data when 20 data points are available.

Template Notes and Reference:

1. This Rare Events G Control Chart template should be used with days or units/opportunities between (typically adverse) rare events.
2. You can replace the **Date/Time** and **Days Between** column headings with any headings that you wish.
3. Enter date in the **Date** column. Days between are automatically calculated and entered into **Days Between** column. Cell **B2** is not used in this case.
4. Alternatively, you may manually enter data in **Days/Units Between**. Note, this will overwrite the cell formula.
5. Click the **Rare Events G Control Chart** button to create a control chart. This will overwrite any existing control chart.
6. After the control chart has been created and additional new **Date** or **Days/Units Between** data entered, click the **Add Data** button to add the data to the existing chart.
7. An "out-of-control" signal above the UCL is desirable, indicating a significant increase in days or units/opportunities between adverse rare events.
8. The calculation of control limits is based on the geometric distribution. The center line is the theoretical median for a geometric distribution = $0.693 * \text{mean}$.
9. Reference: Provost L, Murray S., *The Health Care Data Guide: Learning from Data for Improvement*. San Francisco: Jossey-Bass, 2011, pp. 228-229.

Rare Events Prob G Chart - Introduction

The Rare Events Prob (Probability-Based) G chart (or Geometric chart) is an alternative to a standard attribute chart when the adverse event of interest is rare and discrete opportunities between events are counted (e.g., number of units or days between). The use of probability-based control limits is recommended in order to properly control the Type I (false alarm) error rate.

The calculation of control limits is based on the geometric distribution. Event probability and alpha are used to compute the non-symmetrical limits. An "out-of-control" signal above the UCL is desirable, indicating a significant increase in units/opportunities or days between adverse rare events. See Benneyan, 2001.

g = number of opportunities or units between incidents

\bar{g} = average of g 's

The event probability is user specified or estimated as:

$$\hat{p} = \left(\frac{1}{\bar{g} + 1} \right) \left(\frac{N - 1}{N} \right)$$

This is the minimum variances unbiased estimator.

The probability-based control limits are calculated as:

$$UCL = \frac{\ln(\alpha_{UCL})}{\ln(1 - p)} - 1$$

$$CL = \frac{\ln(0.5)}{\ln(1 - p)} - 1$$

$$LCL = \max \left(0, \frac{\ln(1 - \alpha_{UCL})}{\ln(1 - p)} - 1 \right)$$

where α_{UCL} is user specified.

SigmaXL adjusts the Benneyan CL and LCL by subtracting 1 as done in the UCL and does not round the limit values. Opportunities between events are counted as 0, 1, 2, 3, ... ($a = 0$ in Benneyan Table 3).

Rare Events Prob G Chart - Example

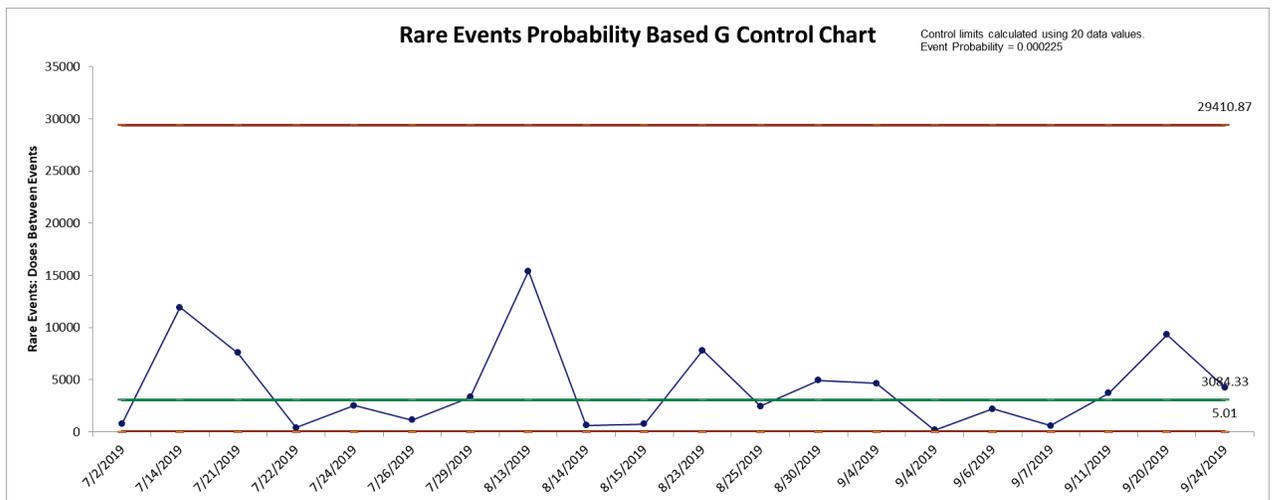
1. Click **SigmaXL > Templates & Calculators > Control Chart Templates > Rare Events > Rare Events Prob G**. This template is also located at **SigmaXL > Control Charts > Control Chart Templates > Rare Events > Rare Events Prob G**.
2. Open **Doses Dispensed Between Adverse Drug Events.xlsx (Sheet 1 tab)**. Adverse Drug Events (ADE) are rare, so Doses Dispensed Between ADEs is used as a health care metric. The data is discrete and distributed geometrically, so the G chart is appropriate to use here. The event dates 7/2/2019 to 9/24/2019 are “Before Improvement” and the Doses Between ADEs will be used to calculate the control limits. Dates 9/28/2019 to 1/17/2020 are “After Improvement” and the Doses Between ADEs will be added to the control chart. The goal is to decrease the occurrence of ADEs which will increase the doses between, with an “out-of-control” signal above the UCL being desirable and confirmation of the improvement effort.

	A	B	C
1	Date ADE	Doses Between Events	Before/After Improvement
2	7/2/2019	758	Before
3	7/14/2019	11911	Before
4	7/21/2019	7593	Before
5	7/22/2019	398	Before
6	7/24/2019	2516	Before
7	7/26/2019	1147	Before
8	7/29/2019	3322	Before
9	8/13/2019	15414	Before
10	8/14/2019	626	Before
11	8/15/2019	752	Before
12	8/23/2019	7819	Before
13	8/25/2019	2445	Before
14	8/30/2019	4937	Before
15	9/4/2019	4663	Before
16	9/4/2019	175	Before
17	9/6/2019	2187	Before
18	9/7/2019	592	Before
19	9/11/2019	3721	Before
20	9/20/2019	9317	Before
21	9/24/2019	4269	Before
22	9/28/2019	3058	After
23	10/22/2019	24638	After
24	10/27/2019	4579	After
25	10/28/2019	1497	After
26	11/1/2019	3761	After
27	11/19/2019	17947	After
28	12/1/2019	11790	After
29	12/3/2019	1725	After
30	1/4/2020	32255	After
31	1/17/2020	12793	After

- Copy the “Before Improvement” data in cells **A1:B21** and Paste Values to the template at **A1**. This will overwrite the formulas in Column B. Adjust the template column widths as necessary.

	A	B	C	D
1	Date ADE	Doses Between Events	Alpha UCL:	0.00135
2	7/2/2019	758	Event Probability:	
3	7/14/2019	11911	Rare Events Prob G Chart	
4	7/21/2019	7593		
5	7/22/2019	398	Add Data	
6	7/24/2019	2516		
7	7/26/2019	1147		
8	7/29/2019	3322		
9	8/13/2019	15414		
10	8/14/2019	626		
11	8/15/2019	752		
12	8/23/2019	7819		
13	8/25/2019	2445		
14	8/30/2019	4937		
15	9/4/2019	4663		
16	9/4/2019	175		
17	9/6/2019	2187		
18	9/7/2019	592		
19	9/11/2019	3721		
20	9/20/2019	9317		
21	9/24/2019	4269		

- We will use the default **Alpha UCL = 0.00135** (this corresponds to the one-sided probability for UCL in a classical Shewhart control chart with 3 sigma limits). Click the **Rare Events Prob G Chart** button to create the Rare Events Probability Based G Control Chart:



- This confirms that the process is “in-control”. Note that the probability based UCL is much larger than the one calculated above for the regular G chart.
- Tip:** If the UCL is considered to be “too large” for practical use, adjust the **Alpha UCL** to a larger value, for example, .01 instead of .00135 and recreate the control chart. Note however that this will increase the Type I (false alarm) rate.

7. Although the process is “in-control”, as a critical health care metric, efforts were made to improve the process and this data will be now added to the chart. Switch back to **Doses Dispensed Between Adverse Drug Events.xlsx (Sheet 1 tab)**. Select and copy the “After Improvement” date/times in cells **A22: B31** as shown.

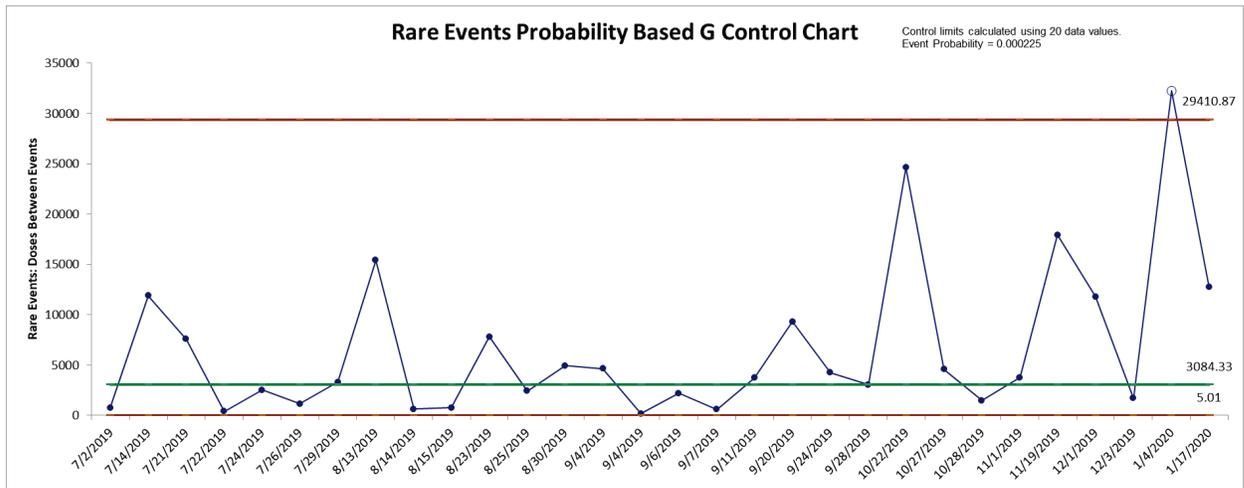
	A	B	C
1	Date ADE	Doses Between Events	Before/After Improvement
2	7/2/2019	758	Before
3	7/14/2019	11911	Before
4	7/21/2019	7593	Before
5	7/22/2019	398	Before
6	7/24/2019	2516	Before
7	7/26/2019	1147	Before
8	7/29/2019	3322	Before
9	8/13/2019	15414	Before
10	8/14/2019	626	Before
11	8/15/2019	752	Before
12	8/23/2019	7819	Before
13	8/25/2019	2445	Before
14	8/30/2019	4937	Before
15	9/4/2019	4663	Before
16	9/4/2019	175	Before
17	9/6/2019	2187	Before
18	9/7/2019	592	Before
19	9/11/2019	3721	Before
20	9/20/2019	9317	Before
21	9/24/2019	4269	Before
22	9/28/2019	3058	After
23	10/22/2019	24638	After
24	10/27/2019	4579	After
25	10/28/2019	1497	After
26	11/1/2019	3761	After
27	11/19/2019	17947	After
28	12/1/2019	11790	After
29	12/3/2019	1725	After
30	1/4/2020	32255	After
31	1/17/2020	12793	After

8. Switch to the Rare Events Prob G Chart template and Paste Values to cell **A22** as shown. This will overwrite the formulas in Column B. Adjust the template column widths as necessary.

	A	B	C	D
1	Date ADE	Doses Between Events	Alpha UCL:	0.00135
2	7/2/2019	758	Event Probability:	
3	7/14/2019	11911	Rare Events Prob G Chart	
4	7/21/2019	7593		
5	7/22/2019	398	Add Data	
6	7/24/2019	2516		
7	7/26/2019	1147		
8	7/29/2019	3322		
9	8/13/2019	15414		
10	8/14/2019	626		
11	8/15/2019	752		
12	8/23/2019	7819		
13	8/25/2019	2445		
14	8/30/2019	4937		
15	9/4/2019	4663		
16	9/4/2019	175		
17	9/6/2019	2187		
18	9/7/2019	592		
19	9/11/2019	3721		
20	9/20/2019	9317		
21	9/24/2019	4269		
22	9/28/2019	3058		
23	10/22/2019	24638		
24	10/27/2019	4579		
25	10/28/2019	1497		
26	11/1/2019	3761		
27	11/19/2019	17947		
28	12/1/2019	11790		
29	12/3/2019	1725		
30	1/4/2020	32255		
31	1/17/2020	12793		

- Notes:
1. This Rare I
 2. You can re
 3. Enter date
 4. Alternative
 5. Enter Alpt

9. Click the **Add Data** button to add the new **Doses Between Events** data to the Rare Events Probability Based G Control Chart:



10. This confirms that the process is now “out-of-control”, so the improvement efforts to decrease the occurrence of ADEs and increase the Doses Between have been successful. Note that the control limits should be recalculated with just the “After Improvement” data when 20 data points are available.

Template Notes and Reference:

1. This Rare Events Prob (Probability-Based) G Control Chart template should be used with days or units/opportunities between (typically adverse) rare events.
2. You can replace the **Date/Time** and **Days Between** column headings with any headings that you wish.
3. Enter date in the **Date** column. Days between are automatically calculated and entered into **Days Between** column. Cell **B2** is not used in this case.
4. Alternatively, you may manually enter data in **Days/Units Between**. Note, this will overwrite the cell formulas.
5. Enter **Alpha UCL**, typically 0.00135, corresponding to the one-sided probability for UCL in a classical Shewhart control chart with 3 sigma limits. This will also be applied to the LCL if greater than 0.
6. The calculation of control limits is based on the geometric distribution. Event probability and alpha are used to compute the non-symmetrical limits. Calculated event probability is noted on the control chart.
7. Optionally enter the historical **Event Probability**.
8. Click the **Rare Events Prob G Chart** button to create a control chart. This will overwrite any existing control chart.

9. After the control chart has been created and additional new **Date** or **Days/Units Between** data entered, click the **Add Data** button to add the data to the existing chart. Control limits will be calculated using the original chart event probability or specified event probability.
10. **Add Data** should only be used if there are at least 20 observations in the original chart or known historical Event Probability has been specified.
11. An "out-of-control" signal above the UCL is desirable, indicating a significant increase in days or units/opportunities between adverse rare events.
12. Reference: Benneyan, J.C. (2001), "Performance of Number-Between g-Type Statistical Control Charts for Monitoring Adverse Events", *Health Care Management Science*, 4, pp. 319–336, Table 3.

Part H – Control Chart Templates: Time-Weighted

Exponentially Weighted Moving Average (EWMA) Chart - Introduction

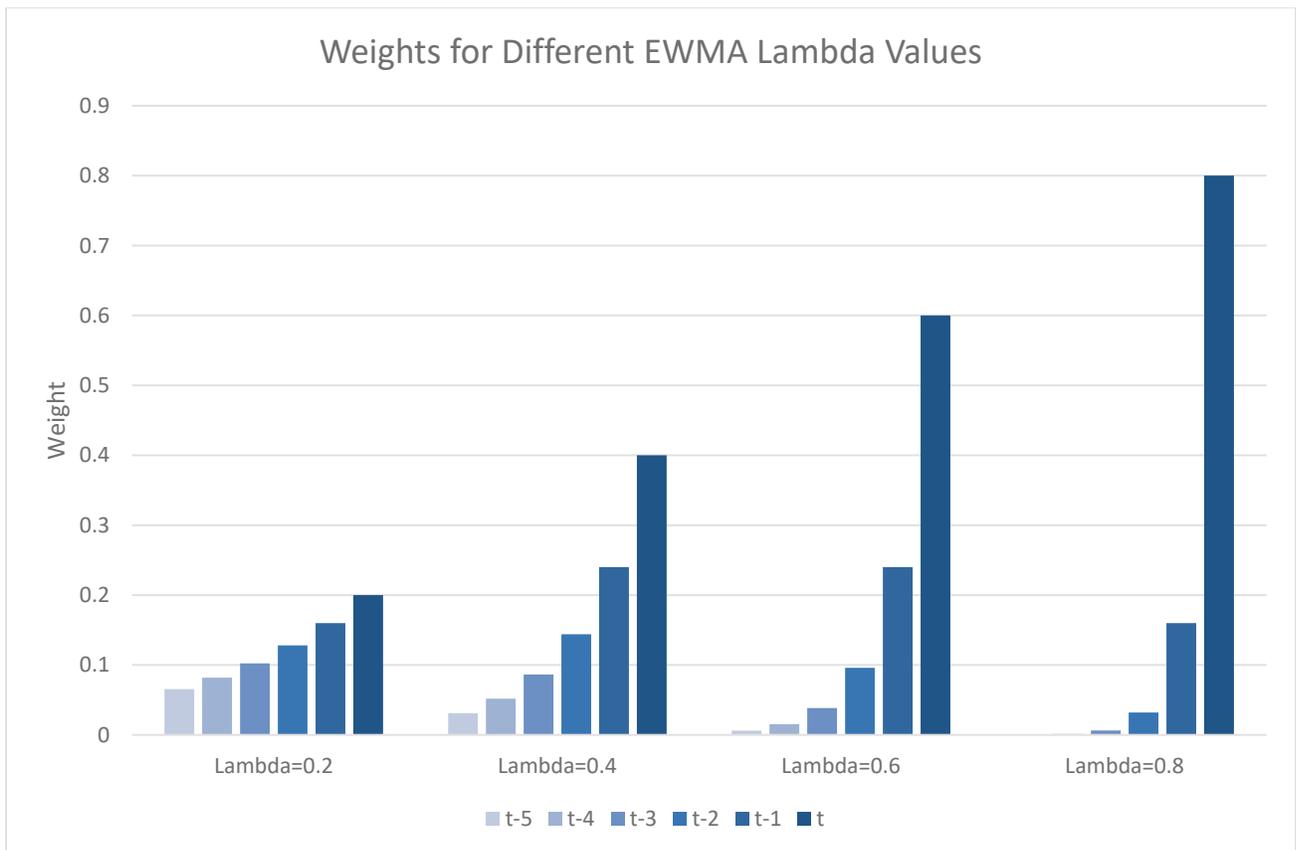
The EWMA control chart uses weighted averages, where the weights decrease exponentially as observations come from further in the past with the smallest weights associated with the oldest observations.

The formula used for the EWMA statistic is given as:

$$EWMA_t = \lambda X_t + (1 - \lambda)EWMA_{t-1}$$

with the starting value $EWMA_0$ estimated as the data mean and the smoothing parameter λ specified according to desired average run length characteristics (see Average Run Length (ARL) Tables).

Weights for different λ values are shown graphically:



The control limits for an EWMA chart are calculated as:

$$\text{UCL} = \mu_0 + K\sigma \sqrt{\frac{\lambda}{(2-\lambda)} [1 - (1-\lambda)^{2t}]}$$

$$\text{CL} = \mu_0$$

$$\text{LCL} = \mu_0 - K\sigma \sqrt{\frac{\lambda}{(2-\lambda)} [1 - (1-\lambda)^{2t}]}$$

where μ_0 is specified as the historical mean or estimated as:

$$\hat{\mu}_0 = \bar{x}$$

and σ is specified as the historical standard deviation or estimated as:

$$\hat{\sigma} = \overline{MR}/d_2.$$

Note, the K multiplier is called L in Montgomery (2013).

Exponentially Weighted Moving Average (EWMA) Chart - Example

1. Click **SigmaXL > Templates & Calculators > Control Chart Templates > Time Weighted > EWMA**. This template is also located at **SigmaXL > Control Charts > Control Chart Templates > Time Weighted > EWMA**.
2. Open **Montgomery Table 9.1.xlsx (Sheet 1 tab)**. This is simulated data from Montgomery, D.C. (2013), *Introduction to Statistical Quality Control*, Seventh Ed., Wiley, pp. 415. Table 9.1. Samples 1 to 20 are drawn from a random normal distribution with population mean $\mu = 10$ and population standard deviation $\sigma = 1$. Samples 21 to 30 have a mean = 11 and standard deviation = 1, so the process has shifted by 1 sigma and is “out-of-control”. While this is a small shift in process mean, it is something that we want to be able to detect and correct as quickly as possible.

	A	B	C	D
1	Sample	X	Process Mean	Process StDev
2	1	9.45	10	1
3	2	7.99	10	1
4	3	9.29	10	1
5	4	11.66	10	1
6	5	12.16	10	1
7	6	10.18	10	1
8	7	8.04	10	1
9	8	11.46	10	1
10	9	9.2	10	1
11	10	10.34	10	1
12	11	9.03	10	1
13	12	11.47	10	1
14	13	10.51	10	1
15	14	9.4	10	1
16	15	10.08	10	1
17	16	9.37	10	1
18	17	10.62	10	1
19	18	10.31	10	1
20	19	8.52	10	1
21	20	10.84	10	1
22	21	10.9	11	1
23	22	9.33	11	1
24	23	12.29	11	1
25	24	11.5	11	1
26	25	10.6	11	1
27	26	11.08	11	1
28	27	10.38	11	1
29	28	11.62	11	1
30	29	11.31	11	1
31	30	10.52	11	1

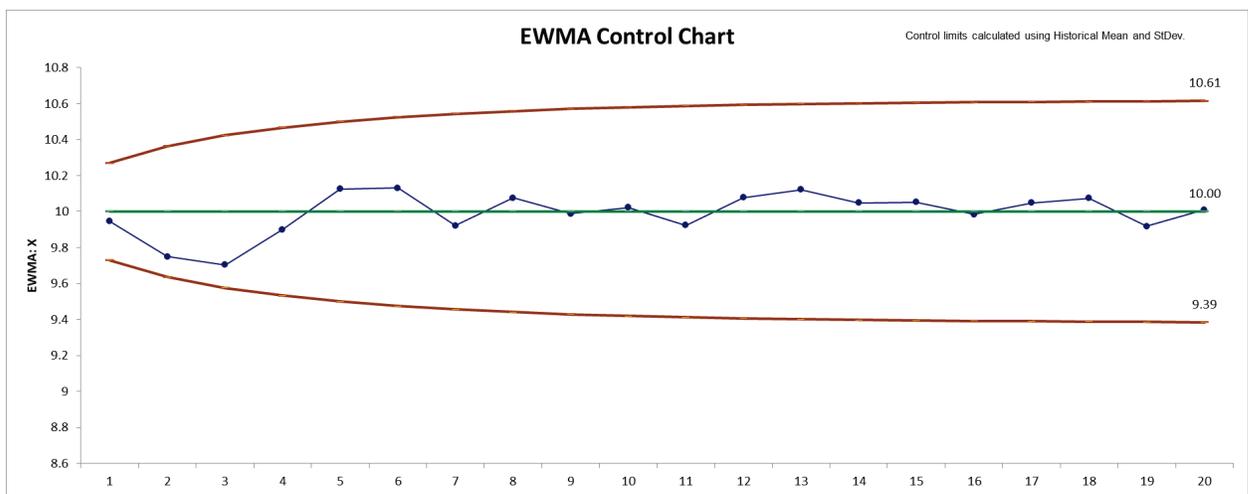
- Copy cells **A1:B21** and Paste Values to the template at **A1**. Use the default **Weight (Lambda) = 0.1**, **K = 2.7**. Specify **Historical Mean = 10** and **Historical StDev = 1** as shown:

	A	B	C	D
1	Sample	X	Weight (Lambda):	0.1
2	1	9.45	K:	2.7
3	2	7.99	Historical Mean:	10
4	3	9.29	Historical StDev:	1
5	4	11.66		
6	5	12.16		
7	6	10.18		
8	7	8.04		
9	8	11.46		
10	9	9.2		
11	10	10.34		
12	11	9.03		
13	12	11.47		
14	13	10.51		
15	14	9.4		
16	15	10.08		
17	16	9.37		
18	17	10.62		
19	18	10.31		
20	19	8.52		
21	20	10.84		
22				

EWMA Control Chart

Add Data

- Click the **EWMA Control Chart** button to create the EWMA Control Chart:



- This confirms that the process is “in-control”.

6. Switch back to **Montgomery Table 9.1.xlsx (Sheet 1 tab)**. Select and copy cells **A22: B31** as shown. This is the data with shifted mean.

	A	B	C	D
1	Sample	X	Process Mean	Process StDev
2	1	9.45	10	1
3	2	7.99	10	1
4	3	9.29	10	1
5	4	11.66	10	1
6	5	12.16	10	1
7	6	10.18	10	1
8	7	8.04	10	1
9	8	11.46	10	1
10	9	9.2	10	1
11	10	10.34	10	1
12	11	9.03	10	1
13	12	11.47	10	1
14	13	10.51	10	1
15	14	9.4	10	1
16	15	10.08	10	1
17	16	9.37	10	1
18	17	10.62	10	1
19	18	10.31	10	1
20	19	8.52	10	1
21	20	10.84	10	1
22	21	10.9	11	1
23	22	9.33	11	1
24	23	12.29	11	1
25	24	11.5	11	1
26	25	10.6	11	1
27	26	11.08	11	1
28	27	10.38	11	1
29	28	11.62	11	1
30	29	11.31	11	1
31	30	10.52	11	1

7. Switch to the EWMA template and Paste Values to cell **A22** as shown.

	A	B	C	D
1	Sample	X	Weight (Lambda):	0.1
2	1	9.45	K:	2.7
3	2	7.99	Historical Mean:	10
4	3	9.29	Historical StDev:	1
5	4	11.66		
6	5	12.16		
7	6	10.18		
8	7	8.04		
9	8	11.46		
10	9	9.2		
11	10	10.34		
12	11	9.03		
13	12	11.47		
14	13	10.51		
15	14	9.4		
16	15	10.08		
17	16	9.37		
18	17	10.62		
19	18	10.31		
20	19	8.52		
21	20	10.84		
22	21	10.9		
23	22	9.33		
24	23	12.29		
25	24	11.5		
26	25	10.6		
27	26	11.08		
28	27	10.38		
29	28	11.62		
30	29	11.31		
31	30	10.52		

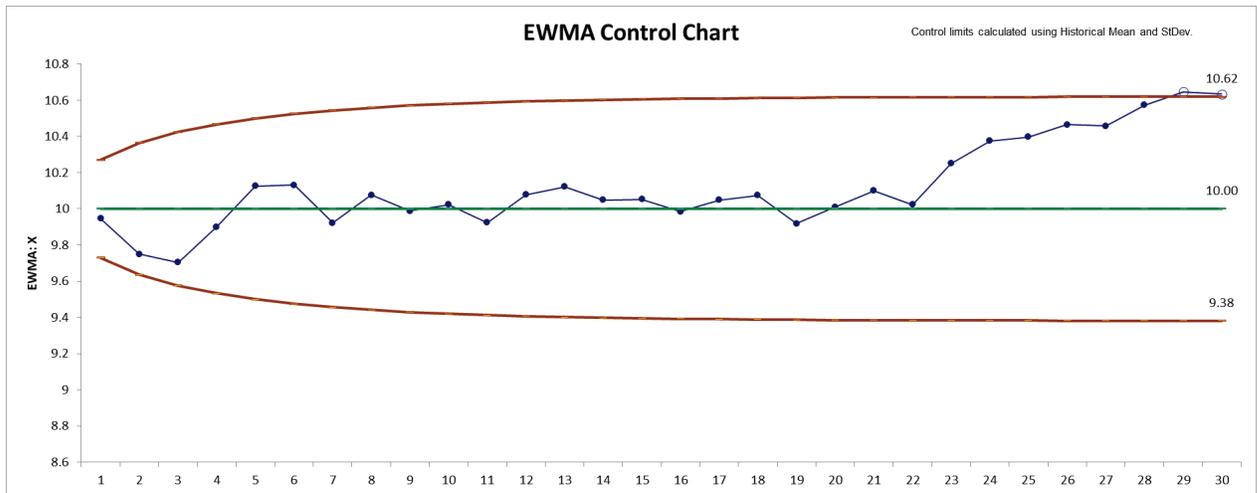
EWMA Control Chart

Add Data

Notes:

1. This Exponer
2. You can replk
3. Enter labels i
4. Enter the EW
5. Enter the EW

8. Click the **Add Data** button to add this data to the EWMA Control Chart:



9. This confirms that the process is now “out-of-control” with signals at samples 29 and 30. This matches Figure 9.7 in Montgomery.

Template Notes and Reference

1. This Exponentially Weighted Moving Average (EWMA) Control Chart template should be used with continuous data. The data must be in chronological time-sequence order.
2. You can replace the **X-Axis Label** and **Data** column headings with any headings that you wish. Enter your data in the **Data** column.
3. Enter labels in **X-Axis Label** column. Labels can be Date, Time, Name, or other text information. These labels are optional and will appear on the horizontal X-Axis of the EWMA Control Chart.
4. Enter the EWMA **Weight (Lambda)** in cell D1. This is a value between 0 and 1 and controls the amount of influence that previous observations have on the current EWMA statistic. A value near 1 puts almost all weight on the current observation, making it resemble a Shewhart chart. For values near 0, a small weight is applied to almost all of the past observations, so the EWMA chart performance is similar to that of a CUSUM chart.
5. Enter the EWMA **K Sigma multiplier** in cell D2. This is a value typically between 2 and 4. It is also referred to as L, but SigmaXL uses K to avoid confusion with Lambda.
6. **Weight (Lambda)** and **K** values affect the Average Run Length (ARL) characteristics. To determine optimal EWMA parameter values and calculate ARL, click **SigmaXL > Templates and Calculators > Control Chart Templates > Average Run Length (ARL) Calculators > EWMA ARL**.
7. **Historical Mean (D3)** and **Historical StDev (D4)** are optional. Enter values if process mean and standard deviation are known.

8. Click the **EWMA Control Chart** button to create an EWMA Control Chart.
9. After the control chart has been created and additional new data entered into the **Data** column, click the **Add Data** button to add the data to the existing chart. Control limits will be calculated using the original chart mean and stdev or specified Historical Mean/StDev.
10. **Add Data** should only be used if there are at least 20 observations in the original chart, or Historical Mean and StDev have been specified.
11. **Weight (Lambda)** and **K** parameters are dynamic. If they are modified, the chart will automatically update with the new parameters. However, they should be selected prior to creating the EWMA chart. Data values and out-of-control formatting are refreshed only when the buttons are used.
12. Reference: Montgomery, D.C. (2013), *Introduction to Statistical Quality Control*, Seventh Ed., Wiley, pp. 433-438.

Tabular Cumulative Sum (CUSUM) Chart - Introduction

The CUSUM chart plots the cumulative sums of deviations of sample values from a target value. Because they combine information from several samples, cumulative sum charts are more effective than Shewhart charts for detecting small process shifts. There are two ways to represent CUSUMs: the tabular (or algorithmic) CUSUM, and the V-mask form of the CUSUM (see Montgomery, 2013). SigmaXL utilizes the Tabular CUSUM.

The formulas used for the Tabular CUSUM statistics are given as:

$$C_1^+ = \max[0, x_t - (\mu_0 + k\sigma) + FIR\sigma]$$

$$C_1^- = \min[0, x_t - (\mu_0 - k\sigma) + FIR\sigma]$$

$$C_t^+ = \max[0, x_t - (\mu_0 + k\sigma) + C_{t-1}^+]$$

$$C_t^- = \min[0, x_t - (\mu_0 - k\sigma) + C_{t-1}^-]$$

where μ_0 is specified as the Target or estimated as:

$$\hat{\mu}_0 = \bar{x}$$

and σ is specified as the historical standard deviation or estimated as:

$$\hat{\sigma} = \overline{MR}/d_2$$

and k and FIR are specified.

k is the reference (or slack) value, typically set to 0.5. It sets the size of mean shift ($2k\sigma$) that you would like to detect quickly, so 0.5 denotes rapid detection of a shift in mean = 1σ .

FIR is the optional fast initial response (or headstart) value. This sets the initial CUSUM statistic so that it improves the sensitivity to a mean shift at startup.

The Tabular CUSUM control limits are calculated as:

$$UCL = h\sigma$$

$$CL = 0$$

$$LCL = -h\sigma$$

where h is the decision interval, typically set to 4 or 5.

If FIR is used, it is typically set to $h/2$.

Tabular Cumulative Sum (CUSUM) Chart - Example

1. Click **SigmaXL > Templates & Calculators > Control Chart Templates > Time Weighted > CUSUM**. This template is also located at **SigmaXL > Control Charts > Control Chart Templates > Time Weighted > CUSUM**.
2. Open **Montgomery Table 9.1.xlsx (Sheet 1 tab)**. This is simulated data from Montgomery, D.C. (2013), *Introduction to Statistical Quality Control*, Seventh Ed., Wiley, pp. 415. Table 9.1. Samples 1 to 20 are drawn from a random normal distribution with population mean $\mu = 10$ and population standard deviation $\sigma = 1$. Samples 21 to 30 have a mean = 11 and standard deviation = 1, so the process has shifted by 1 sigma and is “out-of-control”. While this is a small shift in process mean, it is something that we want to be able to detect and correct as quickly as possible.

	A	B	C	D
1	Sample	X	Process Mean	Process StDev
2	1	9.45	10	1
3	2	7.99	10	1
4	3	9.29	10	1
5	4	11.66	10	1
6	5	12.16	10	1
7	6	10.18	10	1
8	7	8.04	10	1
9	8	11.46	10	1
10	9	9.2	10	1
11	10	10.34	10	1
12	11	9.03	10	1
13	12	11.47	10	1
14	13	10.51	10	1
15	14	9.4	10	1
16	15	10.08	10	1
17	16	9.37	10	1
18	17	10.62	10	1
19	18	10.31	10	1
20	19	8.52	10	1
21	20	10.84	10	1
22	21	10.9	11	1
23	22	9.33	11	1
24	23	12.29	11	1
25	24	11.5	11	1
26	25	10.6	11	1
27	26	11.08	11	1
28	27	10.38	11	1
29	28	11.62	11	1
30	29	11.31	11	1
31	30	10.52	11	1

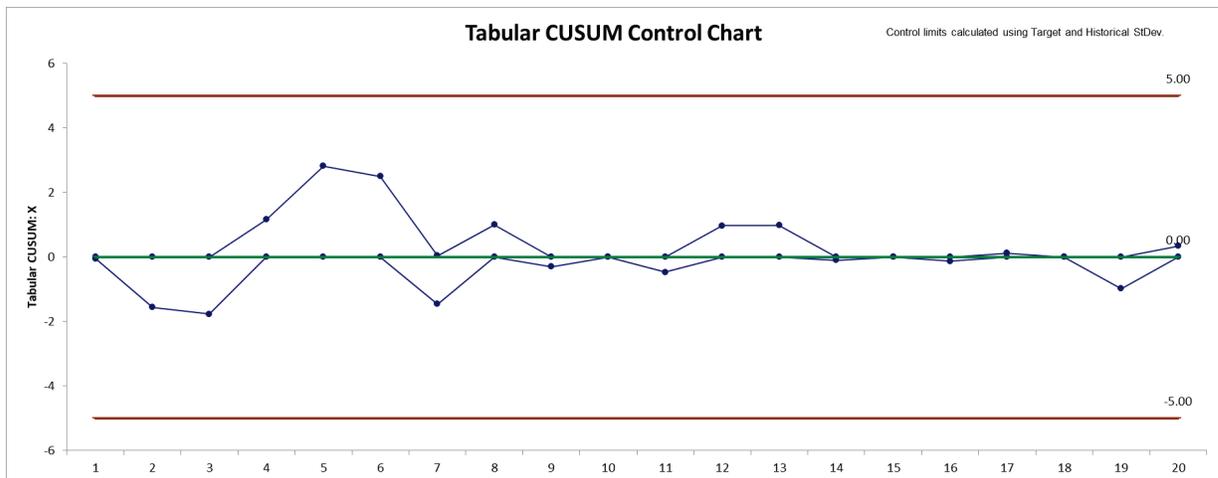
- Copy cells **A1:B21** and Paste Values to the template at **A1**. Use the default **k = 0.5**, **h = 5**, **FIR = 0**. Specify **Target = 10** and **Historical StDev = 1** as shown:

	A	B	C	D
1	Sample	X	k:	0.5
2	1	9.45	h:	5
3	2	7.99	FIR:	0
4	3	9.29	Target:	10
5	4	11.66	Historical StDev:	1
6	5	12.16		
7	6	10.18		
8	7	8.04		
9	8	11.46		
10	9	9.2		
11	10	10.34		
12	11	9.03		
13	12	11.47		
14	13	10.51		
15	14	9.4		
16	15	10.08		
17	16	9.37		
18	17	10.62		
19	18	10.31		
20	19	8.52		
21	20	10.84		

Tabular CUSUM Chart

Add Data

- Click the **Tabular CUSUM Chart** button to create the Tabular CUSUM Control Chart:



- This confirms that the process is “in-control”.

6. Switch back to **Montgomery Table 9.1.xlsx (Sheet 1 tab)**. Select and copy cells **A22: B31** as shown. This is the data with shifted mean.

	A	B	C	D
1	Sample	X	Process Mean	Process StDev
2	1	9.45	10	1
3	2	7.99	10	1
4	3	9.29	10	1
5	4	11.66	10	1
6	5	12.16	10	1
7	6	10.18	10	1
8	7	8.04	10	1
9	8	11.46	10	1
10	9	9.2	10	1
11	10	10.34	10	1
12	11	9.03	10	1
13	12	11.47	10	1
14	13	10.51	10	1
15	14	9.4	10	1
16	15	10.08	10	1
17	16	9.37	10	1
18	17	10.62	10	1
19	18	10.31	10	1
20	19	8.52	10	1
21	20	10.84	10	1
22	21	10.9	11	1
23	22	9.33	11	1
24	23	12.29	11	1
25	24	11.5	11	1
26	25	10.6	11	1
27	26	11.08	11	1
28	27	10.38	11	1
29	28	11.62	11	1
30	29	11.31	11	1
31	30	10.52	11	1

7. Switch to the Tabular CUSUM template and Paste Values to cell **A22** as shown.

	A	B	C	D
1	Sample	X	k:	0.5
2	1	9.45	h:	5
3	2	7.99	FIR:	0
4	3	9.29	Target:	10
5	4	11.66	Historical StDev:	1
6	5	12.16		
7	6	10.18		
8	7	8.04		
9	8	11.46		
10	9	9.2		
11	10	10.34		
12	11	9.03		
13	12	11.47		
14	13	10.51		
15	14	9.4		
16	15	10.08		
17	16	9.37		
18	17	10.62		
19	18	10.31		
20	19	8.52		
21	20	10.84		
22	21	10.9		
23	22	9.33		
24	23	12.29		
25	24	11.5		
26	25	10.6		
27	26	11.08		
28	27	10.38		
29	28	11.62		
30	29	11.31		
31	30	10.52		

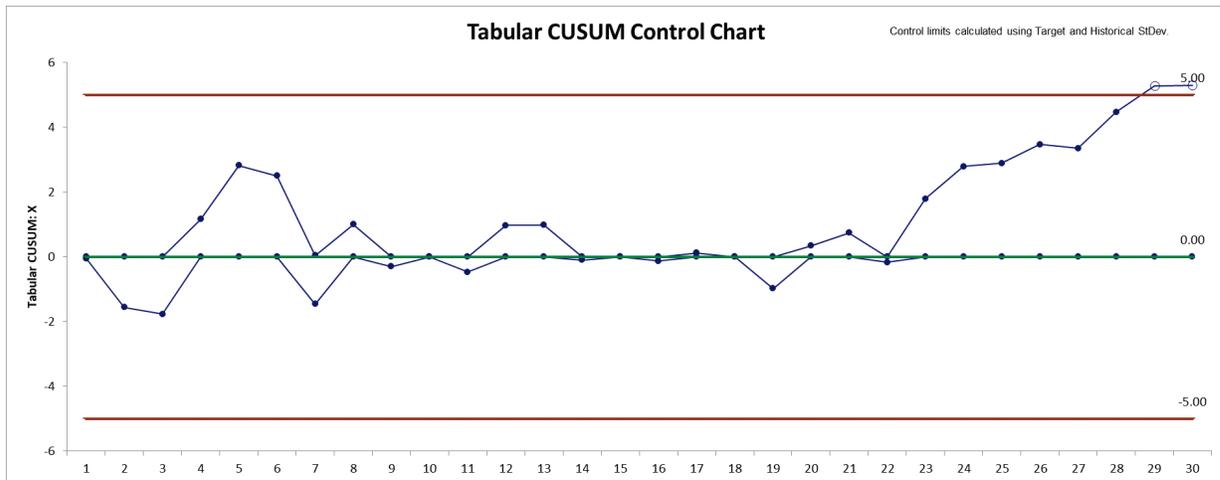
Tabular CUSUM Chart

Add Data

Notes:

1. This Tabular
2. You can re
3. Enter label
4. Enter the C
5. Enter the C

8. Click the **Add Data** button to add this data to the Tabular CUSUM Control Chart:



9. This confirms that the process is now “out-of-control” with signals at samples 29 and 30.

Template Notes and Reference:

1. This Tabular Cumulative Sum (CUSUM) Control Chart template should be used with continuous data. The data must be in chronological time-sequence order.
2. You can replace the **X-Axis Label** and **Data** column headings with any headings that you wish. Enter your data in the **Data** column.
3. Enter labels in the **X-Axis Label** column. Labels can be Date, Time, Name, or other text information. These labels are optional and will appear on the horizontal X-Axis of the Tabular CUSUM Control Chart.
4. Enter the CUSUM **k** parameter in cell D1. This is the reference (or slack) value, typically set to 0.5. It sets the size of mean shift ($2k$ sigma) that you would like to detect quickly, so 0.5 denotes rapid detection of a shift in mean = 1 sigma.
5. Enter the CUSUM **h** parameter in cell D2. This is the decision interval, typically set to 4 or 5. The upper and lower control limits = $\pm h \cdot \text{StDev}$ ($\text{MR-bar}/d2$). The center line = 0.
6. Enter the CUSUM **FIR** parameter in cell D3. This is the fast initial response (or headstart) value, typically set to $h/2$ if used, 0 otherwise. This sets the initial CUSUM statistic so that it improves the sensitivity to a mean shift at startup.
7. Optionally enter the CUSUM **Target** in cell D4. This is your process target value, typically the midpoint of your specification limits or historical mean. If you do not specify a Target, the data average will be used.
8. Optionally enter the CUSUM **Historical StDev** in cell D5. If you do not specify a Historical StDev, it will be estimated using $\text{MR-bar}/d2$.
9. The **h**, **k** and **FIR** parameters affect the Average Run Length (ARL) characteristics. For example, $h=4$ will detect a small shift more quickly than $h=5$, but has a shorter ARL(0) run length (higher false alarm rate).
10. To determine optimal CUSUM parameter values and calculate ARL, click **SigmaXL > Templates and Calculators > Control Chart Templates > Average Run Length (ARL) Calculators > CUSUM ARL**.

11. Click the **Tabular CUSUM Chart** button to create a Tabular CUSUM Control Chart.
12. After the control chart has been created and additional new data entered into the **Data** column, click the **Add Data** button to add the data to the existing chart. Control limits will be calculated using the original chart mean and stdev or specified Target and Historical StDev.
13. **Add Data** should only be used if there are at least 20 observations in the original chart, or Target and Historical StDev have been specified.
14. The **k**, **h** and **FIR** parameters are dynamic. If they are modified, the chart will automatically update with the new parameters. However, they should be selected prior to creating the Tabular CUSUM chart. Data values and out-of-control formatting are refreshed only when the buttons are used.
15. Reference: Montgomery, D.C. (2013), *Introduction to Statistical Quality Control*, Seventh Ed., Wiley, pp. 418-427.

Part I – Control Chart Templates: Trend

Trend Chart - Introduction

The Trend Control Chart template should be used with continuous data. The data must be in chronological time-sequence order and has a consistent positive or negative linear trend which is inherent to the process. This is also known as a Toolwear Control Chart.

Center line (CL) is the linear regression equation. Stdev is estimated using $\overline{MR}/1.128$ of regression residuals.

Note that the regression model estimation error is not included in the calculation of the control limits. R-Square should be at least 50%, preferably greater than 80%.

Trend control limit formulas are given in Provost L, Murray S., *The Health Care Data Guide: Learning from Data for Improvement*. San Francisco: Jossey-Bass, 2011, p. 254.

Alternatively, a regular Individuals Chart may be constructed on the regression model residuals.

Trend Chart - Example

1. Click **SigmaXL > Templates & Calculators > Control Chart Templates > Trend**. This template is also located at **SigmaXL > Control Charts > Control Chart Templates > Trend**.
2. Open **Trend Chart Example.xlsx (Sheet 1 tab)**. We will use Sample 1 to 20 to construct the Trend Control Chart and then add Sample 21 to 30 to this chart.

	A	B
1	Sample	X Trend
2	1	10.45
3	2	9.99
4	3	12.29
5	4	15.66
6	5	17.16
7	6	16.18
8	7	15.04
9	8	19.46
10	9	18.2
11	10	20.34
12	11	20.03
13	12	23.47
14	13	23.51
15	14	23.4
16	15	25.08
17	16	25.37
18	17	27.62
19	18	28.31
20	19	27.52
21	20	30.84
22	21	34.4
23	22	33.83
24	23	37.79
25	24	38
26	25	38.1
27	26	39.58
28	27	39.88
29	28	42.12
30	29	42.81
31	30	43.02

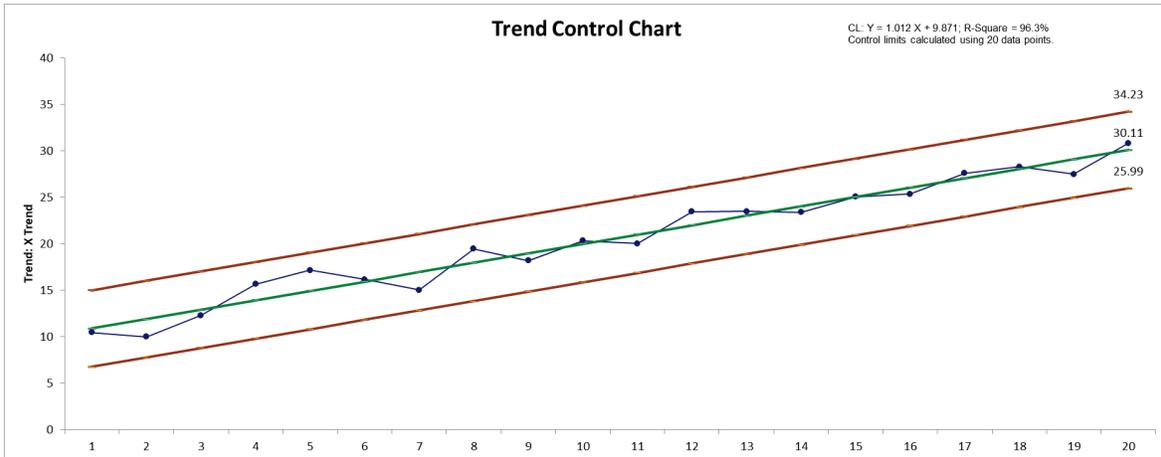
3. Copy the data in cells **A1:B21** and Paste Values to the template at **A1**.

	A	B	C	D
1	Sample	X Trend		
2	1	10.45		
3	2	9.99		
4	3	12.29		
5	4	15.66		
6	5	17.16		
7	6	16.18		
8	7	15.04		
9	8	19.46		
10	9	18.2		
11	10	20.34		
12	11	20.03		
13	12	23.47		
14	13	23.51		
15	14	23.4		
16	15	25.08		
17	16	25.37		
18	17	27.62		
19	18	28.31		
20	19	27.52		
21	20	30.84		

Trend Control Chart

Add Data

4. Click the **Trend Control Chart** button to create the Trend Control:



- This confirms that the process is “in-control” after taking the trend into account. The regression equation and R-Square are reported, with R-Square = 96.3% well above the recommended 80%.
- Switch back to **Trend Chart Example (Sheet 1 tab)**. Select and copy the data **A22: B31** as shown.

	A	B
1	Sample	X Trend
2	1	10.45
3	2	9.99
4	3	12.29
5	4	15.66
6	5	17.16
7	6	16.18
8	7	15.04
9	8	19.46
10	9	18.2
11	10	20.34
12	11	20.03
13	12	23.47
14	13	23.51
15	14	23.4
16	15	25.08
17	16	25.37
18	17	27.62
19	18	28.31
20	19	27.52
21	20	30.84
22	21	34.4
23	22	33.83
24	23	37.79
25	24	38
26	25	38.1
27	26	39.58
28	27	39.88
29	28	42.12
30	29	42.81
31	30	43.02

7. Switch to the Trend Chart template and Paste Values to cell **A23** as shown.

	A	B	C	D	E
1	Sample	X Trend			
2	1	10.45			
3	2	9.99			
4	3	12.29			
5	4	15.66			
6	5	17.16			
7	6	16.18			
8	7	15.04			
9	8	19.46			
10	9	18.2			
11	10	20.34			
12	11	20.03			
13	12	23.47			
14	13	23.51			
15	14	23.4			
16	15	25.08			
17	16	25.37			
18	17	27.62			
19	18	28.31			
20	19	27.52			
21	20	30.84			
22	21	34.4			
23	22	33.83			
24	23	37.79			
25	24	38			
26	25	38.1			
27	26	39.58			
28	27	39.88			
29	28	42.12			
30	29	42.81			
31	30	43.02			
32					

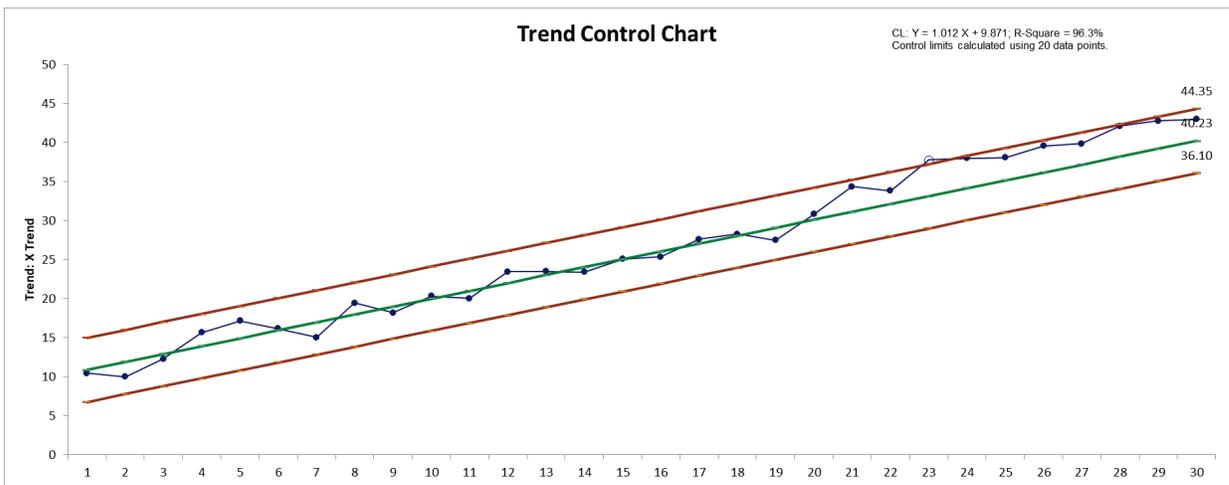
Trend Control Chart

Add Data

Notes:

1. This Trend (ak
2. You can repla
3. Enter your dat
4. Enter labels in
5. Click the Tren
6. After the cont

8. Click the **Add Data** button to add the new data to the Trend Control Chart:



9. This confirms that the process is now “out-of-control”, so corrective action to find and remove the assignable cause is needed.

Template Notes and Reference:

1. This Trend (aka Toolwear) Control Chart template should be used with continuous data. The data must be in chronological time-sequence order and has a consistent positive or negative linear trend which is inherent to the process.
2. You can replace the **X-Axis Label** and **Data** column headings with any headings that you wish.
3. Enter your data in the **Data** column.
4. Enter labels in **X-Axis Label** column. Labels can be Date, Time, Name, or other text information. These labels are optional and will appear on the horizontal X-Axis of the Trend Control Chart.
5. Click the **Trend Control Chart** button to create a Trend Control Chart.
6. After the control chart has been created and additional new data entered into the **Data** column, click the **Add Data** button to add the data to the existing chart. Control limits will be calculated using the original chart slope, intercept and stdev.
7. **Add Data** should only be used if there are at least 20 observations in the original chart.
8. Center line (CL) is the linear regression equation. Stdev. is estimated using $MR\text{-bar}/1.128$ of regression residuals.
9. Note that the regression model estimation error is not included in the calculation of the control limits. R-Square should be at least 50%, preferably greater than 80%.
10. Reference: Trend control limit formulas are given in Provost L, Murray S., *The Health Care Data Guide: Learning from Data for Improvement*. San Francisco: Jossey-Bass, 2011, p. 254.

Part J – Control Chart Templates: Average Run Length (ARL) Calculators

Average Run Length (ARL) - Introduction

Average Run Length (ARL) characteristics are very useful to compare the performance of control charts and determine optimal parameter settings for EWMA and CUSUM time weighted control charts.

The ARL value for a shift in mean = 0 sigma is the “in-control” average run length and is denoted as ARL_0 . ARL_0 is $1/\alpha$, where α is the type I false alarm probability, so this should be as large as possible minimizing the likelihood that an out-of-control signal is a false alarm. In a Shewhart Individuals Control Chart, $ARL_0 = 1/\alpha = 1/(0.00135*2) = 370.4$. On average, we will see a false alarm once every 370 observations. Note that this is a mean of a geometric distribution, so in practice the actual ARL_0 will vary widely with the standard deviation approximately equal to the mean value.

When we have a sustained shift in mean > 0 , the ARL value is the “out-of-control” run length and is denoted as ARL_1 . $ARL_1 = 1/(1-\beta)$, where β is the type II miss probability and $(1-\beta)$ is the power to detect. This should be as small as possible so that a shift in process mean is quickly detected.

The calculations for ARL are quite complex, involving either Exact, Markov Chain approximation or Monte Carlo simulation to solve. The SigmaXL ARL Templates take care of these calculations and are easy to use. If Monte Carlo simulation is used, additional Run Length standard deviation and percentile statistics are reported. Monte Carlo results will take some time and vary slightly every time they are run. With 1000 (1e3) replications it will be fast, approx. 10 seconds, but will have an ARL_0 error of approximately +/- 10%; 10,000 (1e4) replications will take about a minute, with an ARL_0 error of +/- 3.2%; 100,000 (1e5) replications will take about ten minutes, with an ARL_0 error = +/- 1%. More than 1e5 replications will be very time consuming and may run into memory limitations, so is not recommended.

When using EWMA or CUSUM charts, we typically set the parameters to minimize the ARL_1 to give rapid detection for a small shift in mean of 1 sigma. Shewhart charts are typically used when trying to rapidly detect a large shift in mean of ≥ 3 sigma. Tests for special causes may be used with Shewhart to improve the small shift performance, but they give poor ARL_0 performance resulting in frequent false alarms.

Subgroups improve the small shift performance of a Shewhart chart without impacting the ARL_0 rate and, if possible, should be used. The ARL for subgroup averages is adjusted by using the sigma of averages, σ/\sqrt{n} . For example, with a subgroup size of 4, the ARL_1 values at shift in mean = 1 will match the ARL performance of an Individuals chart with shift in mean = 2 sigma.

Note that subgroups for CUSUM and EWMA are not available in SigmaXL.

The problem of robustness to nonnormality can also be considered by using the Pearson family to specify any value of skewness and kurtosis and estimate the ARLs. For further details, see [Appendix: Pearson Family of Distributions](#).

The average run length calculators are for two-sided charts with zero-state, i.e., the shift is assumed to occur at the start. The parameters (mean, standard deviation and proportion) are also assumed to be known. This will not likely be the case in use, but is still useful for determining parameter settings and comparison of ARL across chart types.

Due to the complexity of calculations, SigmaXL must be loaded and appear on the menu in order for the ARL templates to function. Do not add or delete rows or columns in these templates.

Shewhart ARL

1. Click **SigmaXL > Templates & Calculators > Control Chart Templates > Average Run Length (ARL) Calculators > Shewhart ARL**. This template is also located at **SigmaXL > Control Charts > Control Chart Templates > Average Run Length (ARL) Calculators > Shewhart ARL**.
2. The default template settings are **Specify = Exact (Test 1 Only)**, **Subgroup Size = 1**, **Skewness = 0**, **Kurtosis (Normal is 0) = 0**.

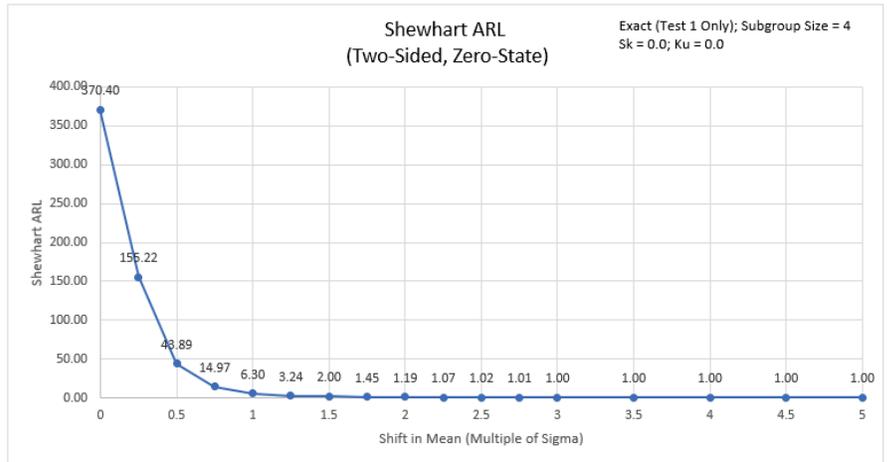
SigmaXL		Shewhart Average Run Length (ARL) Calculator	
User Inputs:			
Specify:	Exact (Test 1 Only)		Calculate Shewhart ARL
Subgroup Size:	1		
Skewness:	0		
Kurtosis (Normal is 0):	0		

Notes: Specify Exact (Test 1 Only) or Monte Carlo using the drop-down list. Parameters to be specified will be shown in yellow highlight, otherwise they are hidden. Exact uses the normal or Pearson cumulative distribution function and is fast. Monte Carlo simulation allows you to assess the ARL performance of all 8 Tests for Special Causes. Test 1 - 1 point more than 3 standard deviations from the center line (CL) is always applied. Monte Carlo simulation also produces the table of Run Length Standard Deviation and Percentiles (scroll right to view). Both methods allow you to assess robustness to nonnormality.

All ARL calculations for Shewhart use a standardized in-control mean=0 and sigma=1.

- Click the **Calculate Shewhart ARL** button to produce the ARL table and chart for these settings:

Shift in Mean (Multiple of Sigma)	Shewhart ARL (Two-Sided, Zero-State)
0	370.40
0.25	155.22
0.5	43.89
0.75	14.97
1	6.30
1.25	3.24
1.5	2.00
1.75	1.45
2	1.19
2.25	1.07
2.5	1.02
2.75	1.01
3	1.00
3.5	1.00
4	1.00
4.5	1.00
5	1.00



The ARL_0 for the Shewhart \bar{x} -bar chart is the same as the Individuals chart, 370.4. The ARL_1 for a small 1 sigma shift in mean is 6.3, so is much more rapid to detect than the Individuals ARL_1 of 43.89, so if possible, subgrouping should always be used.

Note: The ARL for subgroup averages is adjusted by using the sigma of averages, σ/\sqrt{n} . For example, with a subgroup size of 4, the ARL_1 values at shift in mean = 1 will match the ARL performance of an Individuals chart with shift in mean = 2 sigma.

- Now we will assess robustness to nonnormality. Enter **Specify = Exact (Test 1 Only)**, **Subgroup Size = 1**, **Skewness = 2**, **Kurtosis (Normal is 0) = 6**.

Shewhart Average Run Length (ARL) Calculator

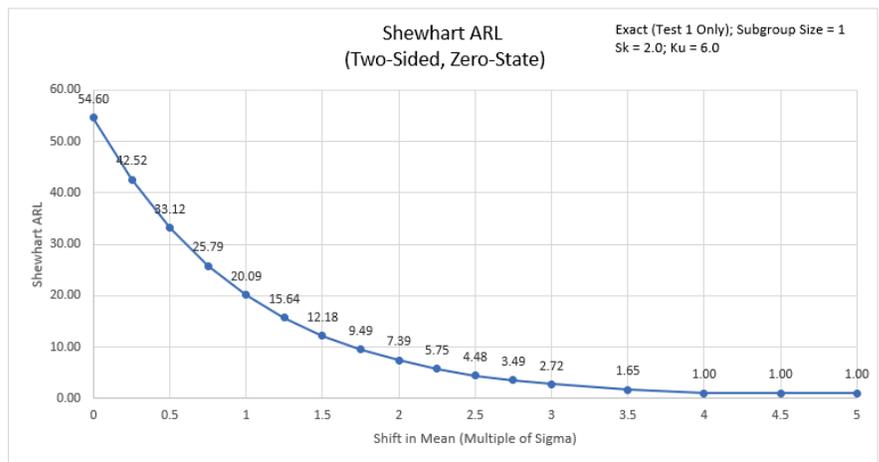
User Inputs:

Specify:	Exact (Test 1 Only)	<input type="button" value="Calculate Shewhart ARL"/>
Subgroup Size:	1	
Skewness:	2	
Kurtosis (Normal is 0):	6	

Note: We are specifying a severe degree of Skewness (Skewness = 0.5 is mild, 1 is moderate, 2 is severe, and > 2 is very severe). The Pearson family is used to create a distribution that matches the specified Skewness and Kurtosis. Skewness = 2 and Kurtosis = 6 corresponds to an Exponential distribution or Gamma distribution with $Shape = 1$ and $Scale = 1$ (for Gamma, $Skewness = 2/\sqrt{Shape}$ and $Kurtosis = 6/Shape$).

7. Click the **Calculate Shewhart ARL** button to produce the ARL table and chart for these settings:

Shift in Mean (Multiple of Sigma)	Shewhart ARL (Two-Sided, Zero-State)
0	54.60
0.25	42.52
0.5	33.12
0.75	25.79
1	20.09
1.25	15.64
1.5	12.18
1.75	9.49
2	7.39
2.25	5.75
2.5	4.48
2.75	3.49
3	2.72
3.5	1.65
4	1.00
4.5	1.00
5	1.00



ARL₀ with these settings is 54.6. This is a very poor performance with a 6.8 x increase (370.4/54.6) in false alarms compared to normal data. The Shewhart Individuals chart is not robust to severe skewness. A Box-Cox Transformation or other Individuals Nonnormal chart should be used (see **SigmaXL > Control Charts > Nonnormal > Individuals Nonnormal**).

Note: ARL₀ = 54.6 matches the result given in Montgomery [2], Table 9.12 for Gam(1,1).

8. Next, we will assess robustness to nonnormality for a Shewhart X-bar chart. Enter **Specify = Exact (Test 1 Only)**, **Subgroup Size = 4**, **Skewness = 2**, **Kurtosis (Normal is 0) = 6**.

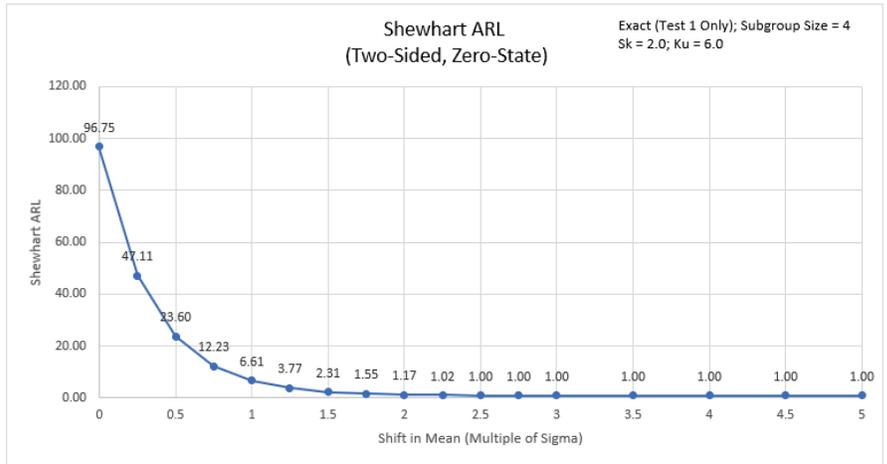
User Inputs:	
Specify:	Exact (Test 1 Only)
Subgroup Size:	4
Skewness:	2
Kurtosis (Normal is 0):	6

Calculate Shewhart ARL

Note: Skewness of averages = $Skewness/\sqrt{n}$. Kurtosis of averages = $Kurtosis/n$. For a subgroup size of 4, the skewness of averages is 1, so is reduced from severe to moderate. Kurtosis of averages is 1.5 (corresponding to a Gamma distribution with *Shape* = 4).

9. Click the **Calculate Shewhart ARL** button to produce the ARL table and chart for these settings:

Shift in Mean (Multiple of Sigma)	Shewhart ARL (Two-Sided, Zero-State)
0	96.75
0.25	47.11
0.5	23.60
0.75	12.23
1	6.61
1.25	3.77
1.5	2.31
1.75	1.55
2	1.17
2.25	1.02
2.5	1.00
2.75	1.00
3	1.00
3.5	1.00
4	1.00
4.5	1.00
5	1.00



ARL₀ with these settings is 96.75. This is an improvement over the Individuals 54.6, but is still a 3.8 x increase (370.4/96.75) in false alarms compared to normal data.

Note: ARL₀ = 96.75 matches the results given in Schilling & Nelson [3] (Table 1, Gamma, shape = 1, n=4), =1/.01034. In Table 2, they point out that a subgroup size of 166 would be required to achieve robustness for this severe skewness.

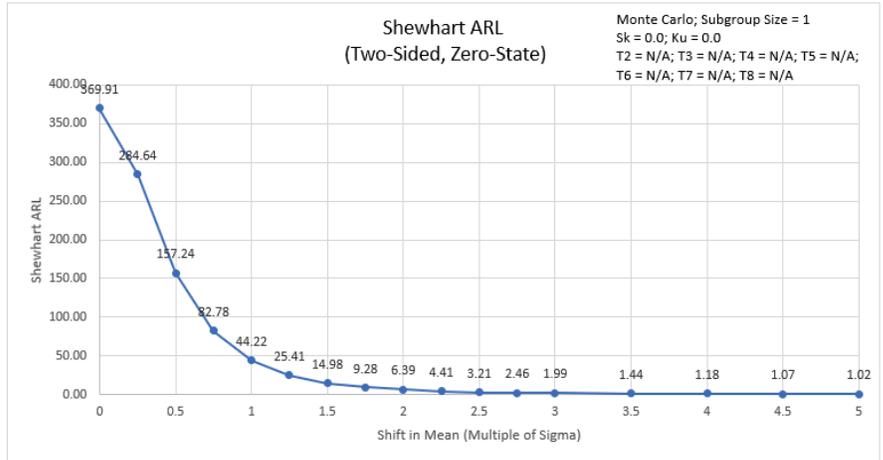
10. Now we will use Monte Carlo simulation to obtain approximate Run Length standard deviation and percentiles for an Individuals Shewhart chart with normal data. Enter **Specify = Monte Carlo**, **Subgroup Size = 1**, **Skewness = 0**, **Kurtosis (Normal is 0) = 0**, **Number of Replications = 1e4**, and **Test 2 to Test 8 = N/A**.

SigmaXL Shewhart Average Run Length (ARL) Calculator	
User Inputs:	
Specify:	Monte Carlo
Subgroup Size:	1
Skewness:	0
Kurtosis (Normal is 0):	0
Number of Replications:	1.00E+04
Test 2 - points in a row on same side of CL:	N/A
Test 3 - points in a row all increasing or decreasing:	N/A
Test 4 - points in a row alternating up and down:	N/A
Test 5 - points more than 2 standard deviations from CL (same side):	N/A
Test 6 - points more than 1 standard deviation from CL (same side):	N/A
Test 7 - points in a row within 1 standard deviation from CL (either side):	N/A
Test 8 - points in a row more than 1 standard deviation from CL (either side):	N/A

Calculate Shewhart ARL

11. Click the **Calculate Shewhart ARL** button to produce the Monte Carlo approximate ARL table, ARL chart and Run Length Standard Deviation and Percentiles table (scroll right to view). Monte Carlo simulation with 10,000 (1e4) replications will take about a minute to run.

Shift in Mean (Multiple of Sigma)	Shewhart ARL (Two-Sided, Zero-State)
0	369.91
0.25	284.64
0.5	157.24
0.75	82.78
1	44.22
1.25	25.41
1.5	14.98
1.75	9.28
2	6.39
2.25	4.41
2.5	3.21
2.75	2.46
3	1.99
3.5	1.44
4	1.18
4.5	1.07
5	1.02



Monte Carlo Simulation Run Length Standard Deviation and Percentiles (Two-Sided, Zero-State)										
Shift in Mean (Multiple of Sigma)	Standard Deviation	1	5th	10th	25th	50th (Median)	75th	90th	95th	99th
0	370.23	5	19	37	105	257	515	863	1106	1738
0.25	284.79	4	16	30	82	199	391	664	851	1274
0.5	157.51	2	9	17	45	111	216	357	469	737
0.75	83.00	1	4	9	24	58	113	190	247	393
1	43.44	1	3	5	13	31	61	101	132	202
1.25	24.89	1	2	3	8	18	35	58	75	114
1.5	14.36	1	1	2	5	11	20	34	44	67
1.75	8.86	1	1	1	3	7	13	20	27	41
2	5.70	1	1	1	2	5	9	14	18	26
2.25	3.97	1	1	1	2	3	6	10	12	19
2.5	2.67	1	1	1	1	2	4	7	8	13
2.75	1.89	1	1	1	1	2	3	5	6	9
3	1.40	1	1	1	1	1	2	4	5	7
3.5	0.79	1	1	1	1	1	2	2	3	4
4	0.47	1	1	1	1	1	1	2	2	3
4.5	0.29	1	1	1	1	1	1	1	2	2
5	0.16	1	1	1	1	1	1	1	1	2

The additional run length statistics show the large variation of run length values. The median $MRL_0 = 257$ (in-control median run length with 0 shift in process mean). The run length percentiles approximately match those given in Chakraborti [4] (Table 1, Standards Known, Shift 0.0).

Note: The results will vary slightly since this is Monte Carlo simulation.

12. Now we will use Monte Carlo simulation to assess the Shewhart Individuals chart with all 8 Tests for Special Causes applied. Enter **Specify = Monte Carlo**, **Subgroup Size = 1**, **Skewness = 0**, **Kurtosis (Normal is 0) = 0**, **Number of Replications = 1e4**, **Test 2 = 9**, **Test 3 = 6**, **Test 4 = 14**, **Test 5 = 2 out of 3**, **Test 6 = 4 out of 5**, **Test 7 = 15** and **Test 8 = 8**.

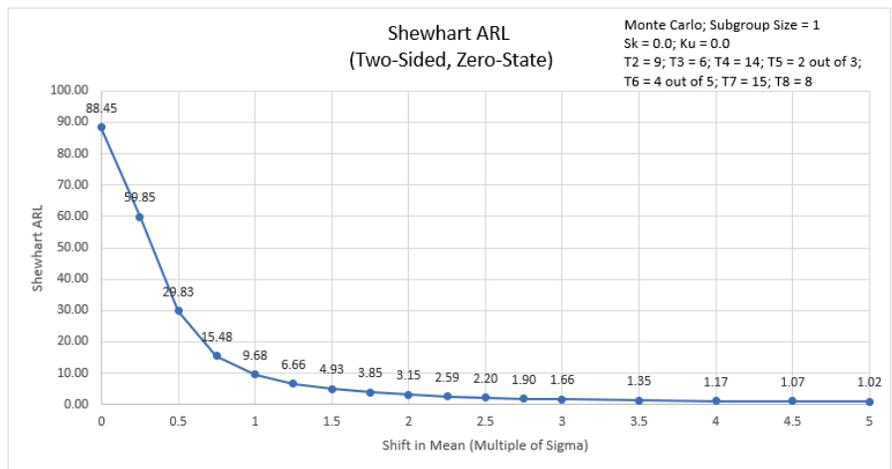
User Inputs:	
Specify:	Monte Carlo
Subgroup Size:	1
Skewness:	0
Kurtosis (Normal is 0):	0
Number of Replications:	1.00E+04
Test 2 - points in a row on same side of CL:	9
Test 3 - points in a row all increasing or decreasing:	6
Test 4 - points in a row alternating up and down:	14
Test 5 - points more than 2 standard deviations from CL (same side):	2 out of 3
Test 6 - points more than 1 standard deviation from CL (same side):	4 out of 5
Test 7 - points in a row within 1 standard deviation from CL (either side):	15
Test 8 - points in a row more than 1 standard deviation from CL (either side):	8

Calculate Shewhart ARL

Note: These are the test settings used as defaults in **SigmaXL > Control Charts > 'Tests for Special Causes' Defaults**. Test 1 is always applied.

13. Click the **Calculate Shewhart ARL** button to produce the Monte Carlo approximate ARL table, ARL chart and Run Length Standard Deviation and Percentiles table.

Shift in Mean (Multiple of Sigma)	Shewhart ARL (Two-Sided, Zero-State)
0	88.45
0.25	59.85
0.5	29.83
0.75	15.48
1	9.68
1.25	6.66
1.5	4.93
1.75	3.85
2	3.15
2.25	2.59
2.5	2.20
2.75	1.90
3	1.66
3.5	1.35
4	1.17
4.5	1.07
5	1.02



Shift in Mean (Multiple of Sigma)	Monte Carlo Simulation Run Length Standard Deviation and Percentiles (Two-Sided, Zero-State)									
	Standard Deviation	1	5th	10th	25th	50th (Median)	75th	90th	95th	99th
0	83.72	3	8	14	29	63	120	197	254	393
0.25	55.52	3	7	10	21	43	82	132	168	260
0.5	26.21	2	5	6	11	22	40	64	82	124
0.75	12.10	1	3	4	7	12	20	31	39	58
1	6.65	1	2	3	5	8	12	18	23	33
1.25	4.21	1	2	2	4	5	9	12	15	21
1.5	2.85	1	1	2	3	4	6	9	10	14
1.75	2.09	1	1	1	2	4	5	6	8	10
2	1.67	1	1	1	2	3	4	5	6	9
2.25	1.35	1	1	1	2	2	3	4	5	7
2.5	1.12	1	1	1	1	2	3	4	4	5
2.75	0.94	1	1	1	1	2	2	3	4	5
3	0.80	1	1	1	1	1	2	3	3	4
3.5	0.56	1	1	1	1	1	2	2	2	3
4	0.40	1	1	1	1	1	1	2	2	2
4.5	0.26	1	1	1	1	1	1	1	2	2
5	0.15	1	1	1	1	1	1	1	1	2

ARL₀ with all 8 tests for special causes is approx. 88.5. This is a poor performance with a 4.2 x increase (370.4/88.5) in false alarms compared to Test 1 only. MRL₀ is approx. 63. On the other hand, ARL₁ for a small 1 sigma shift in mean is approx. 9.7, so is much faster to detect than the Exact Test 1 only ARL₁ of 43.89.

If small shifts are to be detected quickly and subgrouping is not possible, then an EWMA or CUSUM chart is recommended.

Template Notes:

1. Specify **Exact (Test 1 Only)** or **Monte Carlo** using the drop-down list. Parameters to be specified will be shown in yellow highlight, otherwise they are hidden.
2. Exact uses the cumulative distribution function and is fast. Monte Carlo simulation allows you to assess the ARL performance of all 8 Tests for Special Causes and also produces the table of Run Length Standard Deviation and Percentiles (scroll right to view). Both methods allow you to assess robustness to nonnormality.
3. Test 1 - 1 point more than 3 standard deviations from the center line (CL) is always applied.
4. The Pearson Family of distributions is used to match the specified Skewness and Kurtosis.
5. Enter the **Subgroup Size**. Subgroup size = 1 denotes a Shewhart Individuals chart. Subgroup size > 1 is an X-Bar chart.
6. Enter **Skewness**. Skewness = 0 is symmetric.
7. Enter **Kurtosis (Normal is 0)**. Also known as Excess Kurtosis, it must be $\geq \text{Skewness}^2 - 1.48$. This is required to keep the distribution unimodal. If Skewness=0 and Kurtosis = 0, the distribution is normal.
8. If applicable, enter **Number of Replications**. 1000 (1e3) replications will be fast, approx. 10 seconds, but will have an ARL0 error approx. = +/- 10%; 10,000 (1e4) replications will take about a minute, with an ARL0 error = +/- 3.2%; 100,000 (1e5) replications will take about ten minutes, with an ARL0 error = +/- 1%.
9. If applicable, select values for Tests 2 to 8 using the drop-down list. "N/A" indicates that the test is not applied. Tests 2, 3 and 7 provide options that match those provided in SigmaXL's '**Tests for Special Causes' Defaults** dialog.
10. Click the **Calculate Shewhart ARL** button to produce the ARL table and chart. If Monte Carlo was selected, the table of Run Length Standard Deviation and Percentiles will also be produced.
11. The Shewhart ARL is for a two-sided chart with zero-state, i.e., the shift is assumed to occur at the start. The mean and standard deviation are also assumed to be known. This will not likely be the case in use, but is still useful for determining parameter settings and comparison of ARL across chart types.
12. Due to the complexity of calculations, SigmaXL must be loaded and appear on the menu in order for this template to function. Do not add or delete rows or columns in this template.

REFERENCES:

- [1] Champ, C.W. and Woodall, W.H. (1987), "Exact results for Shewhart control charts with supplementary runs rules", *Technometrics* 29, 393-399.
- [2] Montgomery, D.C. (2013), *Introduction to Statistical Quality Control*, Seventh Ed., Wiley.
- [3] Schilling, E. G., and P. R. Nelson (1976), "The Effect of Nonnormality on the Control Limits of \bar{X} Charts," *Journal of Quality Technology*, Vol. 8(4), pp. 183–188.
- [4] Chakraborti, S. (2007), "Run Length Distribution and Percentiles: The Shewhart Chart with Unknown Parameters", *Quality Engineering* 19, 119–127.

Attribute P ARL

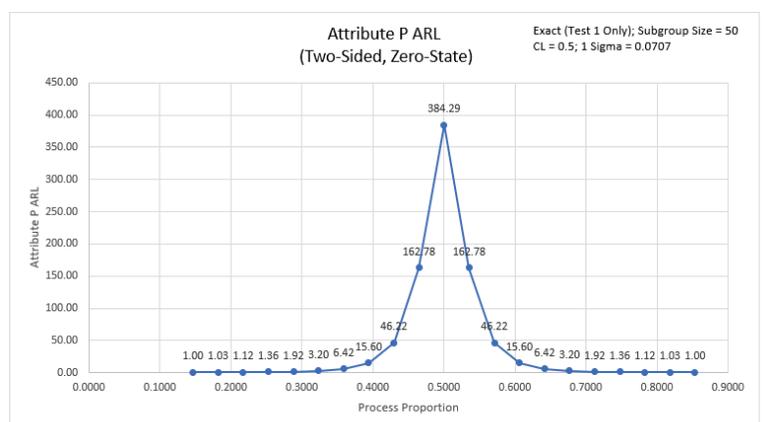
1. Click **SigmaXL > Templates & Calculators > Control Chart Templates > Average Run Length (ARL) Calculators > Attribute P ARL**. This template is also located at **SigmaXL > Control Charts > Control Chart Templates > Average Run Length (ARL) Calculators > Attribute P ARL**.
2. The default template settings are **Specify = Exact (Test 1 Only)**, **Subgroup Size = 50**, **In-Control Historical Proportion (CL) = 0.5**.

Attribute P Average Run Length (ARL) Calculator		
User Inputs:		
Specify:	Exact (Test 1 Only)	Calculate Attribute P ARL
Subgroup Size:	50	
In-Control Historical Proportion (CL):	0.5	

Note: Specify Exact (Test 1 Only) or Monte Carlo using the drop-down list. Parameters to be specified will be shown in yellow highlight, otherwise they are hidden. Exact uses the binomial cumulative distribution function. Monte Carlo simulation uses binomial random data with specified proportion and allows you to assess the ARL performance of all 4 Tests for Special Causes. Test 1 - 1 point more than 3 standard deviations from the center line (CL) is always applied. Monte Carlo simulation also produces the table of Run Length Standard Deviation and Percentiles (scroll right to view).

3. Click the **Calculate Attribute P ARL** button to reproduce the ARL table and chart.

Shift in Proportion (Multiple of Sigma) 1 Sigma = 0.0707	Process Proportion	Attribute P ARL (Two-Sided, Zero-State)
-5.0	0.1464	1.00
-4.5	0.1818	1.03
-4.0	0.2172	1.12
-3.5	0.2525	1.36
-3.0	0.2879	1.92
-2.5	0.3232	3.20
-2.0	0.3586	6.42
-1.5	0.3939	15.60
-1.0	0.4293	46.22
-0.5	0.4646	162.78
0.0	0.5000	384.29
0.5	0.5354	162.78
1.0	0.5707	46.22
1.5	0.6061	15.60
2.0	0.6414	6.42
2.5	0.6768	3.20
3.0	0.7121	1.92
3.5	0.7475	1.36
4.0	0.7828	1.12
4.5	0.8182	1.03
5.0	0.8536	1.00



The ARL_0 (in-control ARL with 0 sigma shift in proportion) for the Attribute P chart is 384.29. The ARL_1 for a small 1 sigma shift is 46.22, so is slow to detect. On the other hand, a large 3 sigma shift in proportion has an $ARL = 1.92$, so is detected rapidly.

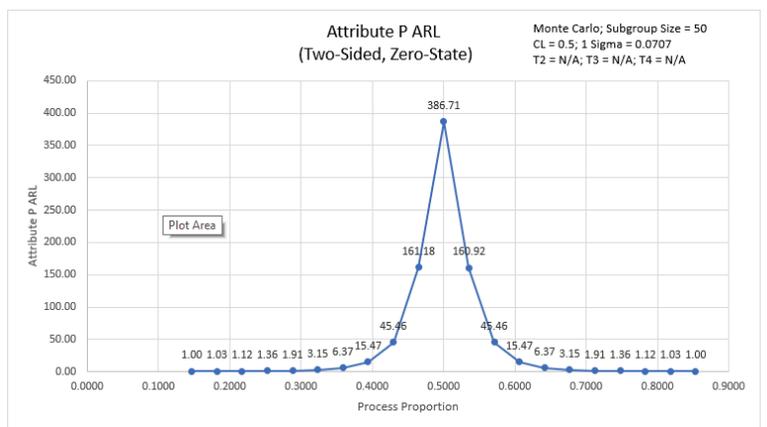
Note: The Process Proportion values are displayed on the ARL chart, but the shift in proportion as a multiple of sigma is also included in the table.

- Now we will use Monte Carlo simulation to obtain approximate Run Length standard deviation and percentiles for the Attribute P chart. Select **Specify** = *Monte Carlo*. Enter **Subgroup Size** = 50, **In-Control Historical Proportion (CL)** = 0.5, **Number of Replications** = 1e4, **Test 2** = N/A, **Test 3** = N/A, **Test 4** = N/A.

Attribute P Average Run Length (ARL) Calculator	
User Inputs:	
Specify:	Monte Carlo
Subgroup Size:	50
In-Control Historical Proportion (CL):	0.5
Number of Replications:	1.00E+04
Test 2 - points in a row on same side of CL:	N/A
Test 3 - points in a row all increasing or decreasing:	N/A
Test 4 - points in a row alternating up and down:	N/A

- Click the **Calculate Attribute P ARL** button to produce the Monte Carlo approximate ARL table, ARL chart and Run Length Standard Deviation and Percentiles table (scroll right to view). Monte Carlo simulation with 10,000 (1e4) replications will take about a minute to run.

Shift in Proportion (Multiple of Sigma) 1 Sigma = 0.0707	Process Proportion	Attribute P ARL (Two-Sided, Zero-State)
-5.0	0.1464	1.00
-4.5	0.1818	1.03
-4.0	0.2172	1.12
-3.5	0.2525	1.36
-3.0	0.2879	1.91
-2.5	0.3232	3.15
-2.0	0.3586	6.37
-1.5	0.3939	15.47
-1.0	0.4293	45.46
-0.5	0.4646	161.18
0.0	0.5000	386.71
0.5	0.5354	160.92
1.0	0.5707	45.46
1.5	0.6061	15.47
2.0	0.6414	6.37
2.5	0.6768	3.15
3.0	0.7121	1.91
3.5	0.7475	1.36
4.0	0.7828	1.12
4.5	0.8182	1.03
5.0	0.8536	1.00



Shift in Proportion (Multiple of Sigma) 1 Sigma = 0.0707	Process Proportion	Standard Deviation	Monte Carlo Simulation Run Length Standard Deviation and Percentiles (Two-Sided, Zero-State)									
			1	5th	10th	25th	50th (Median)	75th	90th	95th	99th	
-5.0	0.1464	0.07	1	1	1	1	1	1	1	1	1	1
-4.5	0.1818	0.18	1	1	1	1	1	1	1	1	1	2
-4.0	0.2172	0.36	1	1	1	1	1	1	1	2	2	3
-3.5	0.2525	0.69	1	1	1	1	1	1	2	2	3	4
-3.0	0.2879	1.32	1	1	1	1	1	1	2	4	5	7
-2.5	0.3232	2.61	1	1	1	1	1	2	4	6	8	12
-2.0	0.3586	5.86	1	1	1	2	5	9	14	18	28	28
-1.5	0.3939	14.85	1	1	2	5	11	21	35	45	70	70
-1.0	0.4293	45.38	1	3	5	13	31	63	105	138	205	205
-0.5	0.4646	161.29	2	8	17	46	113	224	372	480	746	746
0.0	0.5000	384.33	4	19	41	113	267	540	899	1165	1744	1744
0.5	0.5354	162.28	2	8	17	46	111	222	368	480	760	760
1.0	0.5707	45.38	1	3	5	13	31	63	105	138	205	205
1.5	0.6061	14.85	1	1	2	5	11	21	35	45	70	70
2.0	0.6414	5.86	1	1	1	2	5	9	14	18	28	28
2.5	0.6768	2.61	1	1	1	1	2	4	6	8	12	12
3.0	0.7121	1.32	1	1	1	1	1	2	4	5	7	7
3.5	0.7475	0.69	1	1	1	1	1	2	2	3	4	4
4.0	0.7828	0.36	1	1	1	1	1	1	2	2	3	3
4.5	0.8182	0.18	1	1	1	1	1	1	1	1	2	2
5.0	0.8536	0.07	1	1	1	1	1	1	1	1	1	1

The additional run length statistics show the large variation of run length values. The $MRL_0 = 267$ (in-control median run length with 0 sigma shift in process proportion).

Note: The results will vary slightly since this is Monte Carlo simulation.

- Now we will use Monte Carlo simulation to assess the Attribute P chart with all 4 Tests for Special Causes applied. Enter **Specify = Monte Carlo, Subgroup Size = 50, In-Control Historical Proportion (CL) = 0.5, Number of Replications = 1e4, Test 2 = 9, Test 3 = 6, and Test 4 = 14.**

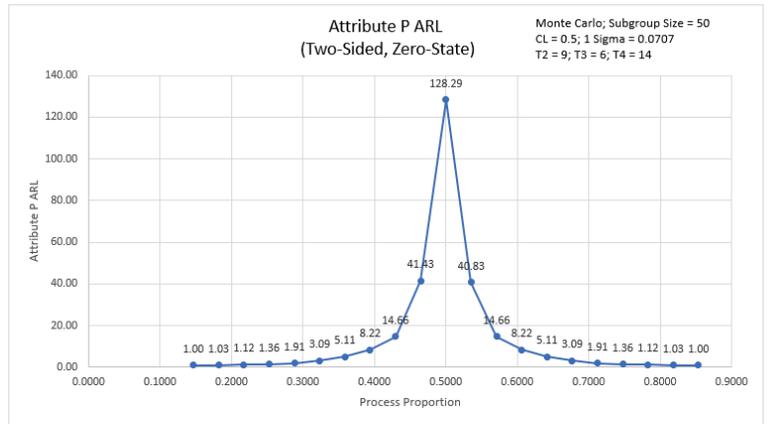
Attribute P Average Run Length (ARL) Calculator	
User Inputs:	
Specify:	Monte Carlo
Subgroup Size:	50
In-Control Historical Proportion (CL):	0.5
Number of Replications:	1.00E+04
Test 2 - points in a row on same side of CL:	9
Test 3 - points in a row all increasing or decreasing:	6
Test 4 - points in a row alternating up and down:	14

Calculate Attribute P ARL

Note: These are the test settings used as defaults in SigmaXL > Control Charts > 'Tests for Special Causes' Defaults. Test 1 is always applied.

- Click the **Calculate Attribute P ARL** button to produce the Monte Carlo approximate ARL table, ARL chart and Run Length Standard Deviation and Percentiles table:

Shift in Proportion (Multiple of Sigma) 1 Sigma = 0.0707	Process Proportion	Attribute P ARL (Two-Sided, Zero-State)
-5.0	0.1464	1.00
-4.5	0.1818	1.03
-4.0	0.2172	1.12
-3.5	0.2525	1.36
-3.0	0.2879	1.91
-2.5	0.3232	3.09
-2.0	0.3586	5.11
-1.5	0.3939	8.22
-1.0	0.4293	14.66
-0.5	0.4646	41.43
0.0	0.5000	128.29
0.5	0.5354	40.83
1.0	0.5707	14.66
1.5	0.6061	8.22
2.0	0.6414	5.11
2.5	0.6768	3.09
3.0	0.7121	1.91
3.5	0.7475	1.36
4.0	0.7828	1.12
4.5	0.8182	1.03
5.0	0.8536	1.00



Shift in Proportion (Multiple of Sigma) 1 Sigma = 0.0707	Process Proportion	Monte Carlo Simulation Run Length Standard Deviation and Percentiles (Two-Sided, Zero-State)									
		Standard Deviation	1	5th	10th	25th	50th (Median)	75th	90th	95th	99th
-5.0	0.1464	0.06	1	1	1	1	1	1	1	1	1
-4.5	0.1818	0.18	1	1	1	1	1	1	1	1	2
-4.0	0.2172	0.37	1	1	1	1	1	1	2	2	3
-3.5	0.2525	0.71	1	1	1	1	1	2	2	3	4
-3.0	0.2879	1.29	1	1	1	1	1	2	4	4	7
-2.5	0.3232	2.27	1	1	1	1	2	4	7	8	9
-2.0	0.3586	3.25	1	1	1	2	5	9	9	9	13
-1.5	0.3939	4.62	1	1	2	5	9	9	14	17	23
-1.0	0.4293	10.00	1	3	5	9	11	18	27	34	50
-0.5	0.4646	35.49	2	9	9	16	31	56	88	112	168
0.0	0.5000	122.77	4	12	19	41	91	176	287	376	568
0.5	0.5354	34.95	2	8	9	16	30	55	87	111	164
1.0	0.5707	10.00	1	3	5	9	11	18	27	34	50
1.5	0.6061	4.62	1	1	2	5	9	9	14	17	23
2.0	0.6414	3.25	1	1	1	2	5	9	9	9	13
2.5	0.6768	2.27	1	1	1	1	2	4	7	8	9
3.0	0.7121	1.29	1	1	1	1	1	2	4	4	7
3.5	0.7475	0.71	1	1	1	1	1	2	2	3	4
4.0	0.7828	0.37	1	1	1	1	1	1	2	2	3
4.5	0.8182	0.18	1	1	1	1	1	1	1	1	2
5.0	0.8536	0.06	1	1	1	1	1	1	1	1	1

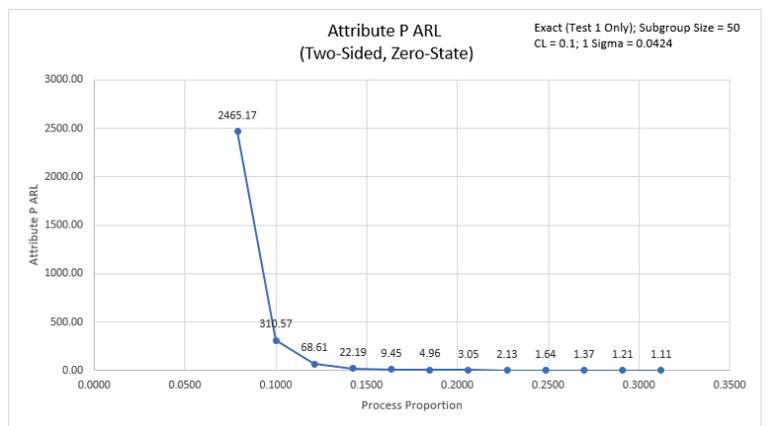
ARL₀ with all 4 tests for special causes is approx. 128.3. This is a poor performance with a 3 x increase (384.3/128.3) in false alarms compared to Exact Test 1 only. MRL₀ is approx. 91. On the other hand, ARL₁ for a small 1 sigma shift in proportion is approx. 14.7, so is much faster to detect than the Exact Test 1 only ARL₁ of 46.22.

8. We will now assess ARL for a low in-control proportion value. Enter **Specify = Exact (Test 1 Only)**, **Subgroup Size = 50**, **In-Control Historical Proportion (CL) = 0.1**.

Attribute P Average Run Length (ARL) Calculator	
User Inputs:	
Specify:	Exact (Test 1 Only)
Subgroup Size:	50
In-Control Historical Proportion (CL):	0.1

9. Click the **Calculate Attribute P ARL** button to produce the ARL table and chart for these settings:

Shift in Proportion (Multiple of Sigma) 1 Sigma = 0.0424	Process Proportion	Attribute P ARL (Two-Sided, Zero-State)
-0.5	0.0788	2465.17
0.0	0.1000	310.57
0.5	0.1212	68.61
1.0	0.1424	22.19
1.5	0.1636	9.45
2.0	0.1849	4.96
2.5	0.2061	3.05
3.0	0.2273	2.13
3.5	0.2485	1.64
4.0	0.2697	1.37
4.5	0.2909	1.21
5.0	0.3121	1.11



The ARL_0 for the Attribute P chart is 310.57. The ARL_1 for a small positive 1 sigma shift in proportion is 22.19, so is slow to detect. On the other hand, a large positive 3 sigma shift in proportion has an $ARL = 2.13$, so is detected rapidly. The ARL for a small negative 0.5 sigma shift in proportion is 2465.7 so negative shifts cannot be detected with these settings.

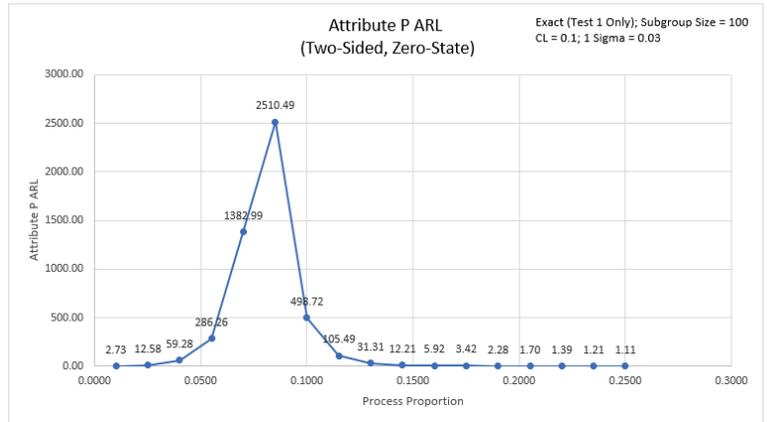
Note: The P Chart LCL is 0 for this CL value.

10. Now we will use Monte Carlo simulation to assess ARL for a low in-control proportion value but with all 4 Tests for Special Causes applied. Enter **Specify = Monte Carlo**, **Subgroup Size = 50**, **In-Control Historical Proportion (CL) = 0.1**, **Number of Replications = 1e4**, **Test 2 = 9**, **Test 3 = 6**, and **Test 4 = 14**.

Attribute P Average Run Length (ARL) Calculator	
User Inputs:	
Specify:	Monte Carlo
Subgroup Size:	50
In-Control Historical Proportion (CL):	0.1
Number of Replications:	1.00E+04
Test 2 - points in a row on same side of CL:	9
Test 3 - points in a row all increasing or decreasing:	6
Test 4 - points in a row alternating up and down:	14

13. Click the **Calculate Attribute P ARL** button to produce the ARL table and chart for these settings:

Shift in Proportion (Multiple of Sigma) 1 Sigma = 0.03	Process Proportion	Attribute P ARL (Two-Sided, Zero-State)
-3.0	0.0100	2.73
-2.5	0.0250	12.58
-2.0	0.0400	59.28
-1.5	0.0550	286.26
-1.0	0.0700	1382.99
-0.5	0.0850	2510.49
0.0	0.1000	498.72
0.5	0.1150	105.49
1.0	0.1300	31.31
1.5	0.1450	12.21
2.0	0.1600	5.92
2.5	0.1750	3.42
3.0	0.1900	2.28
3.5	0.2050	1.70
4.0	0.2200	1.39
4.5	0.2350	1.21
5.0	0.2500	1.11



The ARL_0 for the Attribute P chart is 498.72. The ARL_1 for a small positive 1 sigma shift in mean is 31.31, so is slow to detect. On the other hand, a large positive 3 sigma shift in proportion has an $ARL = 2.28$, so is detected rapidly. Small negative shifts still cannot be detected, but a large negative 3 sigma shift in proportion can be detected with $ARL = 2.73$.

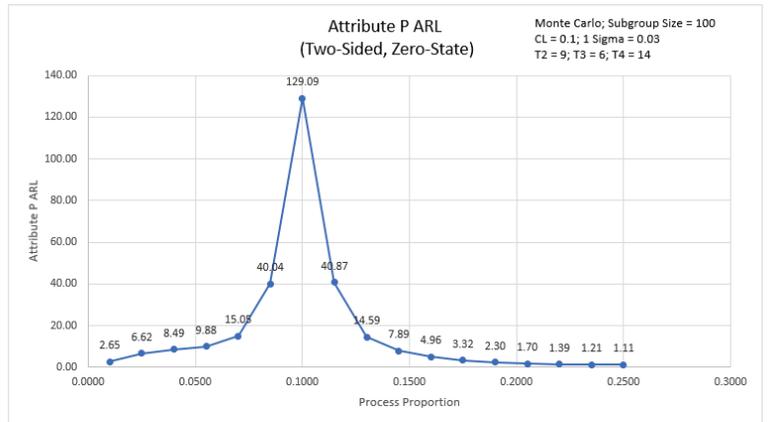
14. Now we will use Monte Carlo simulation to assess ARL for a low in-control proportion value with subgroup size = 100, but with all 4 Tests for Special Causes applied. Enter **Specify = Monte Carlo**, **Subgroup Size = 100**, **In-Control Historical Proportion (CL) = 0.1**. **Number of Replications = 1e4**, **Test 2 = 9**, **Test 3 = 6**, and **Test 4 = 14**.

Attribute P Average Run Length (ARL) Calculator	
User Inputs:	
Specify:	Monte Carlo
Subgroup Size:	100
In-Control Historical Proportion (CL):	0.1
Number of Replications:	1.00E+04
Test 2 - points in a row on same side of CL:	9
Test 3 - points in a row all increasing or decreasing:	6
Test 4 - points in a row alternating up and down:	14

Calculate Attribute P ARL

15. Click the **Calculate Attribute P ARL** button to produce the ARL table and chart for these settings:

Shift in Proportion (Multiple of Sigma) 1 Sigma = 0.03	Process Proportion	Attribute P ARL (Two-Sided, Zero-State)
-3.0	0.0100	2.65
-2.5	0.0250	6.62
-2.0	0.0400	8.49
-1.5	0.0550	9.88
-1.0	0.0700	15.05
-0.5	0.0850	40.04
0.0	0.1000	129.09
0.5	0.1150	40.87
1.0	0.1300	14.59
1.5	0.1450	7.89
2.0	0.1600	4.96
2.5	0.1750	3.32
3.0	0.1900	2.30
3.5	0.2050	1.70
4.0	0.2200	1.39
4.5	0.2350	1.21
5.0	0.2500	1.11



ARL₀ with all 4 tests for special causes is approx. 129.1. This is a poor performance with a 3.9 x increase (498.7/129.1) in false alarms compared to Exact Test 1 only. On the other hand, ARL₁ for a small 1 sigma shift in proportion is approx. 14.6, so is much faster to detect than the Exact Test 1 only ARL₁ of 31.31. Also, it is now possible to detect a negative one, two or three sigma shift in proportion.

Template Notes:

1. Specify **Exact (Test 1 Only)** or **Monte Carlo** using the drop-down list. Parameters to be specified will be shown in yellow highlight, otherwise they are hidden.
2. Exact uses the binomial cumulative distribution function. Monte Carlo simulation uses binomial random data with specified proportion and allows you to assess the ARL performance of all 4 Tests for Special Causes. It also produces the table of Run Length Standard Deviation and Percentiles (scroll right to view).
3. Test 1 - 1 point more than 3 standard deviations from the center line (CL) is always applied.
4. Enter the **Subgroup Size**.
5. Enter the **In-Control Historical Proportion (CL)**. This is a value between 0 and 1.
6. If applicable, enter **Number of Replications**. 1000 (1e3) replications will be fast, approx. 10 seconds, but will have an ARL0 error approx. = +/- 10%; 10,000 (1e4) replications will take about a minute, with an ARL0 error = +/- 3.2%; 100,000 (1e5) replications will take about ten minutes, with an ARL0 error = +/- 1%.
7. If applicable, select values for Tests 2 to 4 using the drop-down list. "N/A" indicates that the test is not applied. Tests 2 and 3 provide options that match those provided in SigmaXL's '**Tests for Special Causes' Defaults** dialog.
8. Click the **Calculate Attribute P ARL** button to produce the ARL table and chart. If Monte Carlo was selected, the table of Run Length Standard Deviation and Percentiles will also be produced.
9. The Attribute P ARL is for a two-sided chart with zero-state, i.e., the shift is assumed to occur at the start. The proportion is also assumed to be known. This will not likely be the case in use, but is still useful for determining parameter settings and comparison of ARL across chart types.
10. The ARL Chart is similar to an Operating Characteristic (OC) Curve, except that the Y axis is ARL rather than Beta probability. Note that the P chart may be ARL biased, with maximum ARL occurring above or below the CL.
11. Due to the complexity of calculations, SigmaXL must be loaded and appear on the menu in order for this template to function. Do not add or delete rows or columns in this template.

REFERENCE:

- [1] Montgomery, D.C. (2013), *Introduction to Statistical Quality Control*, Seventh Ed., Wiley.

Attribute C ARL

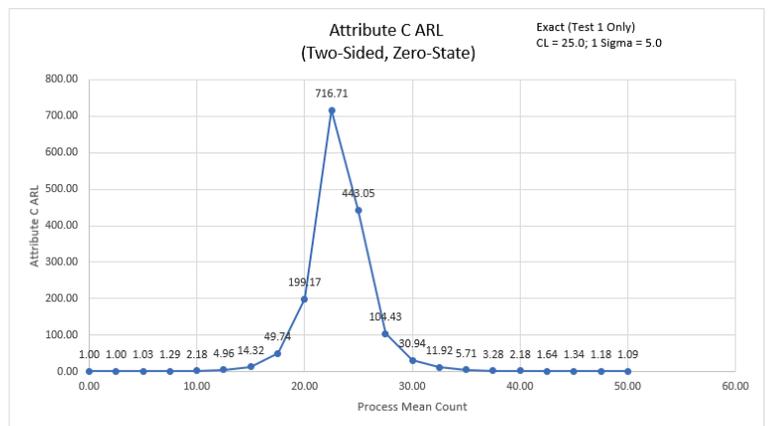
1. Click **SigmaXL > Templates & Calculators > Control Chart Templates > Average Run Length (ARL) Calculators > Attribute C ARL**. This template is also located at **SigmaXL > Control Charts > Control Chart Templates > Average Run Length (ARL) Calculators > Attribute C ARL**.
2. The default template settings are **Specify = Exact (Test 1 Only)**, **In-Control Historical Mean Count (CL) = 25**.

Attribute C Average Run Length (ARL) Calculator	
User Inputs:	
Specify:	Exact (Test 1 Only)
In-Control Historical Mean Count (CL):	25
Calculate Attribute C ARL	

Note: Specify Exact (Test 1 Only) or Monte Carlo using the drop-down list. Parameters to be specified will be shown in yellow highlight, otherwise they are hidden. Exact uses the Poisson cumulative distribution function. Monte Carlo simulation uses Poisson random data with specified proportion and allows you to assess the ARL performance of all 4 Tests for Special Causes. Test 1 - 1 point more than 3 standard deviations from the center line (CL) is always applied. Monte Carlo simulation also produces the table of Run Length Standard Deviation and Percentiles (scroll right to view).

3. Click the **Calculate Attribute C ARL** button to reproduce the ARL table and chart.

Shift in Mean (Multiple of Sigma) 1 Sigma = 5.0	Process Mean Count	Attribute C ARL (Two-Sided, Zero-State)
-5.0	0.00	1.00
-4.5	2.50	1.00
-4.0	5.00	1.03
-3.5	7.50	1.29
-3.0	10.00	2.18
-2.5	12.50	4.96
-2.0	15.00	14.32
-1.5	17.50	49.74
-1.0	20.00	199.17
-0.5	22.50	716.71
0.0	25.00	443.05
0.5	27.50	104.43
1.0	30.00	30.94
1.5	32.50	11.92
2.0	35.00	5.71
2.5	37.50	3.28
3.0	40.00	2.18
3.5	42.50	1.64
4.0	45.00	1.34
4.5	47.50	1.18
5.0	50.00	1.09



The ARL_0 (in-control ARL with 0 shift in mean) for the Attribute C chart is 443.05. The ARL_1 for a small positive 1 sigma shift in mean count is 30.94, so is slow to detect. On the other hand, a large 3 sigma shift in mean count has an $ARL = 2.18$, so is detected rapidly. Small negative shifts of one-half or one sigma cannot be detected.

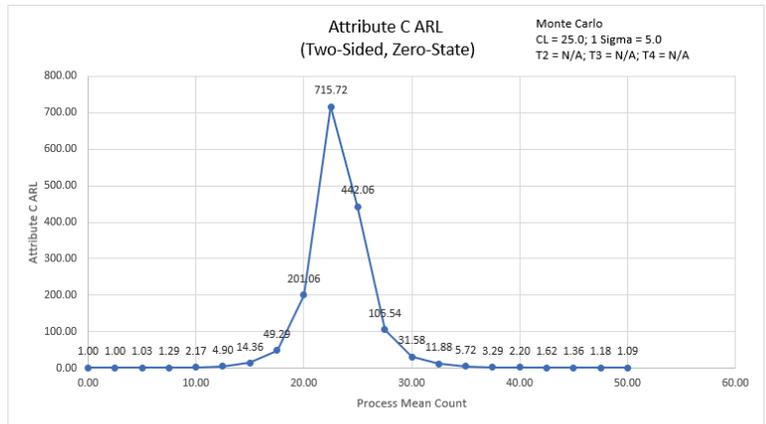
Note: The Process Mean Count values are displayed on the ARL chart, but the shift in mean count as a multiple of sigma is also included in the table.

- Now we will use Monte Carlo simulation to obtain approximate Run Length standard deviation and percentiles for the Attribute C chart. Select **Specify = Monte Carlo**. Enter **In-Control Historical Mean Count (CL) = 25**, **Number of Replications = 1e4**, **Test 2 = N/A**, **Test 3 = N/A**, **Test 4 = N/A**.

Attribute C Average Run Length (ARL) Calculator	
User Inputs:	
Specify:	Monte Carlo
In-Control Historical Mean Count (CL):	25
Number of Replications:	1.00E+04
Test 2 - points in a row on same side of CL:	N/A
Test 3 - points in a row all increasing or decreasing:	N/A
Test 4 - points in a row alternating up and down:	N/A

- Click the **Calculate Attribute C ARL** button to produce the Monte Carlo approximate ARL table, ARL chart and Run Length Standard Deviation and Percentiles table (scroll right to view). Monte Carlo simulation with 10,000 (1e4) replications will take about a minute to run.

Shift in Mean (Multiple of Sigma) 1 Sigma = 5.0	Process Mean Count	Attribute C ARL (Two-Sided, Zero-State)
-5.0	0.00	1.00
-4.5	2.50	1.00
-4.0	5.00	1.03
-3.5	7.50	1.29
-3.0	10.00	2.17
-2.5	12.50	4.90
-2.0	15.00	14.36
-1.5	17.50	49.29
-1.0	20.00	201.06
-0.5	22.50	715.72
0.0	25.00	442.06
0.5	27.50	105.54
1.0	30.00	31.58
1.5	32.50	11.88
2.0	35.00	5.72
2.5	37.50	3.29
3.0	40.00	2.20
3.5	42.50	1.62
4.0	45.00	1.36
4.5	47.50	1.18
5.0	50.00	1.09



Shift in Mean (Multiple of Sigma) 1 Sigma = 5.0	Process Mean Count	Standard Deviation	Monte Carlo Simulation Run Length Standard Deviation and Percentiles (Two-Sided, Zero-State)									
			1	5th	10th	25th	50th (Median)	75th	90th	95th	99th	
-5.0	0.00	0.00	1	1	1	1	1	1	1	1	1	1
-4.5	2.50	0.01	1	1	1	1	1	1	1	1	1	1
-4.0	5.00	0.19	1	1	1	1	1	1	1	1	1	2
-3.5	7.50	0.61	1	1	1	1	1	1	1	2	2	4
-3.0	10.00	1.63	1	1	1	1	2	3	4	5	8	8
-2.5	12.50	4.32	1	1	1	2	4	7	11	14	20	20
-2.0	15.00	13.98	1	1	2	4	10	20	32	43	66	66
-1.5	17.50	49.70	1	3	6	14	34	68	113	147	228	228
-1.0	20.00	200.60	3	12	23	59	140	279	460	588	942	942
-0.5	22.50	720.75	8	37	75	201	491	989	1659	2149	3415	3415
0.0	25.00	438.88	5	24	47	128	308	609	1013	1320	2053	2053
0.5	27.50	105.07	1	6	11	31	73	147	241	312	488	488
1.0	30.00	30.83	1	2	4	10	22	44	71	92	147	147
1.5	32.50	11.30	1	1	2	4	8	16	27	35	52	52
2.0	35.00	5.24	1	1	1	2	4	8	13	16	24	24
2.5	37.50	2.73	1	1	1	1	2	4	7	9	13	13
3.0	40.00	1.60	1	1	1	1	2	3	4	5	8	8
3.5	42.50	1.01	1	1	1	1	1	2	3	4	5	5
4.0	45.00	0.71	1	1	1	1	1	2	2	3	4	4
4.5	47.50	0.46	1	1	1	1	1	1	2	2	3	3
5.0	50.00	0.31	1	1	1	1	1	1	1	2	2	2

The additional run length statistics show the large variation of run length values. The $MRL_0 = 308$ (in-control median run length with 0 sigma shift in process mean count).

Note: The results will vary slightly since this is Monte Carlo simulation

- Now we will use Monte Carlo simulation to assess the Attribute C chart with all 4 Tests for Special Causes applied. Enter **Specify = Monte Carlo, In-Control Historical Mean Count (CL) = 25. Number of Replications = 1e4, Test 2 = 9, Test 3 = 6, and Test 4 = 14.**

Attribute C Average Run Length (ARL) Calculator	
User Inputs:	
Specify:	Monte Carlo
In-Control Historical Mean Count (CL):	25
Number of Replications:	1.00E+04
Test 2 - points in a row on same side of CL:	9
Test 3 - points in a row all increasing or decreasing:	6
Test 4 - points in a row alternating up and down:	14

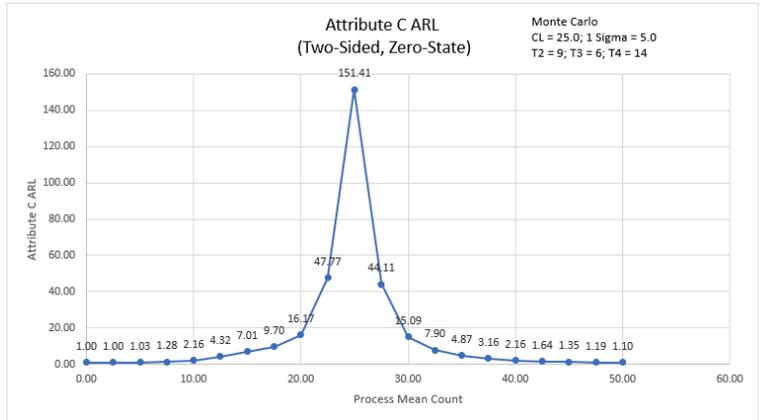
Calculate Attribute C ARL

Note: These are the test settings used as defaults in SigmaXL > Control Charts > 'Tests for Special Causes' Defaults. Test 1 is always applied.

- Click the **Calculate Attribute C ARL** button to produce the Monte Carlo approximate ARL table, ARL chart and Run Length Standard Deviation and Percentiles table.

SigmaXL: Control Phase Tools: SPC

Shift in Mean (Multiple of Sigma) 1 Sigma = 5.0	Process Mean Count	Attribute C ARL (Two-Sided, Zero-State)
-5.0	0.00	1.00
-4.5	2.50	1.00
-4.0	5.00	1.03
-3.5	7.50	1.28
-3.0	10.00	2.16
-2.5	12.50	4.32
-2.0	15.00	7.01
-1.5	17.50	9.70
-1.0	20.00	16.17
-0.5	22.50	47.77
0.0	25.00	151.41
0.5	27.50	44.11
1.0	30.00	15.09
1.5	32.50	7.90
2.0	35.00	4.87
2.5	37.50	3.16
3.0	40.00	2.16
3.5	42.50	1.64
4.0	45.00	1.35
4.5	47.50	1.19
5.0	50.00	1.10



Shift in Mean (Multiple of Sigma) 1 Sigma = 5.0	Process Mean Count	Standard Deviation	Monte Carlo Simulation Run Length Standard Deviation and Percentiles (Two-Sided, Zero-State)								
			1	5th	10th	25th	50th (Median)	75th	90th	95th	99th
-5.0	0.00	0.00	1	1	1	1	1	1	1	1	1
-4.5	2.50	0.02	1	1	1	1	1	1	1	1	1
-4.0	5.00	0.17	1	1	1	1	1	1	1	1	2
-3.5	7.50	0.61	1	1	1	1	1	1	2	2	4
-3.0	10.00	1.53	1	1	1	1	2	3	4	5	8
-2.5	12.50	2.85	1	1	1	2	4	7	9	9	9
-2.0	15.00	3.07	1	1	2	5	9	9	9	9	15
-1.5	17.50	3.91	1	3	6	9	9	9	15	17	24
-1.0	20.00	10.36	2	9	9	9	13	20	29	37	54
-0.5	22.50	41.12	7	9	10	18	35	64	101	131	194
0.0	25.00	146.52	5	13	21	48	106	206	344	448	674
0.5	27.50	39.66	1	6	9	16	32	59	95	122	186
1.0	30.00	11.18	1	2	4	9	12	20	30	37	54
1.5	32.50	5.22	1	1	2	4	8	9	15	18	25
2.0	35.00	3.42	1	1	1	2	4	8	9	10	16
2.5	37.50	2.33	1	1	1	1	2	4	7	9	9
3.0	40.00	1.56	1	1	1	1	2	3	4	5	8
3.5	42.50	1.02	1	1	1	1	1	2	3	4	5
4.0	45.00	0.68	1	1	1	1	1	2	2	3	4
4.5	47.50	0.48	1	1	1	1	1	1	2	2	3
5.0	50.00	0.32	1	1	1	1	1	1	1	2	2

ARL₀ with all 4 tests for special causes is approx. 151.4. This is a poor performance with a 2.9 x increase (443.1/151.4) in false alarms compared to Exact Test 1 only. MRL₀ is approx. 106. On the other hand, ARL₁ for a small 1 sigma shift in proportion is approx. 15.1, so is much faster to detect than the Exact Test 1 only ARL₁ of 30.94. It is now possible to detect a small negative one sigma shift in mean count.

Template Notes:

1. Specify **Exact (Test 1 Only)** or **Monte Carlo** using the drop-down list. Parameters to be specified will be shown in yellow highlight, otherwise they are hidden.
2. Exact uses the Poisson cumulative distribution function. Monte Carlo simulation uses Poisson random data with specified proportion and allows you to assess the ARL performance of all 4 Tests for Special Causes. It also produces the table of Run Length Standard Deviation and Percentiles (scroll right to view).
3. Test 1 - 1 point more than 3 standard deviations from the center line (CL) is always applied.
4. Enter the **Subgroup Size**.
5. Enter the **In-Control Historical Mean Count (CL)**.
6. If applicable, enter **Number of Replications**. 1000 (1e3) replications will be fast, approx. 10 seconds, but will have an ARL0 error approx. = +/- 10%; 10,000 (1e4) replications will take about a minute, with an ARL0 error = +/- 3.2%; 100,000 (1e5) replications will take about ten minutes, with an ARL0 error = +/- 1%.
7. If applicable, select values for Tests 2 to 4 using the drop-down list. "N/A" indicates that the test is not applied. Tests 2 and 3 provide options that match those provided in SigmaXL's '**Tests for Special Causes' Defaults** dialog.
8. Click the **Calculate Attribute C ARL** button to produce the ARL table and chart. If Monte Carlo was selected, the table of Run Length Standard Deviation and Percentiles will also be produced.
9. The Attribute C ARL is for a two-sided chart with zero-state, i.e., the shift is assumed to occur at the start. The proportion is also assumed to be known. This will not likely be the case in use, but is still useful for determining parameter settings and comparison of ARL across chart types.
10. The ARL Chart is similar to an Operating Characteristic (OC) Curve, except that the Y axis is ARL rather than Beta probability. Note that the C chart is ARL biased for small mean counts (<40), with maximum ARL occurring below the CL.
11. Due to the complexity of calculations, SigmaXL must be loaded and appear on the menu in order for this template to function. Do not add or delete rows or columns in this template.

REFERENCE:

- [1] Montgomery, D.C. (2013), *Introduction to Statistical Quality Control*, Seventh Ed., Wiley.

EWMA ARL

1. Click **SigmaXL > Templates & Calculators > Control Chart Templates > Average Run Length (ARL) Calculators > EWMA ARL**. This template is also located at **SigmaXL > Control Charts > Control Chart Templates > Average Run Length (ARL) Calculators > EWMA ARL**.
2. The default template settings are **Specify = Weight (Lambda) & K**, **Weight (Lambda) = 0.1**, **K (Sigma Multiplier) = 2.7**, **Type of Limits = Time-Varying**, **Method = Markov Chain**.

EWMA Average Run Length (ARL) Calculator	
User Inputs:	
Specify:	Weight (Lambda) & K
Weight (Lambda):	0.1
K (Sigma Multiplier):	2.7
Type of Limits:	Time-Varying
Method:	Markov Chain

Notes: Parameters to be specified will be shown in yellow highlight, otherwise they are hidden. The EWMA parameter Weight (Lambda) is a value between 0 and 1 and controls the amount of influence that previous observations have on the current EWMA statistic. A value near 1 puts almost all weight on the current observation, making it resemble a Shewhart chart. For values near 0, a small weight is applied to almost all of the past observations, so the EWMA chart performance is similar to that of a CUSUM chart. The EWMA parameter K (Sigma Multiplier) is a value typically between 2 and 4. It is also referred to as L, but SigmaXL uses K to avoid confusion with Lambda.

The EWMA control chart template has hard coded the Type of Limits as *Time-Varying*, since they improve the sensitivity of the EWMA to detect early changes in the process mean. *Fixed* is included as an option here for comparison of ARL results to published papers. Also, if the process is in control when the EWMA is started but shifts out of control after the control limits have stabilized, the more appropriate ARL for such a case would be Type of Limits = *Fixed*.

The Markov Chain approximation is fast and accurate to compute ARLs. Monte Carlo simulation allows you to assess robustness to nonnormality and also produces the table of Run Length Standard Deviation and Percentiles (scroll right to view).

For further details on the Markov Chain approximation see Lucas [1] for fixed and Steiner [3] for time-varying. Monte Carlo simulation uses the Pearson Family of distributions to match the specified skewness and kurtosis.

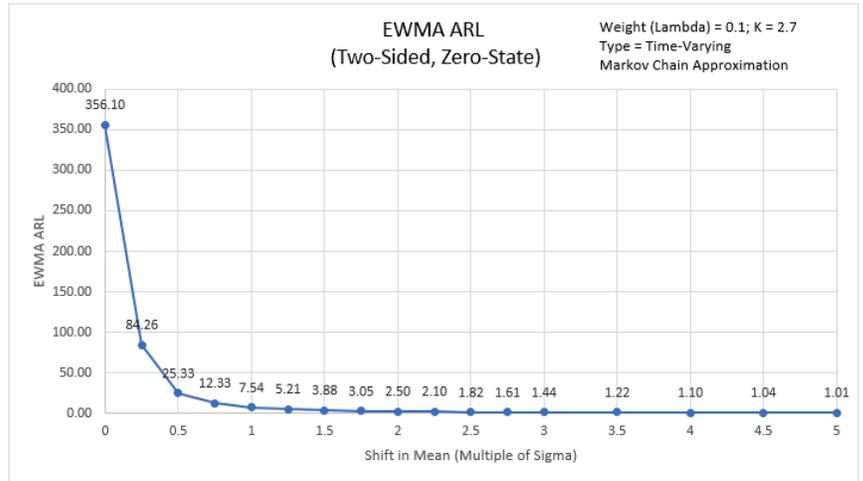
The EWMA ARL is for a two-sided chart with zero-state, i.e., the shift is assumed to occur at the start. The mean and standard deviation are also assumed to be known. This will not likely be

the case in use, but is still useful for determining parameter settings and comparison of ARL across chart types.

All ARL calculations for EWMA use a standardized in-control mean=0 and sigma=1.

- Click the **Calculate EWMA ARL** button to reproduce the ARL table and chart.

Shift in Mean (Multiple of Sigma)	EWMA ARL (Two-Sided, Zero-State)
0	356.10
0.25	84.26
0.5	25.33
0.75	12.33
1	7.54
1.25	5.21
1.5	3.88
1.75	3.05
2	2.50
2.25	2.10
2.5	1.82
2.75	1.61
3	1.44
3.5	1.22
4	1.10
4.5	1.04
5	1.01



The ARL_0 (in-control ARL with 0 shift in mean) for the EWMA chart with these settings is 356.1, which is close to the Shewhart ARL_0 of 370.4. The ARL_1 for a small 1 sigma shift in mean is 7.54, so is much faster to detect than the ARL_1 of 43.89 for [Shewhart Individuals](#) and faster to detect than the Monte Carlo ARL_1 of 9.7 for the [Shewhart Individuals with 8 tests for special causes](#).

- Now we will compare time-varying to fixed limits. Select **Specify** = *Weight (Lambda) & K*. Enter **Weight (Lambda)** = 0.1, **K (Sigma Multiplier)** = 2.7, **Type of Limits** = *Fixed*, **Method** = *Markov Chain*.

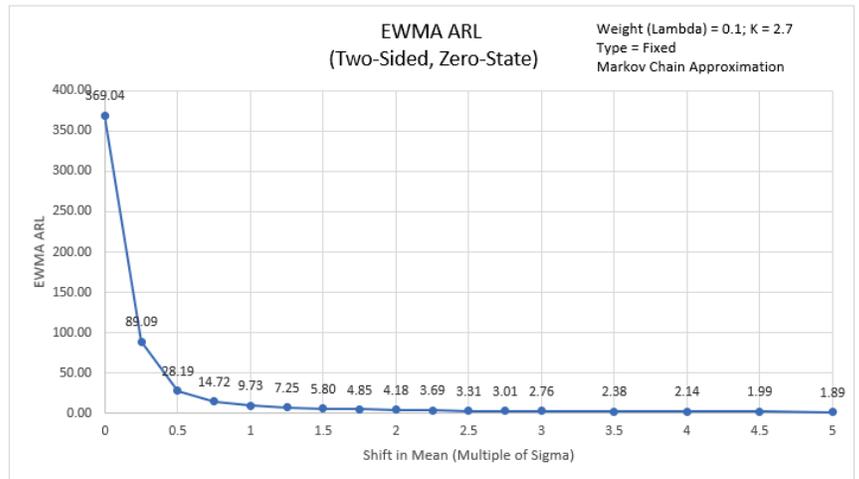
EWMA Average Run Length (ARL) Calculator

User Inputs:

Specify:	Weight (Lambda) & K	Calculate EWMA ARL
Weight (Lambda):	0.1	
K (Sigma Multiplier):	2.7	
Type of Limits:	Fixed	
Method:	Markov Chain	

- Click the **Calculate EWMA ARL** button to produce the ARL table and chart for these settings:

Shift in Mean (Multiple of Sigma)	EWMA ARL (Two-Sided, Zero-State)
0	369.04
0.25	89.09
0.5	28.19
0.75	14.72
1	9.73
1.25	7.25
1.5	5.80
1.75	4.85
2	4.18
2.25	3.69
2.5	3.31
2.75	3.01
3	2.76
3.5	2.38
4	2.14
4.5	1.99
5	1.89



The ARL_0 for the EWMA chart with these settings is 369.04, which is close to the Shewhart ARL_0 of 370.4. The ARL_1 for a small 1 sigma shift in mean with fixed limits is 9.73, which is slower than the ARL_1 with time-varying limits of 7.54.

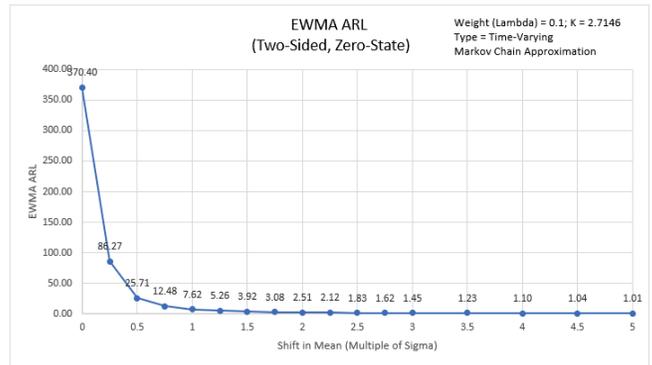
- Next, we will specify Weight (Lambda) = 0.1, the Shewhart ARL_0 value of 370.4 and solve for the K (Sigma Multiplier). Enter **Specify = Weight (Lambda) & ARL0**, **Weight (Lambda) = 0.1**, **In-Control Average Run Length (ARL0) = 370.4**, **Type of Limits = Time-Varying**, **Method = Markov Chain**.

EWMA Average Run Length (ARL) Calculator	
User Inputs:	
Specify:	Weight (Lambda) & ARL0
Weight (Lambda):	0.1
In-Control Average Run Length (ARL0):	370.4
Type of Limits:	Time-Varying
Method:	Markov Chain
Calculate EWMA ARL	

- Click the **Calculate EWMA ARL** button to produce the updated EWMA Parameters, ARL table and chart for these settings:

EWMA Parameters:	
Weight (Lambda)	0.1
K (Sigma Multiplier)	2.7146
Type of Limits	Time-Varying
Method	Markov Chain Approximation

Shift in Mean (Multiple of Sigma)	EWMA ARL (Two-Sided, Zero-State)
0	370.40
0.25	86.27
0.5	25.71
0.75	12.48
1	7.62
1.25	5.26
1.5	3.92
1.75	3.08
2	2.51
2.25	2.12
2.5	1.83
2.75	1.62
3	1.45
3.5	1.23
4	1.10
4.5	1.04
5	1.01



The ARL_0 for the EWMA chart with these settings is 370.4 as specified. The K (Sigma Multiplier) solved to obtain this ARL_0 value is 2.7146. The ARL_1 for a small 1 sigma shift in mean is 7.62.

- Now we will specify Weight (Lambda) = 0.05, the Shewhart ARL_0 value of 370.4 and solve for the K (Sigma Multiplier). Enter **Specify = Weight (Lambda) & ARL₀**, **Weight (Lambda) = 0.05**, **In-Control Average Run Length (ARL₀) = 370.4**, **Type of Limits = Time-Varying**, **Method = Markov Chain**.

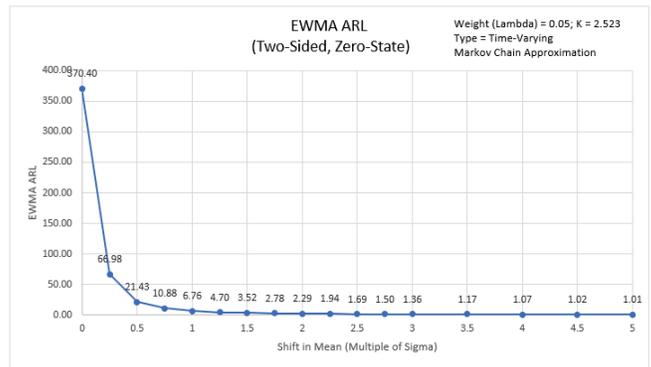
EWMA Average Run Length (ARL) Calculator	
User Inputs:	
Specify:	Weight (Lambda) & ARL ₀
Weight (Lambda):	0.05
In-Control Average Run Length (ARL ₀):	370.4
Type of Limits:	Time-Varying
Method:	Markov Chain

Calculate EWMA ARL

9. Click the **Calculate EWMA ARL** button to produce the updated EWMA Parameters, ARL table and chart for these settings:

EWMA Parameters:	
Weight (Lambda)	0.05
K (Sigma Multiplier)	2.523
Type of Limits	Time-Varying
Method	Markov Chain Approximation

Shift in Mean (Multiple of Sigma)	EWMA ARL (Two-Sided, Zero-State)
0	370.40
0.25	66.98
0.5	21.43
0.75	10.88
1	6.76
1.25	4.70
1.5	3.52
1.75	2.78
2	2.29
2.25	1.94
2.5	1.69
2.75	1.50
3	1.36
3.5	1.17
4	1.07
4.5	1.02
5	1.01



The ARL_0 for the EWMA chart with these settings is 370.4 as specified. The K (Sigma Multiplier) solved to obtain this ARL_0 value is 2.523. The ARL_1 for a small 1 sigma shift in mean is 6.76.

10. Next, we will specify an ARL_0 value of 500 with a desired optimization to detect a 1 sigma shift in mean. The calculator will solve for the optimal Weight (Lambda) and K (Sigma Multiplier). Enter **Specify = ARLO & Shift, In-Control Average Run Length (ARL0) = 500, Shift in Mean to Detect (Multiple of Sigma) = 1, Type of Limits = Time-Varying, Method = Markov Chain.**

EWMA Average Run Length (ARL) Calculator

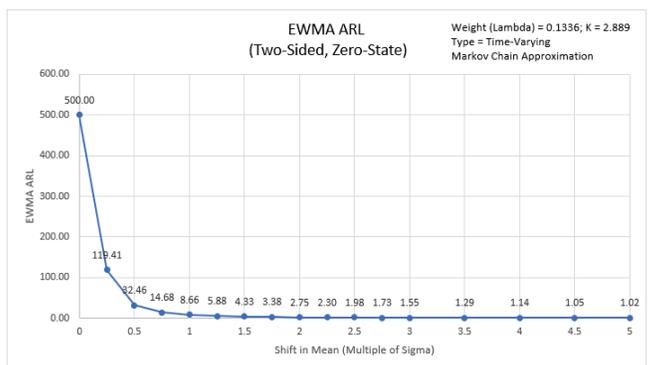
User Inputs:

Specify:	ARLO & Shift	Calculate EWMA ARL
In-Control Average Run Length (ARL0):	500	
Shift in Mean to Detect (Multiple of Sigma):	1	
Type of Limits:	Time-Varying	
Method:	Markov Chain	

11. Click the **Calculate EWMA ARL** button to produce the EWMA Parameters, ARL table and chart for these settings:

EWMA Parameters:	
Weight (Lambda)	0.1336
K (Sigma Multiplier)	2.889
Type of Limits	Time-Varying
Method	Markov Chain Approximation

Shift in Mean (Multiple of Sigma)	EWMA ARL (Two-Sided, Zero-State)
0	500.00
0.25	119.41
0.5	32.46
0.75	14.68
1	8.66
1.25	5.88
1.5	4.33
1.75	3.38
2	2.75
2.25	2.30
2.5	1.98
2.75	1.73
3	1.55
3.5	1.29
4	1.14
4.5	1.05
5	1.02



The ARL_0 for the EWMA chart with these settings is 500.0 as specified. The ARL_1 for a small 1 sigma shift in mean is 8.66. The solved parameters are Weight (Lambda) = 0.1336 and K (Sigma Multiplier) = 2.889.

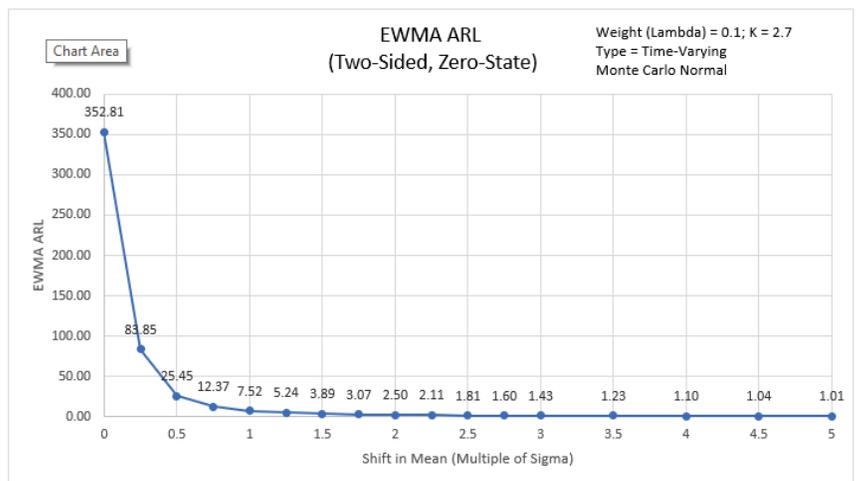
Note: Weight (Lambda) is first solved using fixed limits. This value is then used to solve for K using time-varying limits.

- Now we will use Monte Carlo simulation to obtain approximate Run Length standard deviation and percentiles. Enter **Specify = Weight (Lambda) & K**, **Weight (Lambda) = 0.1**, **K (Sigma Multiplier) = 2.7**, **Type of Limits = Time-Varying**, **Method = Monte Carlo**, **Number of Replications = 1e4**, **Skewness = 0**, **Kurtosis (Normal is 0) = 0**.

EWMA Average Run Length (ARL) Calculator	
User Inputs:	
Specify:	Weight (Lambda) & K
Weight (Lambda):	0.1
K (Sigma Multiplier):	2.7
Type of Limits:	Time-Varying
Method:	Monte Carlo
Number of Replications:	1.00E+04
Skewness:	0
Kurtosis (Normal is 0):	0

- Click the **Calculate EWMA ARL** button to produce the Monte Carlo approximate ARL table, ARL chart and Run Length Standard Deviation and Percentiles table. Monte Carlo simulation with 10,000 (1e4) replications will take about a minute to run.

Shift in Mean (Multiple of Sigma)	EWMA ARL (Two-Sided, Zero-State)
0	352.81
0.25	83.85
0.5	25.45
0.75	12.37
1	7.52
1.25	5.24
1.5	3.89
1.75	3.07
2	2.50
2.25	2.11
2.5	1.81
2.75	1.60
3	1.43
3.5	1.23
4	1.10
4.5	1.04
5	1.01



Monte Carlo Simulation Run Length Standard Deviation and Percentiles (Two-Sided, Zero-State)										
Shift in Mean (Multiple of Sigma)	Standard Deviation	1	5th	10th	25th	50th (Median)	75th	90th	95th	99th
0	361.02	2	14	33	100	238	485	830	1083	1651
0.25	79.98	2	7	12	28	60	114	185	243	388
0.5	20.40	1	3	5	11	20	34	52	65	96
0.75	8.65	1	2	3	6	10	16	24	29	40
1	4.88	1	2	2	4	7	10	14	17	24
1.25	3.19	1	1	2	3	5	7	10	11	15
1.5	2.25	1	1	1	2	3	5	7	8	11
1.75	1.69	1	1	1	2	3	4	5	6	8
2	1.32	1	1	1	2	2	3	4	5	7
2.25	1.05	1	1	1	1	2	3	3	4	5
2.5	0.85	1	1	1	1	2	2	3	3	4
2.75	0.72	1	1	1	1	1	2	3	3	4
3	0.61	1	1	1	1	1	2	2	3	3
3.5	0.45	1	1	1	1	1	1	2	2	3
4	0.31	1	1	1	1	1	1	1	2	2
4.5	0.19	1	1	1	1	1	1	1	1	2
5	0.11	1	1	1	1	1	1	1	1	2

The additional run length statistics show the large variation of run length values. The $MRL_0 = 238$ (in-control median run length with 0 shift in process mean).

Note: The results will vary slightly since this is Monte Carlo simulation.

14. We will now assess robustness to nonnormality using Monte Carlo simulation with Weight (Lambda) = 0.1 and a specified ARL_0 of 370.4 for comparison to Shewhart. Enter **Specify = Weight (Lambda) & ARL_0** , **Weight (Lambda) = 0.1**, **In-Control Average Run Length (ARL_0) = 370.4**, **Type of Limits = Time-Varying**, **Method = Monte Carlo**, **Number of Replications = 1e4**, **Skewness = 2**, **Kurtosis (Normal is 0) = 6**.

EWMA Average Run Length (ARL) Calculator	
User Inputs:	
Specify:	Weight (Lambda) & ARL_0
Weight (Lambda):	0.1
In-Control Average Run Length (ARL_0):	370.4
Type of Limits:	Time-Varying
Method:	Monte Carlo
Number of Replications:	1.00E+04
Skewness:	2
Kurtosis (Normal is 0):	6

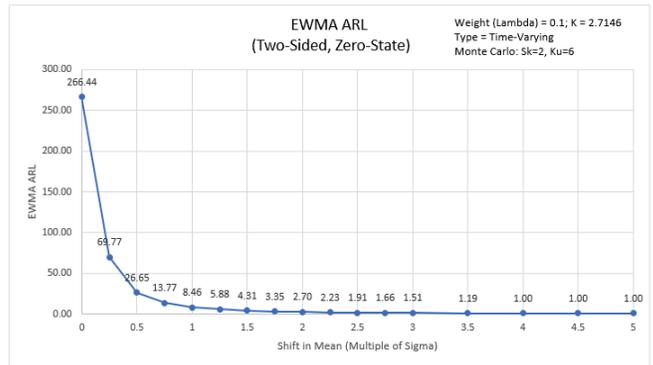
Calculate EWMA ARL

Note: The EWMA parameter K (Sigma Multiplier) will be solved using Markov-Chain approximation and assume a Normal distribution, so will match the value previously calculated above (6).

15. Click the **Calculate EWMA ARL** button to produce the updated EWMA Parameters, Monte Carlo approximate ARL table, ARL chart and Run Length Standard Deviation and Percentiles table:

EWMA Parameters:	
Weight (Lambda)	0.1
K (Sigma Multiplier)	2.7146
Type of Limits	Time-Varying
Method	Monte Carlo: Sk=2, Ku=6

Shift in Mean (Multiple of Sigma)	EWMA ARL (Two-Sided, Zero-State)
0	266.44
0.25	69.77
0.5	26.65
0.75	13.77
1	8.46
1.25	5.88
1.5	4.31
1.75	3.35
2	2.70
2.25	2.23
2.5	1.91
2.75	1.66
3	1.51
3.5	1.19
4	1.00
4.5	1.00
5	1.00



Monte Carlo Simulation Run Length Standard Deviation and Percentiles (Two-Sided, Zero-State)										
Shift in Mean (Multiple of Sigma)	Standard Deviation	1	5th	10th	25th	50th (Median)	75th	90th	95th	99th
0	276.09	1	5	18	68	183	372	618	804	1298
0.25	68.39	1	2	7	21	50	97	157	205	322
0.5	22.63	1	2	4	10	21	36	57	71	103
0.75	10.24	1	1	3	6	12	19	27	33	48
1	5.42	1	1	2	4	8	12	16	19	25
1.25	3.38	1	1	2	3	6	8	10	12	16
1.5	2.24	1	1	1	3	4	6	7	8	10
1.75	1.59	1	1	1	2	3	4	5	6	7
2	1.17	1	1	1	2	3	4	4	5	5
2.25	0.89	1	1	1	2	2	3	3	4	4
2.5	0.71	1	1	1	1	2	2	3	3	3
2.75	0.55	1	1	1	1	2	2	2	2	3
3	0.50	1	1	1	1	2	2	2	2	2
3.5	0.39	1	1	1	1	1	1	2	2	2
4	0.00	1	1	1	1	1	1	1	1	1
4.5	0.00	1	1	1	1	1	1	1	1	1
5	0.00	1	1	1	1	1	1	1	1	1

ARL₀ is approximately 266.4 which is a 1.4 x increase (370.4/266.4) in false alarms compared to Normal but is a much better performance than the ARL₀ = 55 result for [Shewhart Individuals](#). MRL₀ is approx. 183.

16. Next, we will assess robustness to nonnormality using Monte Carlo simulation with a lower Weight (Lambda) = 0.05 and a specified ARL₀ of 370.4. Enter **Specify = Weight (Lambda) & ARL₀**, **Weight (Lambda) = 0.05**, **In-Control Average Run Length (ARL₀) = 370.4**, **Type of Limits = Time-Varying**, **Method = Monte Carlo**, **Number of Replications = 1e4**, **Skewness = 2**, **Kurtosis (Normal is 0) = 6**.

EWMA Average Run Length (ARL) Calculator	
User Inputs:	
Specify:	Weight (Lambda) & ARL ₀
Weight (Lambda):	0.05
In-Control Average Run Length (ARL ₀):	370.4
Type of Limits:	Time-Varying
Method:	Monte Carlo
Number of Replications:	1.00E+04
Skewness:	2
Kurtosis (Normal is 0):	6

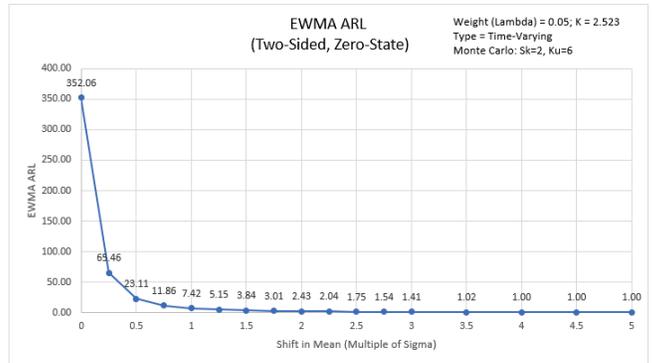
Calculate EWMA ARL

Note: Montgomery [2] (Table 9.12) and Borror, Montgomery & Runger [4] point out that the EWMA becomes more robust with lower values of Lambda.

17. Click the **Calculate EWMA ARL** button to produce the updated EWMA Parameters, Monte Carlo approximate ARL table, ARL chart and Run Length Standard Deviation and Percentiles table:

EWMA Parameters:	
Weight (Lambda)	0.05
K (Sigma Multiplier)	2.523
Type of Limits	Time-Varying
Method	Monte Carlo: Sk=2, Ku=6

Shift in Mean (Multiple of Sigma)	EWMA ARL (Two-Sided, Zero-State)
0	352.06
0.25	65.46
0.5	23.11
0.75	11.86
1	7.42
1.25	5.15
1.5	3.84
1.75	3.01
2	2.43
2.25	2.04
2.5	1.75
2.75	1.54
3	1.41
3.5	1.02
4	1.00
4.5	1.00
5	1.00



Monte Carlo Simulation Run Length Standard Deviation and Percentiles (Two-Sided, Zero-State)										
Shift in Mean (Multiple of Sigma)	Standard Deviation	1	5th	10th	25th	50th (Median)	75th	90th	95th	99th
0	370.14	1	3	15	83	239	499	842	1101	1625
0.25	62.29	1	2	5	20	49	91	146	193	286
0.5	18.80	1	1	3	9	19	33	48	58	86
0.75	8.38	1	1	2	5	10	16	23	28	37
1	4.76	1	1	2	4	7	10	14	16	21
1.25	2.96	1	1	1	3	5	7	9	11	13
1.5	2.01	1	1	1	2	4	5	7	7	9
1.75	1.44	1	1	1	2	3	4	5	5	6
2	1.05	1	1	1	2	2	3	4	4	5
2.25	0.81	1	1	1	1	2	3	3	3	4
2.5	0.63	1	1	1	1	2	2	3	3	3
2.75	0.50	1	1	1	1	2	2	2	2	2
3	0.49	1	1	1	1	1	2	2	2	2
3.5	0.15	1	1	1	1	1	1	1	1	2
4	0.00	1	1	1	1	1	1	1	1	1
4.5	0.00	1	1	1	1	1	1	1	1	1
5	0.00	1	1	1	1	1	1	1	1	1

ARL₀ is approximately 352 which is close to the original specified 370.4. The MRL₀ is approx. 239. If robustness to non-normality is a concern then a Weight (Lambda) = 0.05 is recommended.

Template Notes:

1. Specify EWMA parameters: **Weight (Lambda) & K, Weight (Lambda) & ARL0 or ARL0 & Shift** using the drop-down list. Parameters to be specified will be shown in yellow highlight, otherwise they are hidden.
2. If applicable, enter the EWMA parameter **Weight (Lambda)**. This is a value between 0 and 1 and controls the amount of influence that previous observations have on the current EWMA statistic. A value near 1 puts almost all weight on the current observation, making it resemble a Shewhart chart. For values near 0, a small weight is applied to almost all of the past observations, so the EWMA chart performance is similar to that of a CUSUM chart.
3. If applicable, enter the EWMA parameter **K (Sigma Multiplier)**. This is a value typically between 2 and 4. It is also referred to as L, but SigmaXL uses K to avoid confusion with Lambda.
4. If applicable, enter the desired **In-Control Average Run Length (ARL0)**. This will be the target ARL for mean shift = 0 and should be a large value to minimize false alarms, typically 370 to 500. The K (Sigma Multiplier) will be solved to achieve this ARL0, given a specified Weight (Lambda) value.
5. If applicable, enter the desired **Shift in Mean to Detect (Multiple of Sigma)**. The Weight (Lambda) value that minimizes ARL for the specified shift will be solved.
6. Select Type of Limits: **Time-Varying** or **Fixed** using the drop-down list.
7. The EWMA control chart template uses time-varying control limits since they improve the sensitivity of the EWMA to detect early changes in the process mean. Published ARL tables typically use fixed limits, so providing both allows comparison between the two types.
8. Select Method: **Markov Chain** or **Monte Carlo** using the drop-down list. Markov Chain approximation is fast and accurate to compute ARLs. Monte Carlo simulation allows you to assess robustness to non-normality and also produces the table of Run Length Standard Deviation and Percentiles (scroll right to view).
9. For further details on the Markov Chain approximation see Lucas [1] for fixed and Steiner [3] for time-varying. Monte Carlo simulation uses the Pearson Family of distributions to match the specified skewness and kurtosis.
10. If applicable, enter **Number of Replications**. 1000 (1e3) replications will be fast, approx. 10 seconds, but will have an ARL0 error approx. = +/- 10%; 10,000 (1e4) replications will take about a minute, with an ARL0 error = +/- 3.2%; 100,000 (1e5) replications will take about ten minutes, with an ARL0 error = +/- 1%.
11. If applicable, enter **Skewness**. Skewness must be ≥ 0 . Skewness = 0 is symmetric.
12. If applicable, enter **Kurtosis (Normal is 0)**. Also known as Excess Kurtosis, it must be $\geq \text{Skewness}^2 - 1.48$. This is required to keep the distribution unimodal. If Skewness=0 and Kurtosis = 0, the distribution is normal.
13. Click the **Calculate EWMA ARL** button to produce the ARL table and chart. If Monte Carlo was selected, the table of Run Length Standard Deviation and Percentiles will also be produced.
14. The EWMA ARL is for a two-sided chart with zero-state, i.e., the shift is assumed to occur at the start. The mean and standard deviation are also assumed to be known.

This will not likely be the case in use, but is still useful for determining parameter settings and comparison of ARL across chart types.

15. Due to the complexity of calculations, SigmaXL must be loaded and appear on the menu in order for this template to function. Do not add or delete rows or columns in this template.

REFERENCES:

- [1] Lucas J.M. and Saccucci M.S. (1990), "Exponentially weighted moving average control schemes: Properties and enhancements", *Technometrics* 32, 1-12.
- [2] Montgomery, D.C. (2013), *Introduction to Statistical Quality Control*, Seventh Ed., Wiley.
- [3] Steiner, S. H. (1999), "EWMA control charts with time-varying control limits and fast initial response", *Journal of Quality Technology* 31(1), 75-86.
- [4] Borrer, C. M., Montgomery, D.C. and Runger G. C. (1999). "Robustness of the EWMA Control Chart to Nonnormality," *Journal of Quality Technology*, 31(3), 309–316.

CUSUM ARL

1. Click **SigmaXL > Templates & Calculators > Control Chart Templates > Average Run Length (ARL) Calculators > CUSUM ARL**. This template is also located at **SigmaXL > Control Charts > Control Chart Templates > Average Run Length (ARL) Calculators > CUSUM ARL**.
2. The default template settings are **Specify = k & h** , **k parameter = 0.5**, **h parameter = 5**, **Fast Initial Response $FIR = 0$** , **Method = Markov Chain**.

User Inputs:	
Specify:	k & h
k parameter:	0.5
h parameter:	5
Fast Initial Response:	FIR = 0
Method:	Markov Chain

Calculate CUSUM ARL

Notes: Parameters to be specified will be shown in yellow highlight, otherwise they are hidden. The CUSUM parameter k is the reference (or slack) value, typically set to 0.5. It sets the size of mean shift ($2k$ sigma) that you would like to detect quickly, so 0.5 denotes rapid detection of a shift in mean = 1 sigma. Alternatively, Woodall & Faltin [4] recommend larger k values (e.g., $k = 0.9$) to avoid false alarms and detect shifts of practical significance.

The CUSUM parameter h is the decision interval, typically set to 4 or 5.

FIR sets the initial CUSUM statistic so that it improves the sensitivity to a mean shift at startup. Note that if the process is in control when the CUSUM is started but shifts out of control later, the more appropriate ARL for such a case would be $FIR=0$. See Montgomery [3], pages 426-427.

Markov Chain approximation is fast and accurate to compute ARLs. Monte Carlo simulation allows you to assess robustness to nonnormality and also produces the table of Run Length Standard Deviation and Percentiles (scroll right to view).

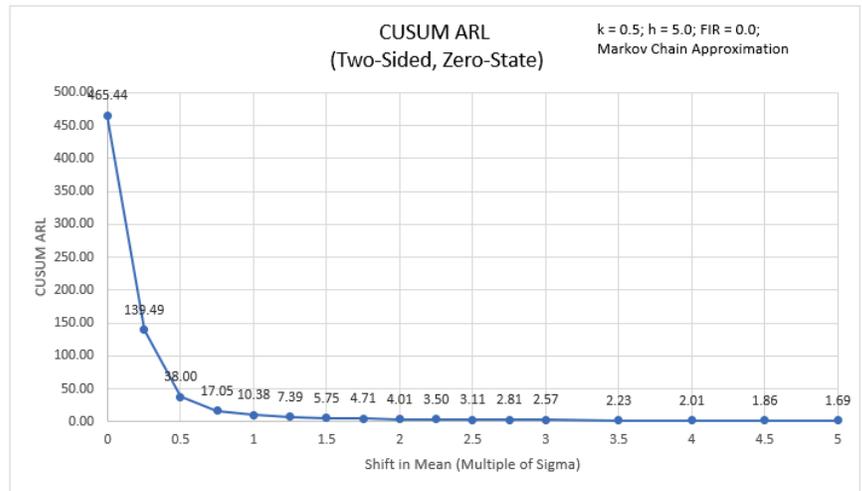
For further details on the Markov Chain approximation see Hawkins [1] and Lucas [2]. Monte Carlo simulation uses the Pearson Family of distributions to match the specified skewness and kurtosis.

The CUSUM ARL is for a two-sided chart with zero-state, i.e., the shift is assumed to occur at the start. The mean and standard deviation are also assumed to be known. This will not likely be the case in use, but is still useful for determining parameter settings and comparison of ARL across chart types.

All ARL calculations for CUSUM use a standardized in-control mean=0 and sigma=1.

3. Click the **Calculate CUSUM ARL** button to reproduce the ARL table and chart.

Shift in Mean (Multiple of Sigma)	CUSUM ARL (Two-Sided, Zero-State)
0	465.44
0.25	139.49
0.5	38.00
0.75	17.05
1	10.38
1.25	7.39
1.5	5.75
1.75	4.71
2	4.01
2.25	3.50
2.5	3.11
2.75	2.81
3	2.57
3.5	2.23
4	2.01
4.5	1.86
5	1.69



The ARL_0 (in-control ARL with 0 shift in mean) for the CUSUM chart with these settings is 465.44. The ARL_1 for a small 1 sigma shift in mean is 10.38.

4. Now we will evaluate CUSUM with the same parameters, but use the Fast Initial Response option. Select **Specify = k & h**. Enter **k parameter = 0.5**, **h parameter = 5**, **Fast Initial Response FIR = h/2**, **Method = Markov Chain**.

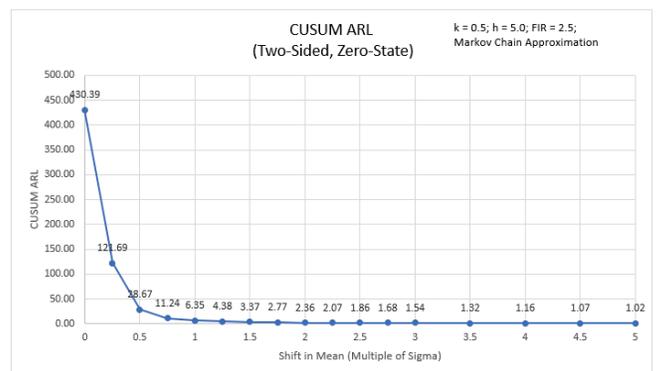
User Inputs:	
Specify:	k & h
k parameter:	0.5
h parameter:	5
Fast Initial Response:	FIR = h/2
Method:	Markov Chain

Calculate CUSUM ARL

5. Click the **Calculate CUSUM ARL** button to produce the ARL table and chart for these settings.

Specified CUSUM Parameters:	
k	0.5
h	5.0
FIR	2.5
Method	Markov Chain Approximation

Shift in Mean (Multiple of Sigma)	CUSUM ARL (Two-Sided, Zero-State)
0	430.39
0.25	121.69
0.5	28.67
0.75	11.24
1	6.35
1.25	4.38
1.5	3.37
1.75	2.77
2	2.36
2.25	2.07
2.5	1.86
2.75	1.68
3	1.54
3.5	1.32
4	1.16
4.5	1.07
5	1.02



The ARL_0 for the CUSUM chart with these settings is 430.39 which is slightly lower than the $FIR = 0$ setting, so slightly higher false alarm rate. The ARL_1 for a small 1 sigma shift in mean is 6.35 which is faster than the 10.38 for $FIR = 0$.

- Now we will specify the CUSUM k parameter = 0.5 with a Shewhart ARL_0 value of 370.4 and solve for the h parameter. Enter **Specify = k & ARL_0 , k parameter = 0.5, In-Control Average Run Length (ARL_0) = 370.4, Fast Initial Response $FIR = 0$, Method = Markov Chain**.

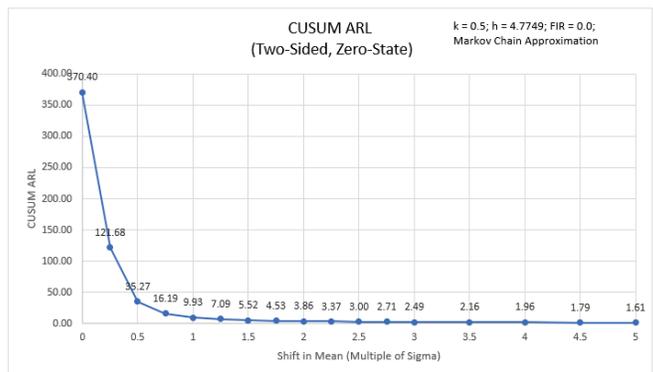
User Inputs:		
Specify:	k & ARL_0	
k parameter:	0.5	
In-Control Average Run Length (ARL_0):	370.4	
Fast Initial Response:	FIR = 0	
Method:	Markov Chain	

Calculate CUSUM ARL

- Click the **Calculate CUSUM ARL** button to produce the CUSUM Parameters, ARL table and chart for these settings.

CUSUM Parameters:	
k	0.5
h	4.7749
FIR	0.0
Method	Markov Chain Approximation

Shift in Mean (Multiple of Sigma)	CUSUM ARL (Two-Sided, Zero-State)
0	370.40
0.25	121.68
0.5	35.27
0.75	16.19
1	9.93
1.25	7.09
1.5	5.52
1.75	4.53
2	3.86
2.25	3.37
2.5	3.00
2.75	2.71
3	2.49
3.5	2.16
4	1.96
4.5	1.79
5	1.61



The ARL_0 for the CUSUM chart with these settings is 370.4 as specified. The h parameter solved to obtain this ARL_0 value is 4.7749. The ARL_1 for a small 1 sigma shift in mean is 9.93 so is much faster to detect than the ARL_1 of 43.89 for [Shewhart Individuals](#) and close to the Monte Carlo ARL_1 of 9.7 for [Shewhart Individuals with 8 tests for special causes](#).

8. Next, we will specify a Shewhart ARL_0 value of 370.4, with a desired optimization to detect a 1 sigma shift in mean and use Fast Initial Response. The calculator will solve for the optimal k and h parameters. Enter **Specify = ARL_0 & Shift**, **In-Control Average Run Length (ARL_0) = 370.4**, **Shift in Mean to Detect (Multiple of Sigma) = 1**, **Fast Initial Response $FIR = h/2$** , **Method = Markov Chain**.

User Inputs:	
Specify:	ARL0 & Shift
In-Control Average Run Length (ARL_0):	370.4
Shift in Mean to Detect (Multiple of Sigma):	1
Fast Initial Response:	FIR = h/2
Method:	Markov Chain

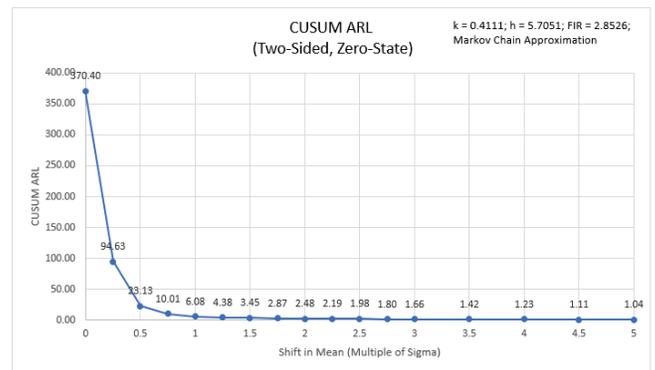
Calculate CUSUM ARL

Note: Since both k and h are solved, this takes about 20-30 seconds to compute.

9. Click the **Calculate CUSUM ARL** button to produce the CUSUM Parameters, ARL table and chart for these settings.

CUSUM Parameters:	
k	0.4111
h	5.7051
FIR	2.8526
Method	Markov Chain Approximation

Shift in Mean (Multiple of Sigma)	CUSUM ARL (Two-Sided, Zero-State)
0	370.40
0.25	94.63
0.5	23.13
0.75	10.01
1	6.08
1.25	4.38
1.5	3.45
1.75	2.87
2	2.48
2.25	2.19
2.5	1.98
2.75	1.80
3	1.66
3.5	1.42
4	1.23
4.5	1.11
5	1.04



The ARL_0 for the CUSUM chart with these settings is 370.4 as specified. ARL_1 for a small 1 sigma shift in mean is 6.08 so is much faster to detect than the ARL_1 of 43.89 for [Shewhart Individuals](#) and faster than the Monte Carlo ARL_1 of 9.7 for [Shewhart Individuals with 8 tests for special causes](#). It is slightly faster to detect than the ARL_1 of 6.76 for [Time-Varying EWMA with Weight\(Lambda\) = .05](#). The solved parameters are k = 0.4111, h = 5.7051 and FIR = 2.8526.

10. As an alternative to using the CUSUM to rapidly detect small shifts in mean, Woodall & Faltin [4] recommend larger k values to avoid false alarms and detect shifts of practical significance. Enter **Specify** = k & h , **k parameter** = 0.9, **h parameter** = 4.65, **Fast Initial Response FIR** = 0, **Method** = *Markov Chain*.

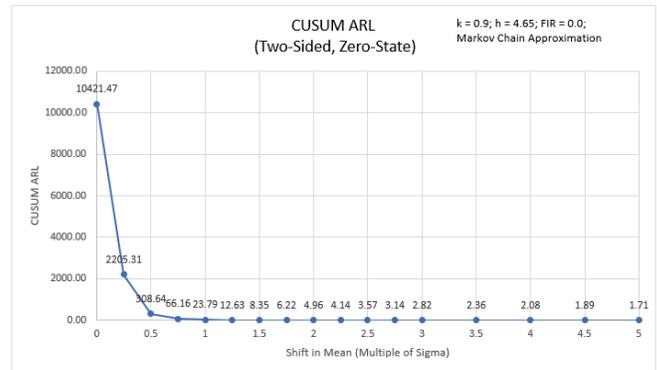
User Inputs:	
Specify:	k & h
k parameter:	0.9
h parameter:	4.65
Fast Initial Response:	FIR = 0
Method:	Markov Chain

Calculate CUSUM ARL

11. Click the **Calculate CUSUM ARL** button to produce the CUSUM Parameters, ARL table and chart for these settings.

Specified CUSUM Parameters:	
k	0.9
h	4.65
FIR	0.0
Method	Markov Chain Approximation

Shift in Mean (Multiple of Sigma)	CUSUM ARL (Two-Sided, Zero-State)
0	10421.47
0.25	2205.31
0.5	308.64
0.75	66.16
1	23.79
1.25	12.63
1.5	8.35
1.75	6.22
2	4.96
2.25	4.14
2.5	3.57
2.75	3.14
3	2.82
3.5	2.36
4	2.08
4.5	1.89
5	1.71



This gives large ARL values for shift in mean ≤ 1 sigma and small ARL values for a shift in mean ≥ 1.5 sigma.

12. Now we will use Monte Carlo simulation to obtain approximate Run Length standard deviation and percentiles using the CUSUM k parameter = 0.5 with a Shewhart ARL_0 value of 370.4. Enter **Specify = k & ARL_0** , **k parameter = 0.5**, **In-Control Average Run Length (ARL_0) = 370.4**, **Fast Initial Response $FIR = 0$** , **Method = Monte Carlo**, **Number of Replications = 1e4**, **Skewness = 0**, **Kurtosis (Normal is 0) = 0**.

User Inputs:	
Specify:	k & ARL_0
k parameter:	0.5
In-Control Average Run Length (ARL_0):	370.4
Fast Initial Response:	FIR = 0
Method:	Monte Carlo
Number of Replications:	1.00E+04
Skewness:	0
Kurtosis (Normal is 0):	0

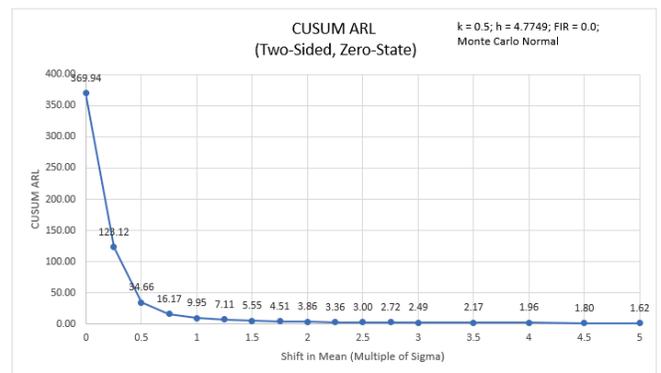
Calculate CUSUM ARL

Note: The CUSUM h parameter will be solved first using the Markov Chain approximation and assumes a Normal distribution, so will match the 4.7749 value previously calculated above (6).

13. Click the **Calculate CUSUM ARL** button to produce the CUSUM Parameters, Monte Carlo approximate ARL table, ARL chart and Run Length Standard Deviation and Percentiles table. Monte Carlo simulation with 10,000 (1e4) replications will take about a minute to run.

CUSUM Parameters:	
k	0.5
h	4.7749
FIR	0.0
Method	Monte Carlo Normal

Shift in Mean (Multiple of Sigma)	CUSUM ARL (Two-Sided, Zero-State)
0	369.94
0.25	123.12
0.5	34.66
0.75	16.17
1	9.95
1.25	7.11
1.5	5.55
1.75	4.51
2	3.86
2.25	3.36
2.5	3.00
2.75	2.72
3	2.49
3.5	2.17
4	1.96
4.5	1.80
5	1.62



Monte Carlo Simulation Run Length Standard Deviation and Percentiles (Two-Sided, Zero-State)										
Shift in Mean (Multiple of Sigma)	Standard Deviation	1	5th	10th	25th	50th (Median)	75th	90th	95th	99th
0	366.35	9	24	45	110	255	513	860	1098	1662
0.25	115.82	6	12	18	40	89	167	276	359	541
0.5	28.17	4	7	9	15	26	45	72	90	136
0.75	10.50	3	5	6	9	13	21	30	37	53
1	5.30	3	4	5	6	9	12	17	20	28
1.25	3.22	2	3	4	5	6	9	11	13	18
1.5	2.22	2	3	3	4	5	7	8	10	13
1.75	1.60	2	2	3	3	4	5	7	8	9
2	1.27	2	2	2	3	4	5	5	6	8
2.25	1.03	2	2	2	3	3	4	5	5	6
2.5	0.87	2	2	2	2	3	3	4	5	5
2.75	0.73	2	2	2	2	3	3	4	4	5
3	0.64	1	2	2	2	2	3	3	4	4
3.5	0.48	1	2	2	2	2	2	3	3	3
4	0.40	1	1	2	2	2	2	2	3	3
4.5	0.43	1	1	1	2	2	2	2	2	3
5	0.49	1	1	1	1	2	2	2	2	2

The additional run length statistics show the large variation of run length values. The $MRL_0 = 255$ (in-control median run length with 0 shift in process mean).

Note: The results will vary slightly since this is Monte Carlo simulation.

14. We will now assess robustness to nonnormality using Monte Carlo simulation and compare to Shewhart and EWMA charts. Enter **Specify = k & ARL0**, **k parameter = 0.5**, **In-Control Average Run Length (ARL0) = 370.4**, **Fast Initial Response FIR = 0**, **Method = Monte Carlo**, **Number of Replications = 1e4**, **Skewness = 2**, **Kurtosis (Normal is 0) = 6**.

CUSUM Average Run Length (ARL) Calculator

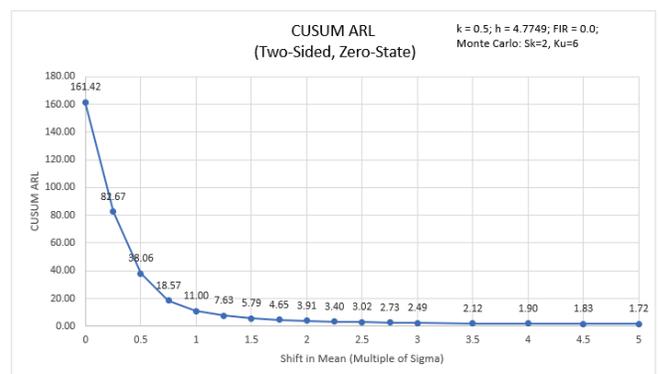
User Inputs:

Specify:	k & ARL0	Calculate CUSUM ARL
k parameter:	0.5	
In-Control Average Run Length (ARL0):	370.4	
Fast Initial Response:	FIR = 0	
Method:	Monte Carlo	
Number of Replications:	1.00E+04	
Skewness:	2	
Kurtosis (Normal is 0):	6	

15. Click the **Calculate CUSUM ARL** button to produce the CUSUM Parameters, Monte Carlo approximate ARL table, ARL chart and Run Length Standard Deviation and Percentiles table:

CUSUM Parameters:	
k	0.5
h	4.7749
FIR	0.0
Method	Monte Carlo: Sk=2, Ku=6

Shift in Mean (Multiple of Sigma)	CUSUM ARL (Two-Sided, Zero-State)
0	161.42
0.25	82.67
0.5	38.06
0.75	18.57
1	11.00
1.25	7.63
1.5	5.79
1.75	4.65
2	3.91
2.25	3.40
2.5	3.02
2.75	2.73
3	2.49
3.5	2.12
4	1.90
4.5	1.83
5	1.72



Shift in Mean (Multiple of Sigma)	Standard Deviation	Monte Carlo Simulation Run Length Standard Deviation and Percentiles (Two-Sided, Zero-State)								
		1	5th	10th	25th	50th (Median)	75th	90th	95th	99th
0	161.72	3	10	18	47	112	223	363	480	757
0.25	80.11	3	7	12	26	57	113	187	243	372
0.5	33.56	2	5	8	15	28	50	82	106	157
0.75	13.19	2	4	6	9	15	24	36	44	63
1	5.96	2	3	4	7	10	14	19	22	30
1.25	3.37	2	3	4	5	7	10	12	14	17
1.5	2.18	2	3	3	4	6	7	9	10	11
1.75	1.54	1	2	3	4	5	6	7	7	8
2	1.16	1	2	2	3	4	5	5	6	6
2.25	0.93	1	2	2	3	3	4	5	5	5
2.5	0.77	1	2	2	3	3	4	4	4	4
2.75	0.66	1	2	2	2	3	3	3	4	4
3	0.58	1	2	2	2	3	3	3	3	3
3.5	0.49	1	1	2	2	2	2	3	3	3
4	0.31	1	1	1	2	2	2	2	2	2
4.5	0.38	1	1	1	2	2	2	2	2	2
5	0.45	1	1	1	1	2	2	2	2	2

ARL₀ is approximately 161.4 which is a 2.3 x increase (370.4/161.4) in false alarms compared to Normal but is a much better performance than the ARL₀ = 54.6 result for [Shewhart Individuals](#).

It is, however, less robust to nonnormality than [EWMA chart with Weight \(Lambda\) = 0.1](#) which had a Monte Carlo ARL₀ = 266.4.

MRL₀ is approximately 112.

Stoumbos and Reynolds [5] recommend setting h=6.148 (with k=0.5) as a way to improve the CUSUM robustness to non-normality.

Template Notes:

1. Specify CUSUM parameters: **k & h**, **k & ARL0** or **ARL0 & Shift** using the drop-down list. Parameters to be specified will be shown in yellow highlight, otherwise they are hidden.
2. If applicable, enter the CUSUM parameter **k**. This is the reference (or slack) value, typically set to 0.5. It sets the size of mean shift ($2k$ sigma) that you would like to detect quickly, so 0.5 denotes rapid detection of a shift in mean = 1 sigma. Alternatively, Woodall & Faltin [4] recommend larger k values (e.g., $k = 0.9$) to avoid false alarms and detect shifts of practical significance.
3. If applicable, enter the CUSUM parameter **h**. This is the decision interval, typically set to 4 or 5.
4. If applicable, enter the desired **In-Control Average Run Length (ARL0)**. This will be the target ARL for mean shift = 0 and should be a large value to minimize false alarms, typically 370 to 500. The h parameter will be solved to achieve this ARL0, given a specified k value.
5. If applicable, enter the desired **Shift in Mean to Detect (Multiple of Sigma)**. This will minimize ARL for the given shift. If $FIR=0$, then k will be $shift/2$. If $FIR=h/2$, then k will be optimized, requiring about 20-30 seconds to compute.
6. Select **FIR=0** or **FIR=h/2** using the drop-down list. This is the fast initial response (or headstart) value.
7. FIR sets the initial CUSUM statistic so that it improves the sensitivity to a mean shift at startup. Note that if the process is in control when the CUSUM is started but shifts out of control later, the more appropriate ARL for such a case would be $FIR=0$. See Montgomery [3], pages 426-427.
8. Select Method: **Markov Chain** or **Monte Carlo** using the drop-down list. Markov Chain approximation is fast and accurate to compute ARLs. Monte Carlo simulation allows you to assess robustness to nonnormality and also produces the table of Run Length Standard Deviation and Percentiles (scroll right to view).
9. For further details on the Markov Chain approximation see Hawkins [1] and Lucas [2]. Monte Carlo simulation uses the Pearson Family of distributions to match the specified skewness and kurtosis.
10. If applicable, enter **Number of Replications**. 1000 ($1e3$) replications will be fast, approx. 10 seconds, but will have an ARL0 error approx. = +/- 10%; 10,000 ($1e4$) replications will take about a minute, with an ARL0 error = +/- 3.2%; 100,000 ($1e5$) replications will take about ten minutes, with an ARL0 error = +/- 1%.
11. If applicable, enter **Skewness**. Skewness must be ≥ 0 . Skewness = 0 is symmetric.
12. If applicable, enter **Kurtosis (Normal is 0)**. Also known as Excess Kurtosis, it must be $\geq Skewness^2 - 1.48$. This is required to keep the distribution unimodal. If Skewness=0 and Kurtosis = 0, the distribution is normal.
13. Click the **Calculate CUSUM ARL** button to produce the ARL table and chart. If Monte Carlo was selected, the table of Run Length Standard Deviation and Percentiles will also be produced.
14. The CUSUM ARL is for a two-sided chart with zero-state, i.e., the shift is assumed to occur at the start. The mean and standard deviation are also assumed to be known. This will not likely be the case in use, but is still useful for determining parameter settings and comparison of ARL across chart types.

15. Due to the complexity of calculations, SigmaXL must be loaded and appear on the menu in order for this template to function. Do not add or delete rows or columns in this template.

REFERENCES:

- [1] Hawkins, D. M. and Olwell, D. H. (1998), *Cumulative Sum Charts and Charting for Quality Improvement (Information Science and Statistics)*, Springer, New York.
- [2] Lucas, J.M. and Crosier R.B. (1982), "Fast Initial Response for CUSUM Quality-Control Schemes: Give Your CUSUM A Headstart", *Technometrics* 24, 199-205.
- [3] Montgomery, D.C. (2013), *Introduction to Statistical Quality Control*, Seventh Ed., Wiley.
- [4] Woodall, W. H. and Faltin, F.W. (2019), "Rethinking control chart design and evaluation", *Quality Engineering* 31, 596-605.
- [5] Stoumbos, Z. G. and Reynolds, M.R. Jr. (2004), "The Robustness and Performance of CUSUM Control Charts Based on the Double-Exponential and Normal Distributions", In: Lenz, H. J., Wilrich, P. T. (eds) *Frontiers in Statistical Quality Control 7*, Physica, Heidelberg, 79-100.

Part K – Time Series Forecasting and Control Charts for Autocorrelated Data

Introduction – Time Series Forecasting

From https://en.wikipedia.org/wiki/Time_series:

A time series is a series of data points indexed (or listed or graphed) in time order. Most commonly, a time series is a sequence taken at successive equally spaced points in time.

Time series analysis comprises methods for analyzing time series data in order to extract meaningful statistics and other characteristics of the data. Time series forecasting is the use of a model to predict future values based on previously observed values.

The free online book by Rob Hyndman and George Athanasopoulos, “Forecasting: Principles and Practice,” (<https://otexts.com/fpp2/>) is an excellent introductory resource to time series analysis (further denoted as “fpp2”). The book references the R forecast package (<https://cran.r-project.org/web/packages/forecast/index.html>), but the results given by SigmaXL will closely match those in the book. Note that SigmaXL is independent of R. Slight differences in results are due to minor differences in algorithm options and optimization methods.

The following Wikipedia articles are also helpful resources:

https://en.wikipedia.org/wiki/Exponential_smoothing

https://en.wikipedia.org/wiki/Autoregressive_integrated_moving_average

Box and Jenkins, et al., 2016, is the classic book on ARIMA modeling (further denoted as “Box and Jenkins”):

Box, G.E.P., Jenkins, G.M., Reinsel, G.C. and Ljung, G.M. (2016). *Time Series Analysis, Forecasting and Control*, 5th edition, Wiley.

Datasets to be used in the examples include:

1. Chemical Process Concentration – Series A (Box and Jenkins)
2. Monthly Airline Passengers – Series G (Box and Jenkins)
3. Monthly Airline Passengers - Missing Values
4. Monthly Airline Passengers - Modified for Control Charts
5. Daily Electricity Demand with Predictors – ElecDaily (fpp2)
6. Sales with Indicator – Modified Series M (Box and Jenkins)
7. Half-Hourly Multiple Seasonal Electricity Demand – Taylor (R forecast)

SigmaXL provides the following tools for exploratory data analysis of time series data:

- Run Chart
- Autocorrelation Function (ACF)/Partial Autocorrelation (PACF) Plots
- Cross Correlation (CCF) Plots with Pre-Whiten Data option
- Seasonal Trend Decomposition Plots
- Spectral Density Plot with Detection of Seasonal Frequency

SigmaXL provides the following methods for time series forecasting:

- Exponential Smoothing
- Exponential Smoothing – Multiple Seasonal Decomposition (MSD)
- ARIMA – Box-Jenkins Autoregressive Integrated Moving Average
- ARIMA with Predictors
- ARIMA – MSD

Typically, either Exponential Smoothing or ARIMA may be used. It may be useful to try both to see which one gives a better model or use the average of the forecast from both methods. If the data has negative autocorrelation, ARIMA is recommended. If the data includes continuous or categorical predictors, use ARIMA with Predictors.

ARIMA assumes that the time series is stationary, i.e., it has the property that the mean, variance and autocorrelation structure do not change over time. If a time series mean is not stationary (e.g. trending), this can be corrected by differencing, computing the differences between consecutive observations for nonseasonal and between consecutive periods for seasonal data (e.g., Jan 2019 – Jan 2018, etc.). If the variance changes over time, a Box-Cox transformation may be applied to achieve constant variance. Exponential Smoothing does not require stationarity.

If the data are seasonal (i.e., influenced by seasonal factors), SigmaXL requires that the seasonal frequency be specified. Frequency is the number of observations per “cycle” unit of time, so monthly sales would be specified as seasonal frequency = 12 (observations per year). Quarterly revenue would be specified as seasonal frequency = 4. Hourly data would be 24 (observations per day). Note that “seasonal frequency” is also referred to as “seasonal period”, but SigmaXL follows the convention as used in Hyndman and Athanasopoulos (fpp2).

Exponential Smoothing is limited to a maximum seasonal frequency of 24. For higher frequencies use Exponential Smoothing – Multiple Seasonal Decomposition (MSD). In MSD the seasonal component is first removed through decomposition, a nonseasonal exponential smooth model fitted to the remainder (+trend), and then the seasonal component is added back in. For forecasting, a naïve seasonal forecast is used on the seasonal component. As the name implies, Multiple Seasonal Decomposition (MSD) also accommodates multiple seasonality.

ARIMA does not have a theoretical frequency limit, but for computational efficiency and to minimize the potential loss of observations through differencing, we recommend using ARIMA – MSD for seasonal frequency greater than 52 (or with multiple frequencies). Note, ARIMA with Predictors – MSD is not available.

Details of the time series forecast methods and formulas used in SigmaXL are given in the Appendix: [Time Series Forecasting and Control Charts for Autocorrelated Data](#). References are given at: [References for Time Series Forecasting and Control Charts for Autocorrelated Data](#).

Introduction – Control Charts for Autocorrelated Data

Statistical process control for autocorrelated processes have typically used the EWMA (Exponentially Weighted Moving Average) one-step-ahead forecast model. The time series model forecasts the motion in the mean and an Individuals control chart is plotted of the residuals to detect assignable causes. SigmaXL extends this concept to include all of the forecast methods mentioned above.

Woodall and Faltin give some helpful guidelines on dealing with autocorrelation [\[20\]](#):

- If possible, one should first attempt to remove the source of the autocorrelation.
- If the source of autocorrelation cannot be removed directly, then it may be possible to model the autocorrelation and use a feedback control scheme to reduce variability about a specified target value.
- If the source of the autocorrelation cannot be removed directly, and feedback control is not a viable option, then it is important to monitor the process with control charts which do not repeatedly give signals due to presence of the autocorrelation.

Failure to account for positive autocorrelation will produce limits that are too narrow resulting in excessive false alarms. Failure to account for negative autocorrelation will produce limits that are too wide resulting in misses. An accurate forecast for your time series means that the residuals will most often have the right properties to correctly apply a control chart, thus leading to an improved control chart with reduced false alarms and misses.

In addition to creating an Individuals control chart for residuals, a Moving Limits chart is included, which uses the one step prediction as the center line, so the control limits move with the center line.

The popular “Add Data”, “Show Last 30” and “Scroll” features in SigmaXL Chart Tools are available for these control charts. For “Add Data”, the time series models are not refitted, but used to compute the residual values for the new data.

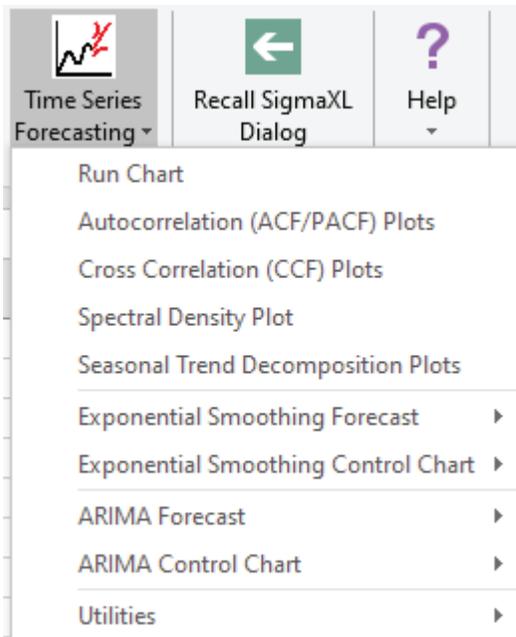
For further details and references, see the Appendix: [Control Charts for Autocorrelated Data](#).

Summary of Features in Time Series Forecasting and Control Charts for Autocorrelated Data

- Run Chart
- Autocorrelation Function (ACF)/Partial Autocorrelation (PACF) Plots
- Cross Correlation (CCF) Plots with Pre-Whiten Data option
- Seasonal Trend Decomposition Plots
- Spectral Density Plot with Detection of Seasonal Frequency
- Exponential Smoothing:
 - Forecast with Prediction Intervals
 - Exponential Smoothing models use Rob Hyndman's taxonomy:
 - Additive/Multiplicative Error
 - Additive/Additive Damped Trend
 - Additive/Multiplicative Seasonal
 - This includes all of the classical exponential smoothing models such as Simple/Single/EWMA, Double and Holt-Winters
 - Multiple Seasonal Decomposition (MSD) option
 - Useful for high frequency and/or multiple frequency data, such as Monthly with frequency = 12, Daily with frequency = 7 and Hourly with frequency = 24
- Exponential Smoothing Residuals Control Chart for autocorrelated data:
 - Individuals and Moving Limits (with One-Step Ahead Forecast) Charts
 - Add Data, Show Last 30 Data Points, Enable Scroll options
 - MSD option
- Autoregressive Integrated Moving Average (ARIMA):
 - Forecast with Prediction Intervals
 - ARIMA Forecast with Predictors (Continuous and/or Categorical)
 - MSD option
- ARIMA Residuals Control Chart for autocorrelated data:
 - Individuals and Moving Limits (with One-Step Ahead Forecast) Charts
 - ARIMA Control Chart with Predictors
 - Add Data, Show Last 30 Data Points, Enable Scroll options
 - MSD option
- Utilities:
 - Difference Data
 - Lag Data
 - Interpolate Missing Values (seasonally adjusted linear interpolation)
- Time Series Forecasting Model Features:
 - ARIMA and Exponential Smoothing models are fully automatic or user specified
 - Utilizes modern State Space and Kalman Filter models for accurate parameter estimation
 - ARIMA estimates missing values with Kalman Filter; Exponential Smoothing uses seasonally adjusted linear interpolation
 - Automatic Box-Cox Transformation
 - Automatic seasonal frequency detection

- Model Diagnostics:
 - ACF/PACF Plots
 - Ljung-Box p-values
 - Log-Likelihood, AIC, AICc, BIC, Residual StDev
 - Residual plots (histogram, normal probability, residual versus fits, residuals versus order)
- Forecast Accuracy:
 - In-Sample (Estimation) one-step-ahead forecast errors (RMSE, MAE, MASE, MAPE)
 - Out-of-Sample (Withhold) one-step-ahead forecast errors
 - Out-of-Sample (Withhold) multi-step-ahead forecast errors
 - Evaluated using the benchmark standard M4 forecast competition data, a total of 100,000 data sets with Yearly, Quarterly, Monthly, Weekly, Daily and Hourly data. Using a hybrid average of automatic Exponential Smoothing and ARIMA, SigmaXL (unofficially) ranked 10th out of 60 in the Overall Weighted Average forecast accuracy score, ahead of three well known commercial forecast software packages.

Time Series Forecasting Menu

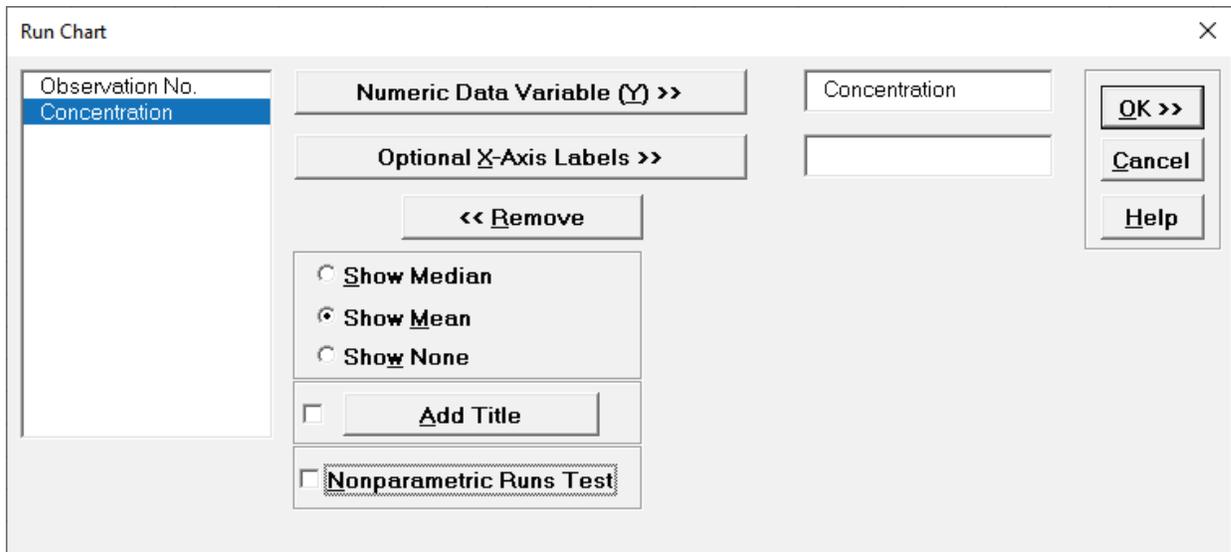


Run Chart

The Run Chart in Time Series Forecasting is the same tool as used in Graphical Tools > Run Chart (see [Part H – Run Charts](#)). We will demonstrate the use of Run Charts for initial exploratory data analysis of the time series examples, as well as the use of the difference and interpolate utilities, and use of SigmaXL graphical tools for data with predictors.

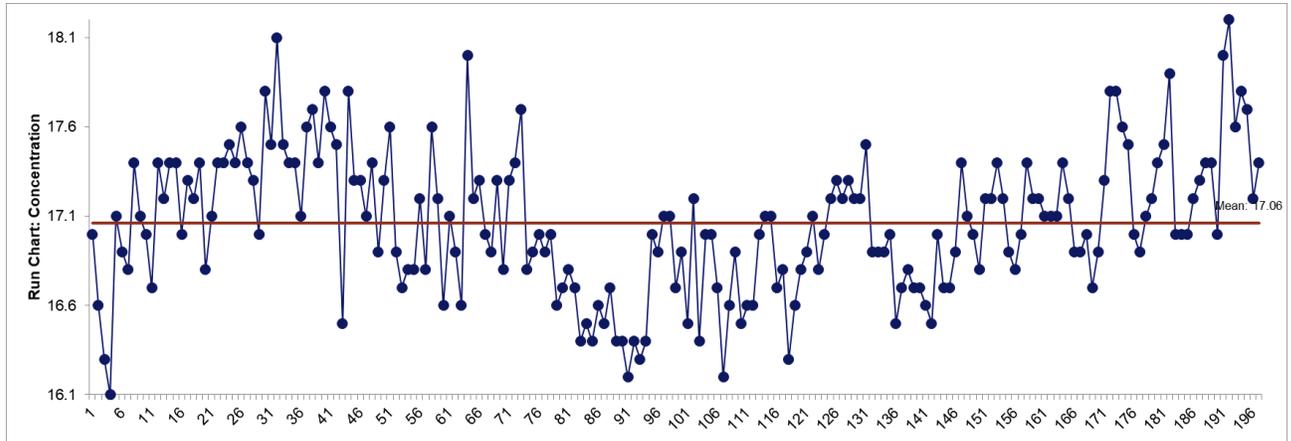
Chemical Process Concentration - Series A

1. Open **Chemical Process Concentration – Series A.xlsx (Sheet 1 tab)**. This is the Series A data from Box and Jenkins, a set of 197 concentration values from a chemical process taken at two-hour intervals.
2. Click **SigmaXL > Time Series Forecasting > Run Chart**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
3. Select *Concentration*, click **Numeric Data Variable (Y) >>**. Select **Show Mean**. Uncheck **Nonparametric Runs Test**.



Tip: Show Mean is selected for consistency with the use of the mean in forecasting. If you wish to use the **Nonparametric Runs Test**, then use the default **Show Median** to display the Median rather than the Mean.

4. Click **OK**. A Run Chart of Concentration with Mean center line is produced.

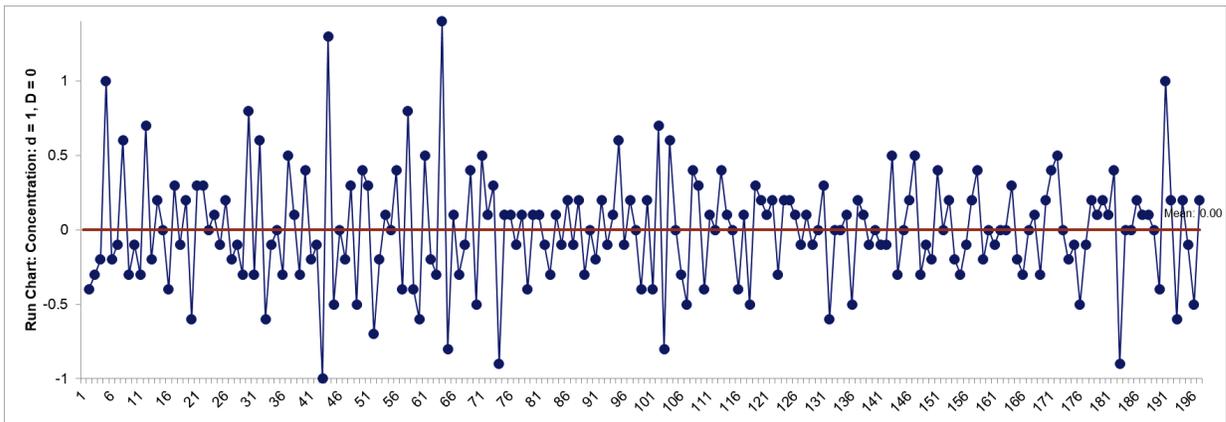


5. Clearly, this is a non-stationary process with a “wandering mean”. Exponential smoothing can model this, but ARIMA requires a stationary process. We will now demonstrate how differencing can be used to make the process stationary. (Note that ARIMA Forecast will automatically take care of this).
6. Click the **Sheet 1** tab. Click **SigmaXL > Time Series Forecasting > Utilities > Difference Data**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
7. Select *Concentration*, click **Numeric Time Series Data (Y) >>**. Enter 1 for **Nonseasonal Differencing (d)**.

8. Click **OK**. A new sheet is created with the order $d = 1$ differenced data ($Y_2 - Y_1, Y_3 - Y_2, \dots$).

Concentration	Concentration: d = 1, D = 0
17	
16.6	-0.4
16.3	-0.3
16.1	-0.2
17.1	1
16.9	-0.2
16.8	-0.1
17.4	0.6
17.1	-0.3
17	-0.1
16.7	-0.3
17.4	0.7

9. Redo the Run Chart for the differenced data:



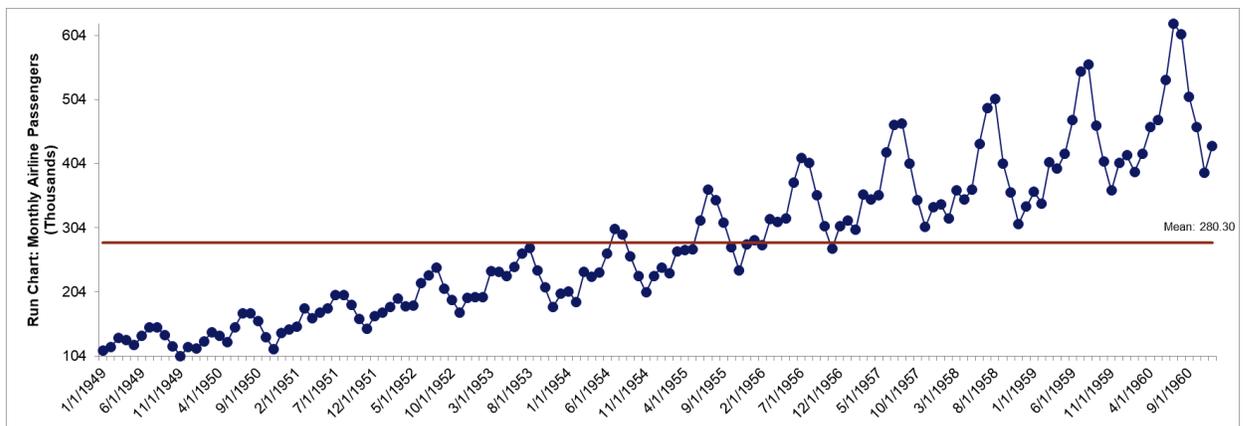
This now appears to be a stationary process and can be modeled using ARIMA. Save the workbook with the differenced data for use with Autocorrelation Plots.

Monthly Airline Passengers - Series G

10. Open **Monthly Airline Passengers - Series G.xlsx (Sheet 1 tab)**. This is the Series G data from Box and Jenkins, monthly total international airline passengers for January 1949 to December 1960 and is one of the most popular datasets used in introductory time series forecasting.
11. Click **SigmaXL > Time Series Forecasting > Run Chart**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
12. Select *Monthly Airline Passengers (Thousands)*, click **Numeric Data Variable (Y) >>**; select *Date*, click **Optional X-Axis Labels >>**. Select **Show Mean**. Uncheck **Nonparametric Runs Test**.

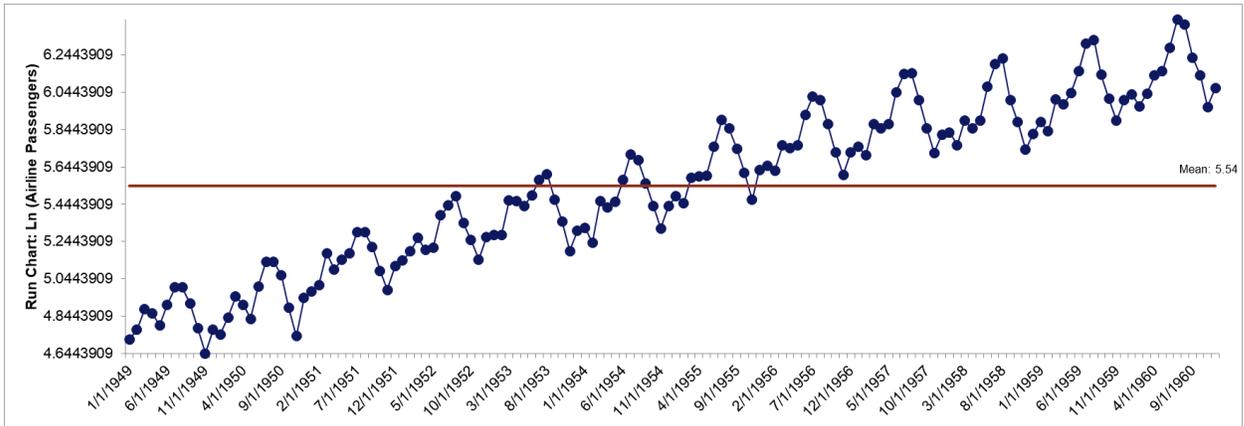
The screenshot shows the 'Run Chart' dialog box. On the left, a list of variables includes 'Obs. No.', 'Date', 'Monthly Airline Passengers', and 'Ln (Airline Passengers)'. The 'Ln (Airline Passengers)' variable is selected. In the center, there are two main sections: 'Numeric Data Variable (Y) >>' with 'Monthly Airline Passengers' entered, and 'Optional X-Axis Labels >>' with 'Date' entered. Below these are radio buttons for 'Show Median', 'Show Mean' (which is selected), and 'Show None'. There is also an 'Add Title' checkbox and a 'Nonparametric Runs Test' checkbox which is unchecked. On the right side, there are buttons for 'OK >>', 'Cancel', and 'Help'.

13. Click **OK**. A Run Chart of Monthly Airline Passengers with Mean center line is produced.



14. Here we see a strong positive trend as well as a monthly seasonal effect (yearly cycle of 12 months), with an increase in the variance over time. This non-stationarity in the variance can be addressed by using a Ln transformation of the data. Later, we will demonstrate the use of a Box-Cox Transformation to automatically determine the best transformation, but for now we will use the Ln(Airline Passenger) data.

15. Click the **Sheet 1** tab. Redo the Run Chart for *Ln (Airline Passengers)*:



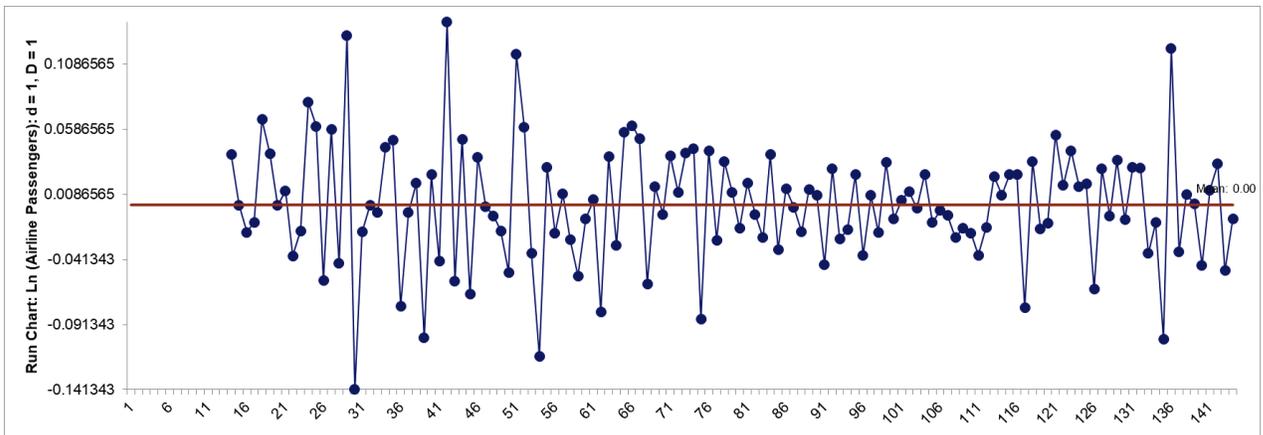
- 16. This addresses the change in variance over time, but still needs to be differenced in order to make it stationary for ARIMA. (Note that ARIMA Forecast will automatically take care of this).
- 17. Click the **Sheet 1** tab. Click **SigmaXL > Time Series Forecasting > Utilities > Difference Data**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
- 18. Select *Ln(Airline Passengers)*, click **Numeric Time Series Data (Y) >>**. Enter 1 for **Nonseasonal Differencing (d)**; enter 1 for **Seasonal Differencing (D)**; **Seasonal Frequency** is specified as 12.

19. Click **OK**. A new sheet is created with the seasonal order $D = 1$ differenced data ($Y_{13} - Y_1, Y_{14} - Y_2, \dots$) and nonseasonal order $d=1$ ($D_2 - D_1, D_3 - D_2, \dots$).

Ln (Airline Passengers)	Ln (Airline Passengers): d = 1, D = 1
4.718498871	
4.770684624	
4.882801923	
4.859812404	
4.795790546	
4.905274778	
4.997212274	
4.997212274	
4.912654886	
4.779123493	
4.644390899	
4.770684624	
4.744932128	
4.836281907	0.039164025
4.94875989	0.000360685
4.905274778	-0.020495594
4.828313737	-0.012939182

Tip: For an example of manual seasonal and nonseasonal differencing in Excel see https://faculty.fuqua.duke.edu/~rnau/Decision411_2007/Class10notes.htm.

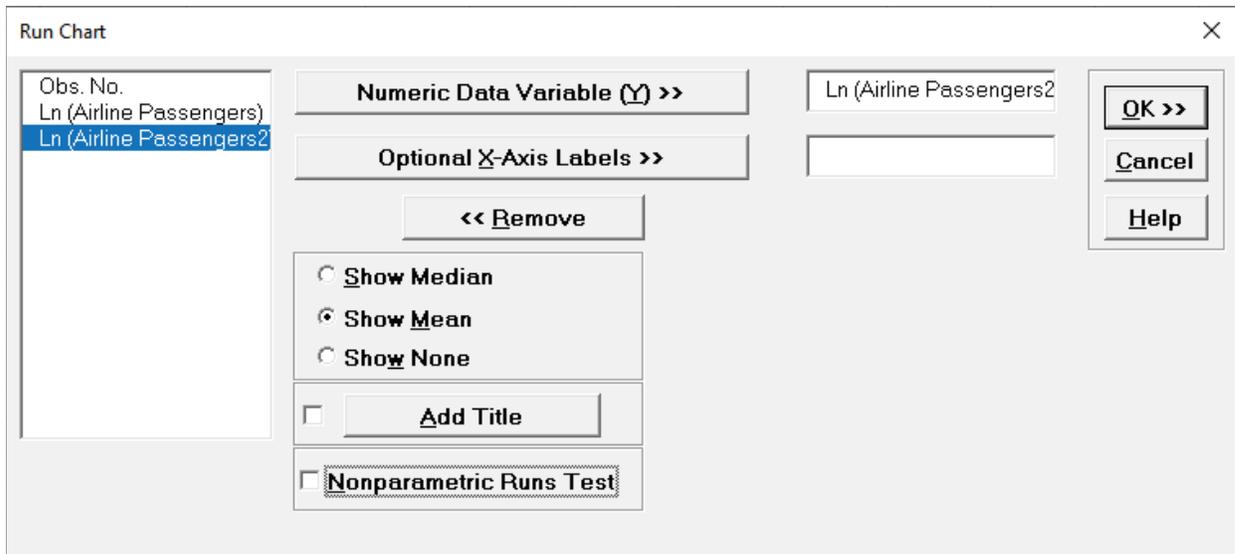
20. Redo the Run Chart for the differenced data:



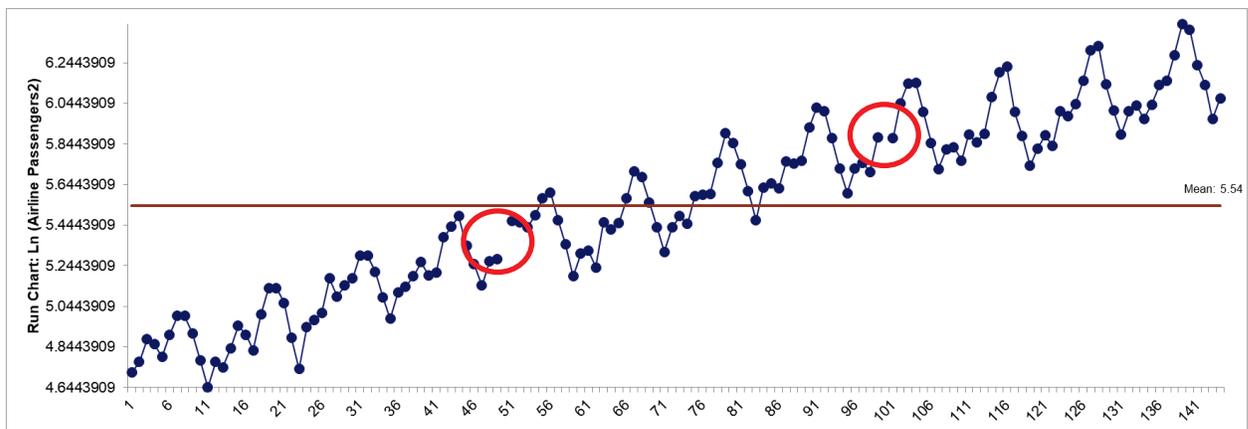
This now appears to be a stationary process and can be modeled using ARIMA. Save the workbook with the differenced data for use with Autocorrelation Plots.

Monthly Airline Passengers – Missing Values

21. Open **Monthly Airline Passengers – Missing Values.xlsx** (Sheet 1 tab). Ln(Airline Passengers2) have missing values at observations 50 and 100.
22. Click **SigmaXL > Time Series Forecasting > Run Chart**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
23. Select Ln(Airline Passengers2), click **Numeric Data Variable (Y) >>**. Select **Show Mean**. Uncheck **Nonparametric Runs Test**.

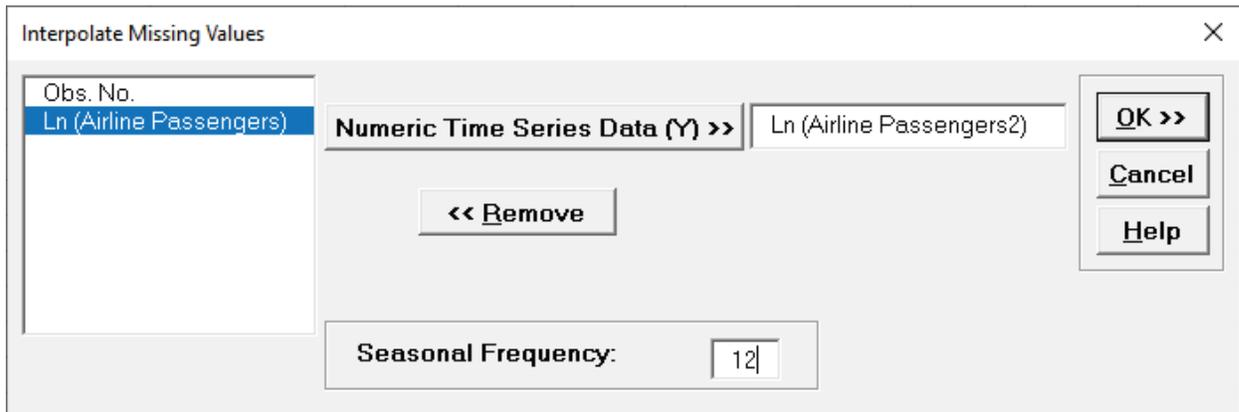


24. Click **OK**. A Run Chart of Ln(Airline Passengers2) with Mean center line is produced.



25. The missing values result in a broken line on the run chart.
26. Click the **Sheet 1** tab. Click **SigmaXL > Time Series Forecasting > Utilities > Interpolate Missing Values**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.

27. Select *Ln(Airline Passengers2)*, click **Numeric Time Series Data (Y) >>**. **Seasonal Frequency** is specified as 12.



28. Click **OK**. A new sheet is created with the seasonally adjusted linear interpolated values highlighted in yellow.

Ln (Airline Passengers2)	Y: Interpolated (12)
4.718498871	4.718498871
4.770684624	4.770684624
4.882801923	4.882801923
4.859812404	4.859812404
	5.254047791
	5.846836235

Y50 is estimated as 5.254 (versus original value of 5.278). Y100 is estimated as 5.847 (versus original value of 5.852).

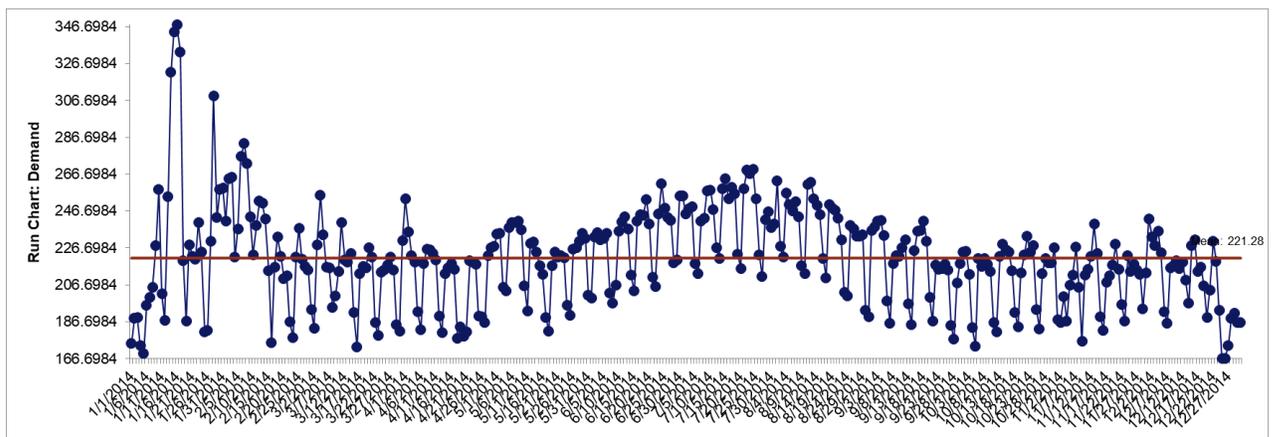
Note, while there is robustness against some missing values, if the number of missing values is large, then model estimation and forecast accuracy will be degraded.

Daily Electricity Demand with Predictors – ElecDaily

29. Open **Daily Electricity Demand with Predictors – ElecDaily.xlsx (Sheet 1 tab)**. This is daily electricity demand (GW) for the state of Victoria, Australia, every day during 2014 (Hyndman, fpp2). Temp (C) is the maximum daily temperature in degrees Celsius for the city of Melbourne. TempSq is Temperature squared. WorkDay takes on the value 1 on work days and 0 otherwise. This data has a seasonal frequency = 7 (observations per week).
30. Click **SigmaXL > Time Series Forecasting > Run Chart**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
31. Select *Demand*, click **Numeric Data Variable (Y) >>**; select *Date*, click **Optional X-Axis Labels >>**. Select **Show Mean**. Uncheck **Nonparametric Runs Test**.

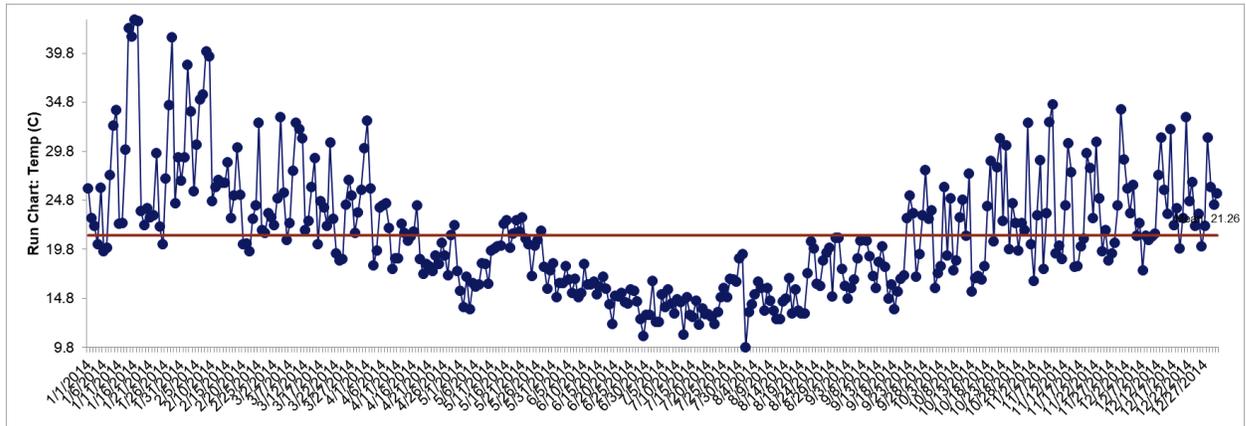
The screenshot shows the 'Run Chart' dialog box. On the left, a list of variables includes 'Date', 'Demand', 'Temp (C)', 'TempSq', and 'WorkDay', with 'Temp (C)' highlighted. In the center, there are two main sections: 'Numeric Data Variable (Y) >>' with 'Demand' entered, and 'Optional X-Axis Labels >>' with 'Date' entered. Below these are three radio buttons: 'Show Median' (unselected), 'Show Mean' (selected), and 'Show None' (unselected). There is also an 'Add Title' checkbox (unchecked) and a 'Nonparametric Runs Test' checkbox (unchecked). On the right side, there are three buttons: 'OK >>', 'Cancel', and 'Help'.

32. Click **OK**. A Run Chart of Electricity Demand with Mean center line is produced.



33. We can see the daily seasonality with 7-day cycle, but there is also a summer winter effect, with the peak demand occurring mid-January, which is summer in Australia. Later we will use a Spectral Density Plot to confirm the seasonal frequency value of 7.

34. Click the **Sheet 1** tab. Redo the Run Chart for Temp (C):

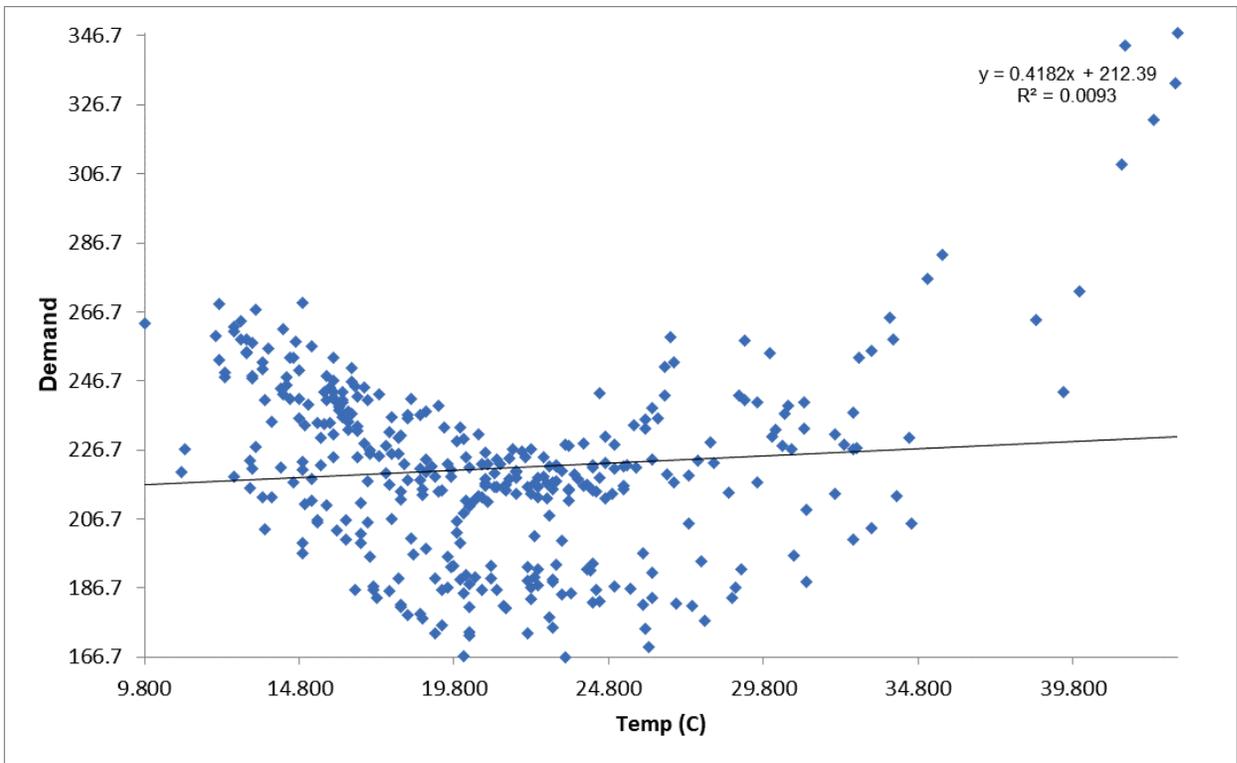


35. Now we will explore the relationship between Demand and Temperature. Click the **Sheet 1** tab. Click **SigmaXL > Graphical Tools > Scatterplots**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.

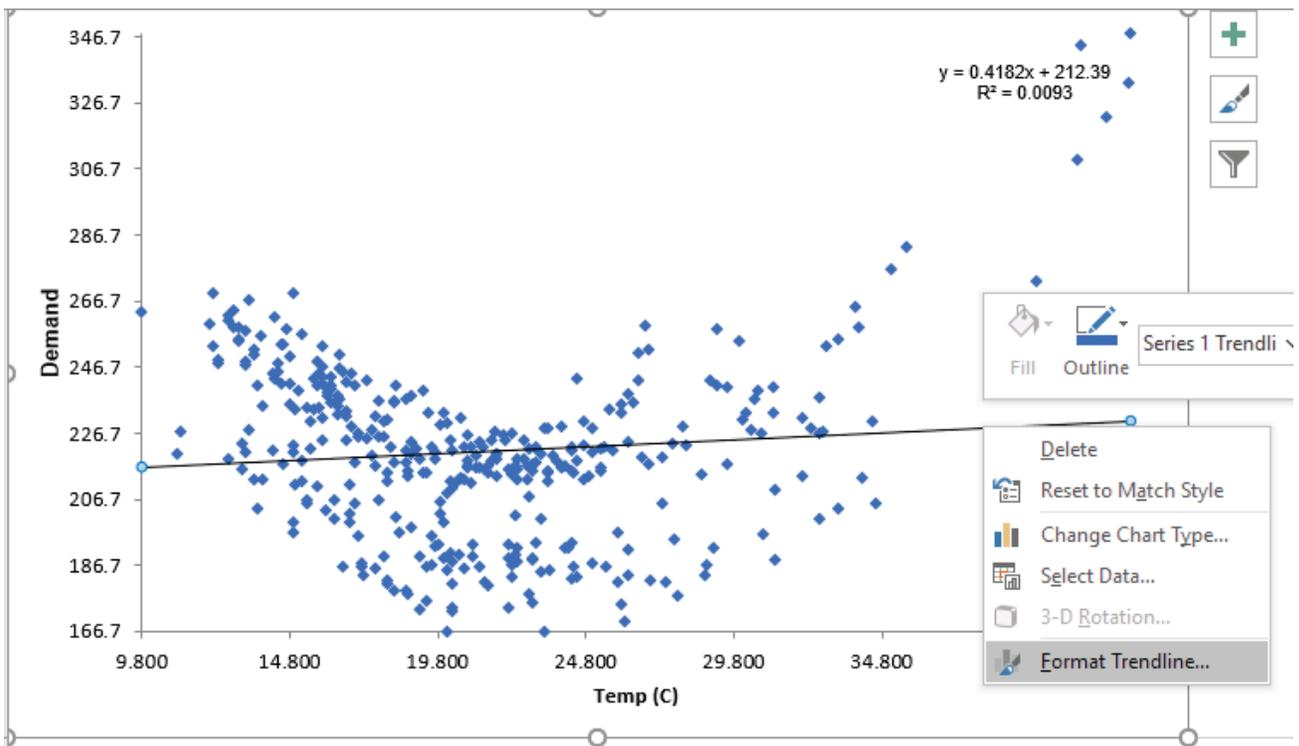
36. Select *Demand*, click **Numeric Response (Y) >>**; select *Temp (C)*, click **Numeric Predictor (X1) >>**. Check **Trendline**.

The image shows the 'Scatter Plots' dialog box in SigmaXL. On the left, a list of variables includes 'Date', 'TempSq', and 'WorkDay', with 'TempSq' selected. In the center, there are three main sections: 'Numeric Response (Y) >>' with 'Demand' entered, 'Numeric Predictor (X1) >>' with 'Temp (C)' entered, and 'Group Category (X2) >>' which is empty. Below these is a '<< Remove' button. On the right, there are 'OK >>', 'Cancel', and 'Help' buttons. At the bottom right, there is a 'Display Options' section with three checkboxes: 'Trendline' (checked), '95% Confidence Interval' (unchecked), and '95% Prediction Interval' (unchecked). Below this is an 'Add Title' checkbox which is also unchecked.

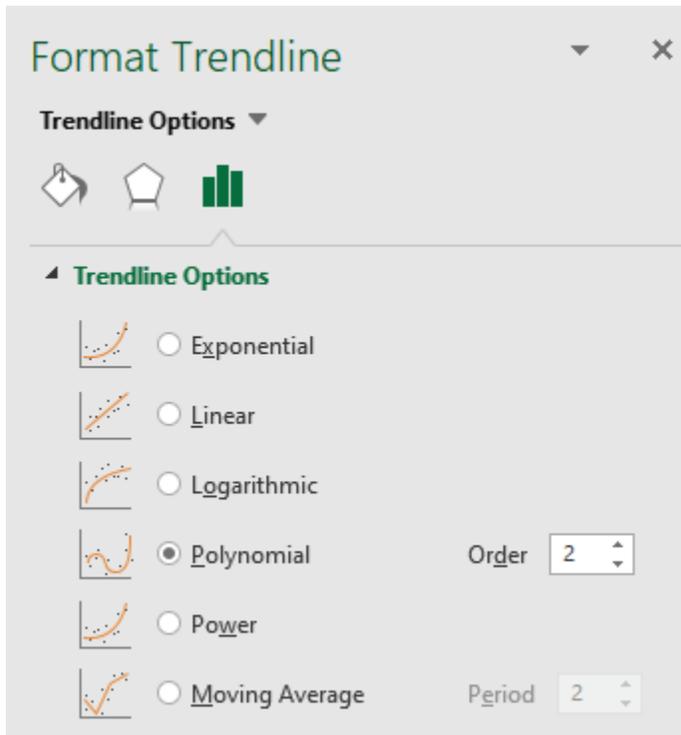
37. Click **OK**. A Scatterplot of Electricity Demand versus Temperature is produced.



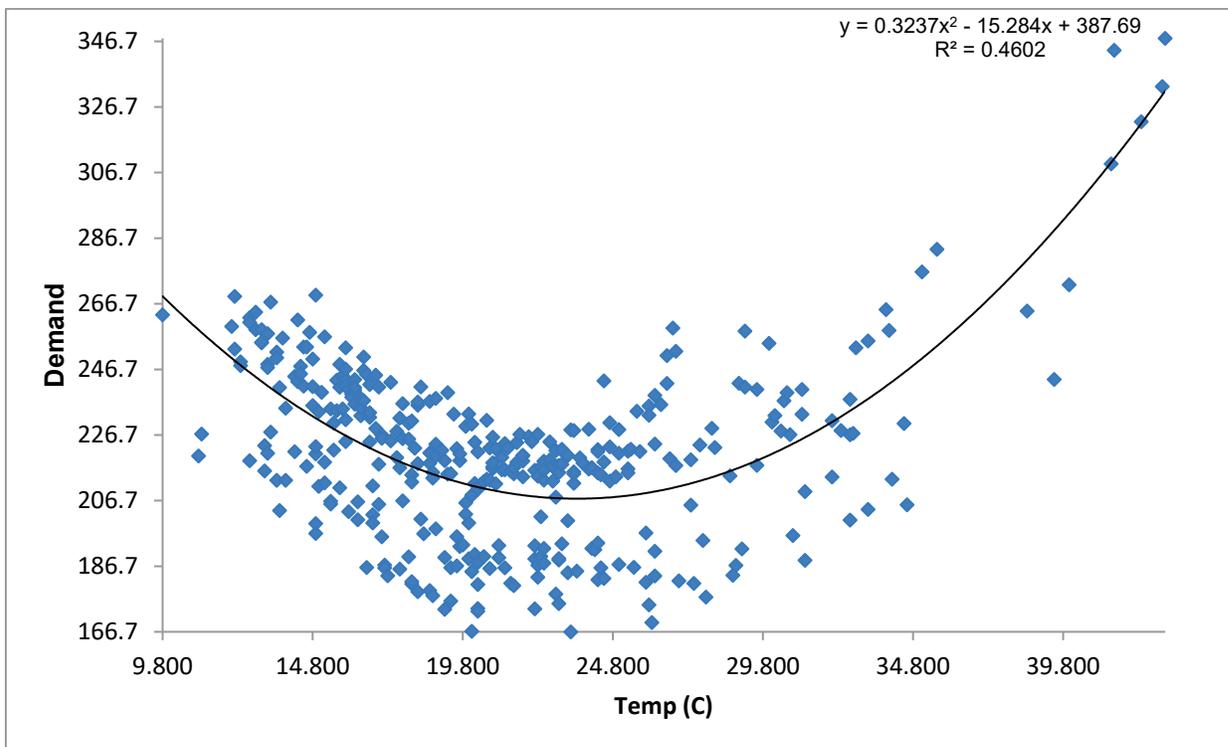
38. This shows a quadratic relationship: high temperatures in the summer cause electricity demand for air conditioning, low temperatures in the winter cause demand for heating. We will modify the trendline in Excel to a quadratic fit. Click on the Trendline, right click and select Format Trendline as shown:



The **Format Trendline** options are given. Select **Polynomial** with **Order 2** as shown.

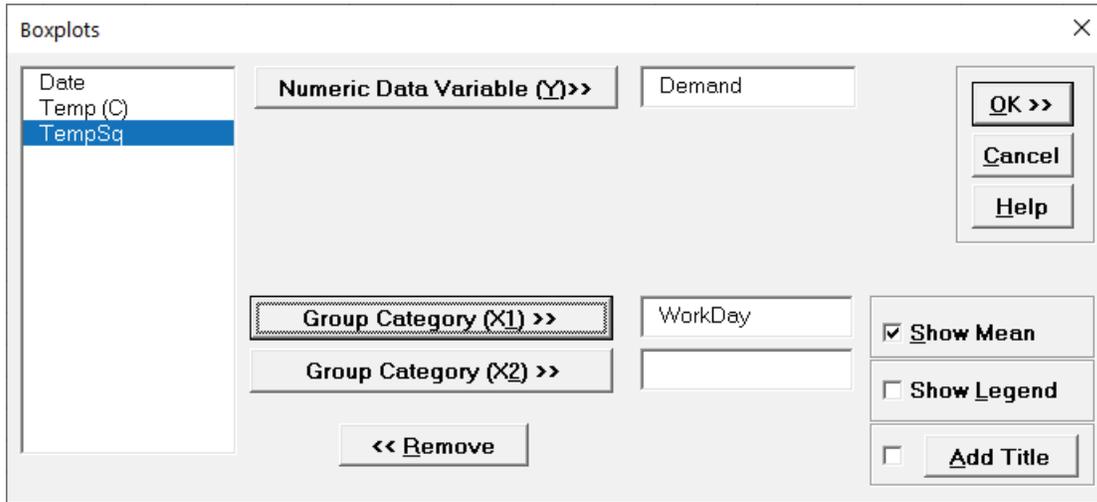


39. The Trend Line is now a quadratic function as shown:

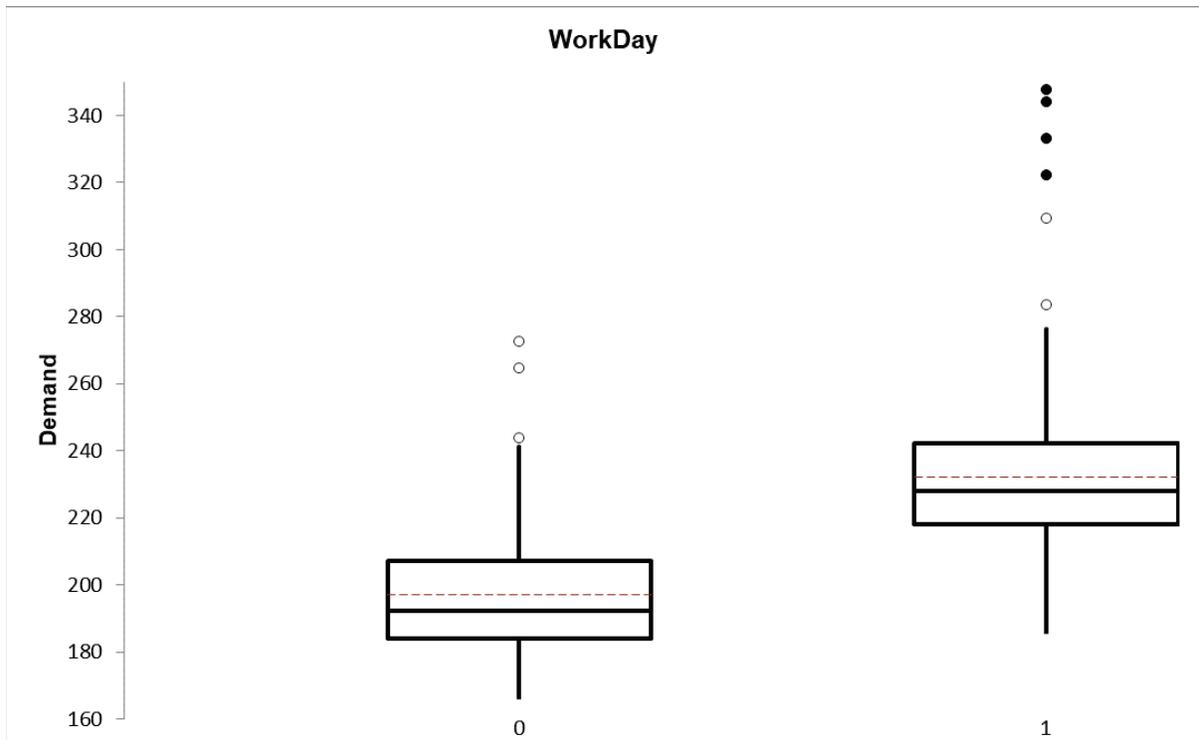


40. TempSq (Temperature Squared) has been added to the data so that it may be included when we do ARIMA forecast with predictors.

- Now we will explore the relationship between Demand and WorkDay. Click the **Sheet 1** tab. Click **SigmaXL > Graphical Tools > Boxplots**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
- Select *Demand*, click **Numeric Data Variable (Y) >>**; select *WorkDay*, click **Group Category (X1) >>**. Check **Show Mean**.



- Click **OK**. A Boxplot of Electricity Demand versus Work Day is produced.



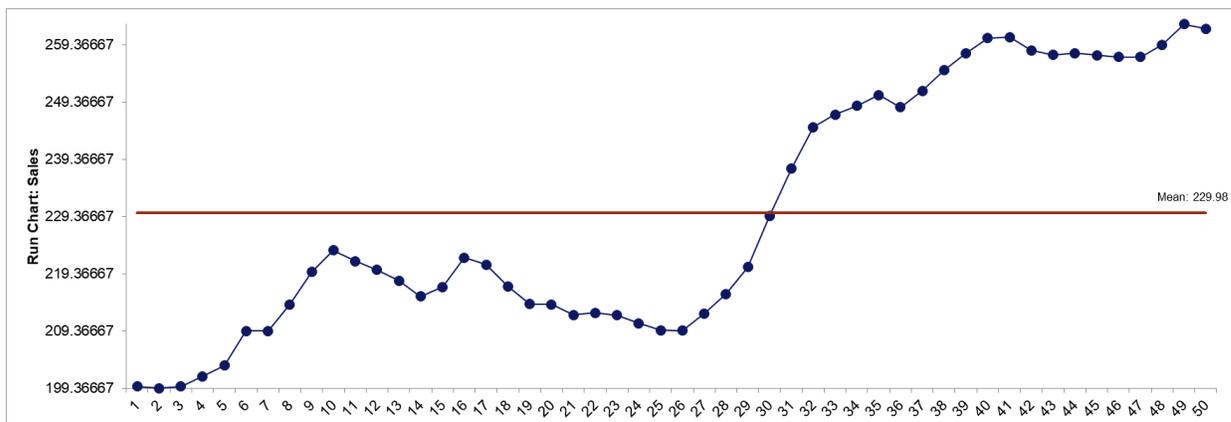
As expected, a work day (1) has higher electricity demand than a non-work day (0). Later, Temp (C), TempSq and WorkDay will be included as predictors in the ARIMA forecast model.

Sales with Indicator - Modified Series M

44. Open **Sales with Indicator - Modified Series M.xlsx**. (Sheet 1 tab). This is modified Series M data from Box and Jenkins. Originally, a set of 150 monthly corporate sales values along with a leading indicator, the data was modified by converting it to quarterly values by averaging every three months, so 50 quarters. This was done in order to simplify the analysis of the leading indicator. Although the data was monthly and summarized to quarterly, it will be treated as nonseasonal, as done in Box and Jenkins.
45. Click **SigmaXL > Time Series Forecasting > Run Chart**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
46. Select *Sales*, click **Numeric Data Variable (Y) >>**. Select **Show Mean**. Uncheck **Nonparametric Runs Test**.

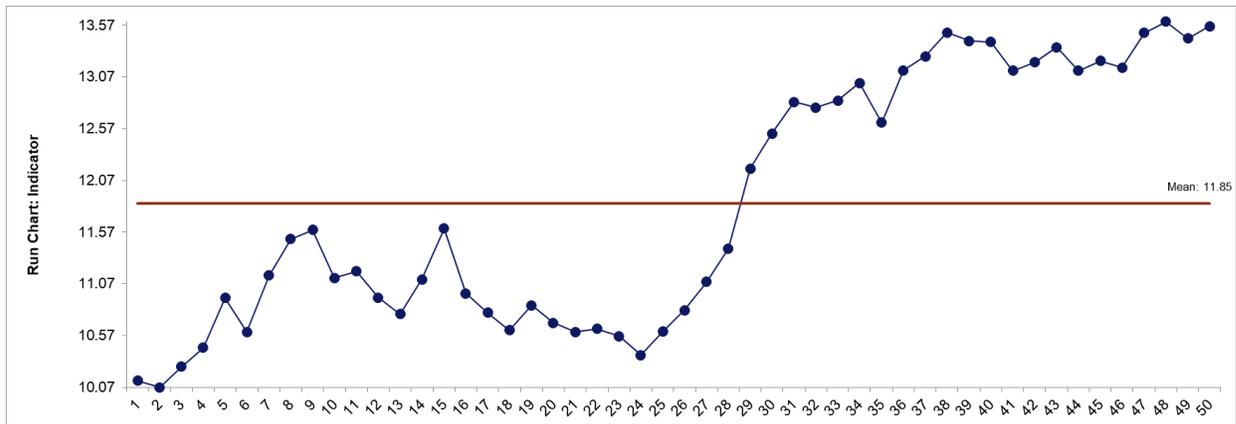
The screenshot shows the 'Run Chart' dialog box. On the left, a list of variables includes 'Qtr. Sales' and 'Indicator'. The 'Numeric Data Variable (Y) >>' button is highlighted, and the text 'Sales' is entered in the adjacent field. Below this, the 'Optional X-Axis Labels >>' button is visible. A '<< Remove' button is also present. In the center, three radio buttons are shown: 'Show Median', 'Show Mean' (which is selected), and 'Show None'. Below these, there is an 'Add Title' checkbox and a 'Nonparametric Runs Test' checkbox, both of which are unchecked. On the right side of the dialog, there are three buttons: 'OK >>', 'Cancel', and 'Help'.

47. Click **OK**. A Run Chart of Sales with Mean center line is produced.



As discussed above, there is no obvious seasonality in the data. Later we will use a Spectral Density Plot to confirm the seasonal frequency value of 1.

48. Click the **Sheet 1** tab. Redo the Run Chart for Indicator:

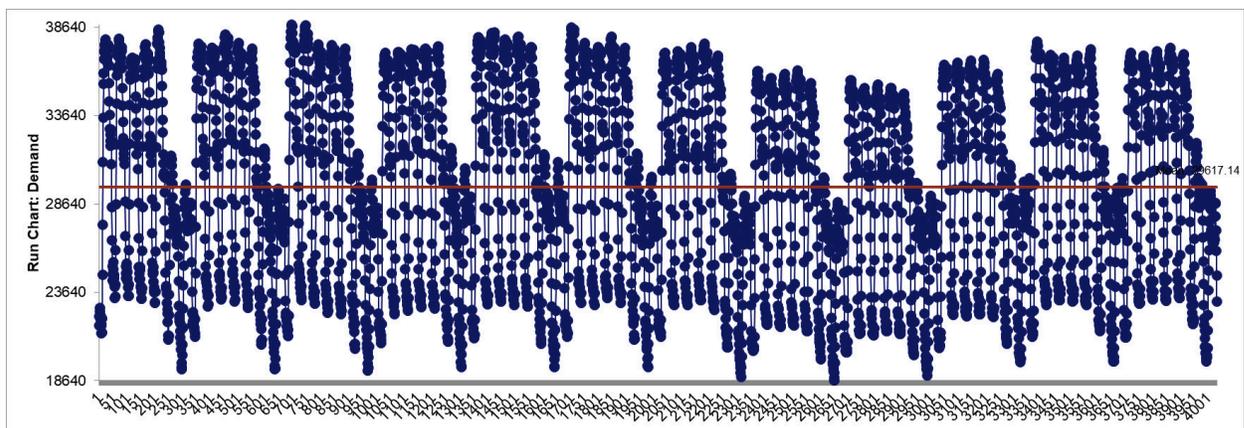


Clearly there is a close relationship between Sales and Indicator. We will analyze this relationship later with the Cross Correlation Function (CCF) Plot and with an ARIMA forecast model that includes Indicator as a predictor.

Half-Hourly Multiple Seasonal Electricity Demand – Taylor

49. Open **Half-Hourly Multiple Seasonal Electricity Demand - Taylor.xlsx (Sheet 1 tab)**. This is half-hourly electricity demand (MW) in England and Wales from Monday, June 5, 2000 to Sunday, August 27, 2000 (taylor, R forecast). This data has multiple seasonality with frequency = 48 (observations per day) and 336 (observations per week), with a total of 4032 observations. Predictors such as temperature and work day are not available for this data. Reference: Taylor, J.W. (2003) Short-term electricity demand forecasting using double seasonal exponential smoothing. *Journal of the Operational Research Society*, **54**, 799-805.
50. Click **SigmaXL > Time Series Forecasting > Run Chart**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
51. Select *Demand*, click **Numeric Data Variable (Y) >>**. Select **Show Mean**. Uncheck **Nonparametric Runs Test**.

52. Click **OK**. A Run Chart of Electricity Demand with Mean center line is produced.



Here we see the strong half-hourly seasonal effect with a daily cycle of 48 observations and a weekly cycle of $48 \times 7 = 336$ observations. Since this data has multiple frequencies (and frequencies > 24), multiple seasonal decomposition (MSD) will be required to model it. Later, we will use a Spectral Density Plot to confirm the multiple seasonal frequency values of 48 and 336

Autocorrelation (ACF/PACF) Plots

Just as correlation measures the extent of a linear relationship between two variables, autocorrelation measures the linear relationship between lagged values of data. A plot of the data vs. the same data at lag k may show a positive or negative trend. If the slope is positive, the autocorrelation is positive; if there is a negative slope, the autocorrelation is negative. We will use the lag utility and scatterplots to demonstrate this for 3 lags using the Series A data.

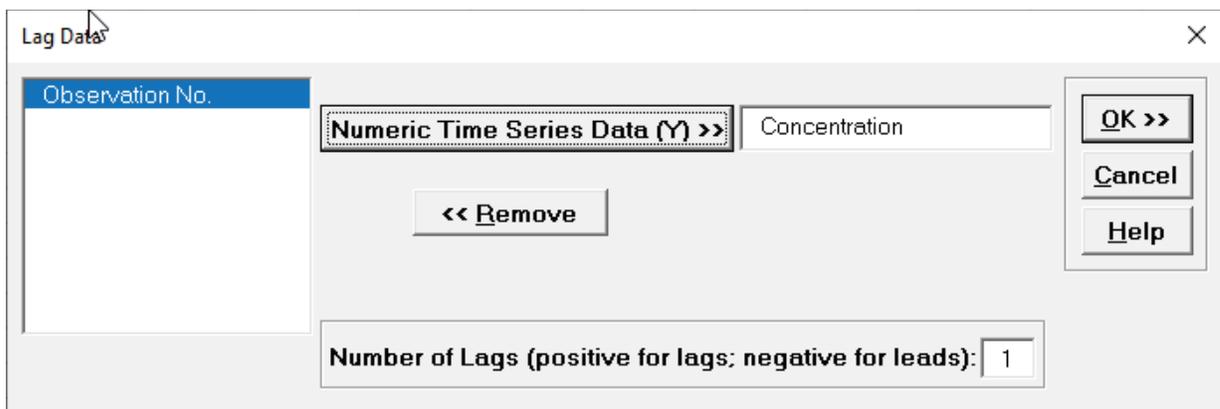
ACF denotes the AutoCorrelation Function plot (sometimes called a Correlogram). PACF denotes the Partial AutoCorrelation Function plot. For further details and formulas, see Appendix: [Autocorrelation \(ACF\), Partial Autocorrelation \(PACF\) and Cross Correlation \(CCF\)](#).

We will use the ACF and PACF to get a general idea of what models should be used, but let the automatic algorithms do the heavy lifting of determining which model is optimal. Later, ACF on model residuals are used to assist in determining how well a model fits the autocorrelation structure, ideally with the residuals showing no statistically significant autocorrelation.

We will also examine how differencing results in stationarity with the ACF and PACF plots.

Chemical Process Concentration - Series A

1. Open **Chemical Process Concentration – Series A.xlsx (Sheet 1 tab)**. This is the Series A data from Box and Jenkins, a set of 197 concentration values from a chemical process taken at two-hour intervals.
2. Click **SigmaXL > Time Series Forecasting > Utilities > Lag Data**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
3. Select *Concentration*, click **Numeric Time Series Data (Y) >>**. Use the default **Number of Lags = 1**.



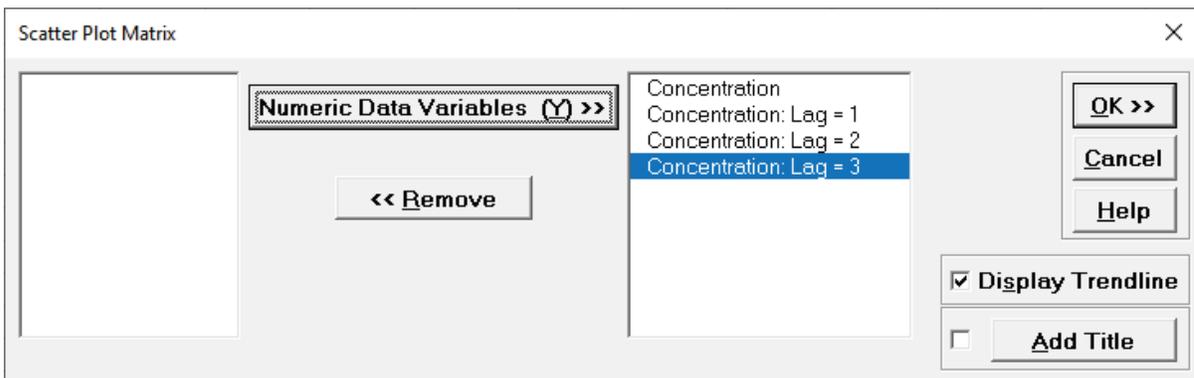
4. Click **OK**. A new sheet is created with the Lag 1 data.

Concentration	Concentration: Lag = 1
17	
16.6	17
16.3	16.6
16.1	16.3
17.1	16.1
16.9	17.1

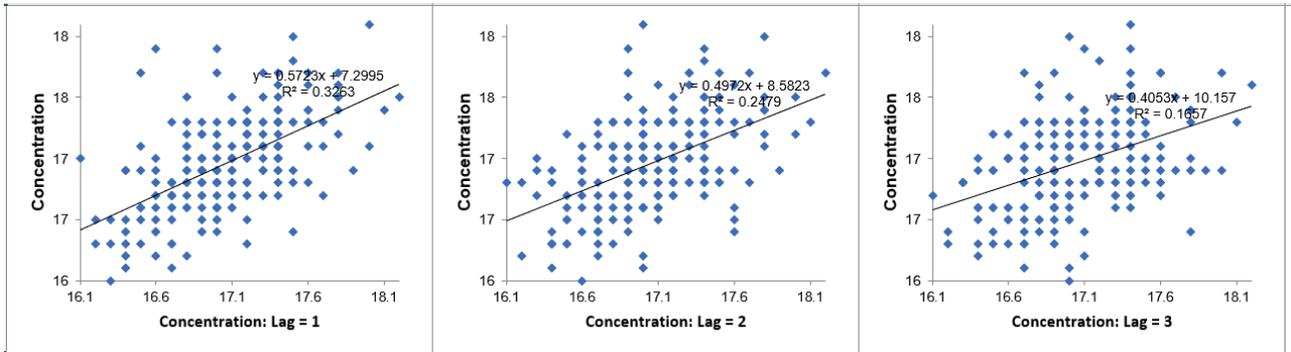
- Click **Recall SigmaXL Dialog** menu or press **F3** to recall last dialog. Enter **Lag = 2**. Click **OK**. Repeat for **Lag = 3**. The Lag columns are appended as shown:

Concentration	Concentration: Lag = 1	Concentration: Lag = 2	Concentration: Lag = 3
17			
16.6	17		
16.3	16.6	17	
16.1	16.3	16.6	17
17.1	16.1	16.3	16.6
16.9	17.1	16.1	16.3
16.8	16.9	17.1	16.1

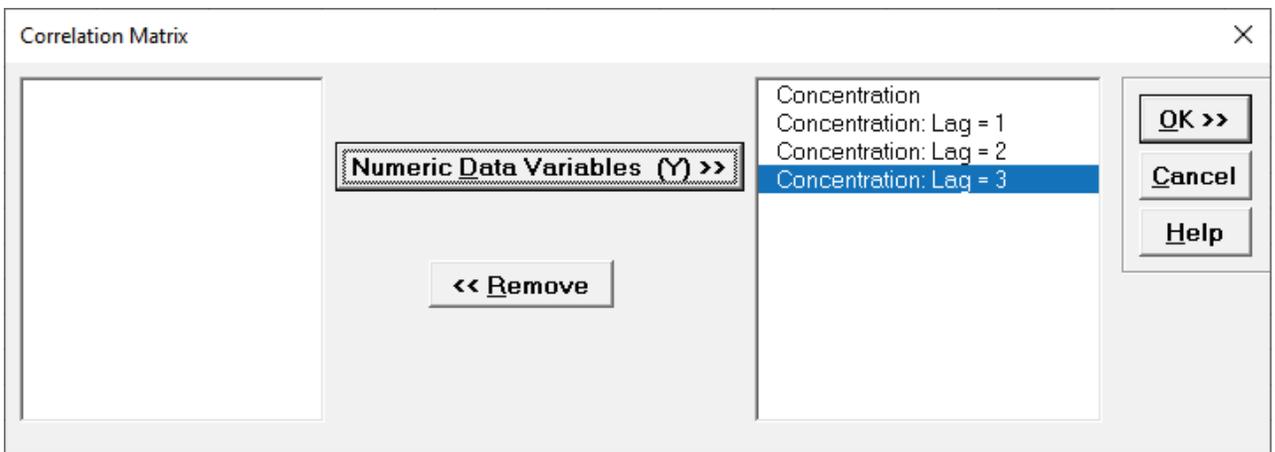
- Click **SigmaXL > Graphical Tools > Scatter Plot Matrix**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
- Select all four columns, click **Numeric Data Variables (Y) >>**. Check **Display Trendline**.



- Click **OK**. The Scatterplots for Lag 1 to 3 are produced, showing a positive trend, so positive autocorrelation.



- Click the **Lag Data** sheet tab. Click **SigmaXL > Statistical Tools > Correlation Matrix**. Ensure that the entire data table is selected. If not, check Use Entire Data Table. Click Next.
- Select all four columns, click **Numeric Data Variables (Y) >>**.



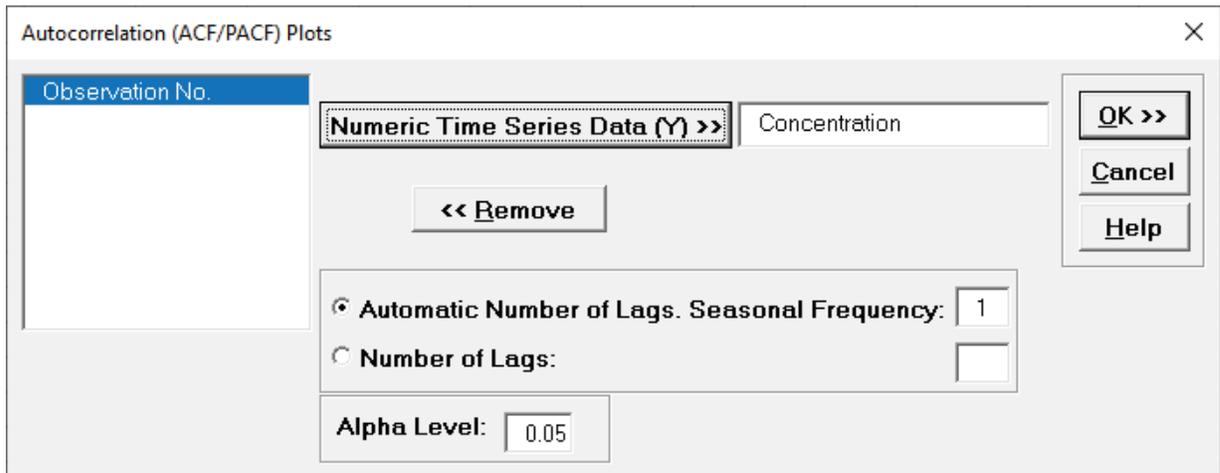
- Click **OK**. The Correlation Coefficient Matrix is produced.

Pearson Correlations	Concentration	Concentration: Lag = 1	Concentration: Lag = 2	Concentration: Lag = 3
Concentration	1	0.5713	0.4979	0.4070
Concentration: Lag = 1		1	0.5710	0.4952
Concentration: Lag = 2			1	0.5721
Concentration: Lag = 3				1

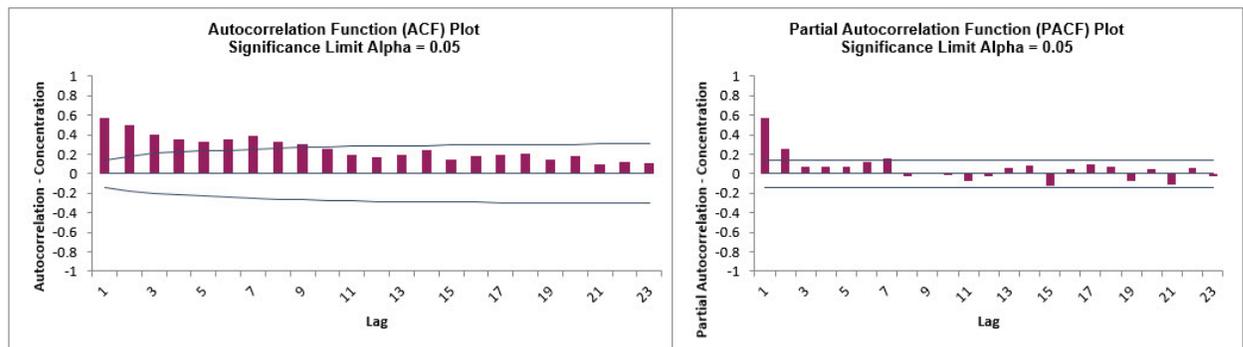
Lag 1 correlation is .57, Lag 2 is .5, and Lag 3 is .41. The correlations in red (and bolded P-Values) indicates that the correlations are significant.

- Now we will create the ACF/PACF plots. Click the **Sheet 1** tab. Click **SigmaXL > Time Series Forecasting > Autocorrelation (ACF/PACF) Plots**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.

13. Select *Concentration*, click **Numeric Time Series Data (Y) >>**. Use the default **Automatic Number of Lags**. **Seasonal Frequency = 1**. **Alpha Level = 0.05**.



14. Click **OK**. The ACF and PACF plots are produced.

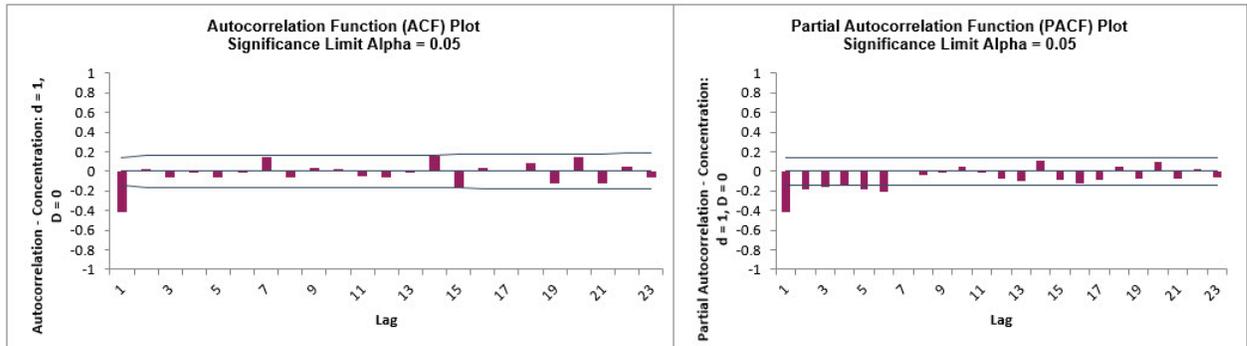


Hover the mouse cursor on the Lag 1 ACF bar to view the correlation value = .57, Lag 2 = 0.5, and Lag 3 = .4. The values are approximately the same as those obtained manually (with some minor differences in how the correlations are calculated).

The lags that exceed the significance limit lines are statistically significant at alpha = 0.05, so there is an autocorrelation structure in the data which can be modeled using Exponential Smoothing or ARIMA.

15. The slow decrease in the ACF as the lags increase is due to the non-stationarity, so we will now compare to the differenced data. Select the **Difference Data** sheet (or recreate using **SigmaXL > Time Series Forecasting > Utilities > Difference Data**, with **Nonseasonal Differencing (d) = 1**).

16. Redo the ACF/PACF Plots for the differenced data:



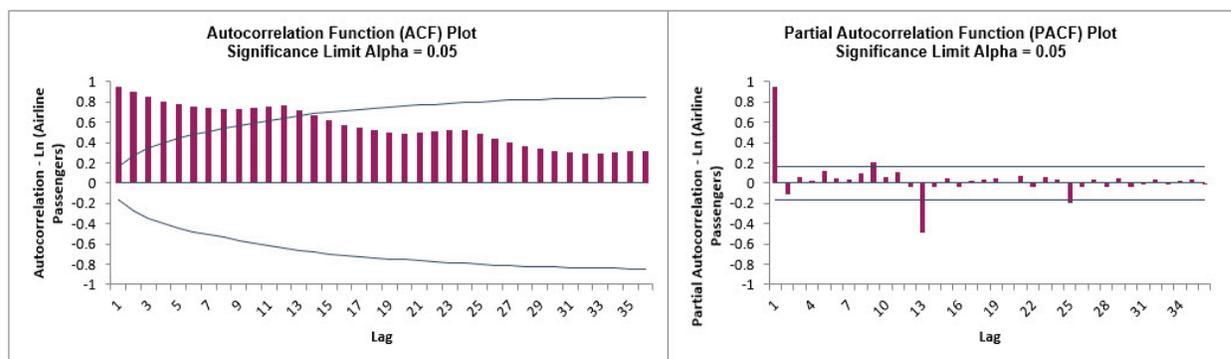
After differencing Lag 1 now shows a negative autocorrelation, but Lag 2 and following have insignificant autocorrelations and small partial autocorrelations. This is in agreement with the results in Box and Jenkins (2016, pp.185-186) who also suggest that, after differencing, the model might be a moving average of order 1. We will examine this later in ARIMA Forecasting.

Monthly Airline Passengers - Series G

- Open **Monthly Airline Passengers - Series G.xlsx (Sheet 1 tab)**. This is the Series G data from Box and Jenkins, monthly total international airline passengers for January 1949 to December 1960.
- Click **SigmaXL > Time Series Forecasting > Autocorrelation (ACF/PACF) Plots**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
- Select *Ln(Airline Passengers)*, click **Numeric Time Series Data (Y) >>**. Select **Automatic Number of Lags**. Specify **Seasonal Frequency = 12** and **Alpha Level = 0.05**.

The automatic number of lags will be (a minimum of) $\text{Seasonal Frequency} * 3 = 36$ allowing us to see seasonal patterns in the ACF.

- Click **OK**. The ACF and PACF plots are produced.

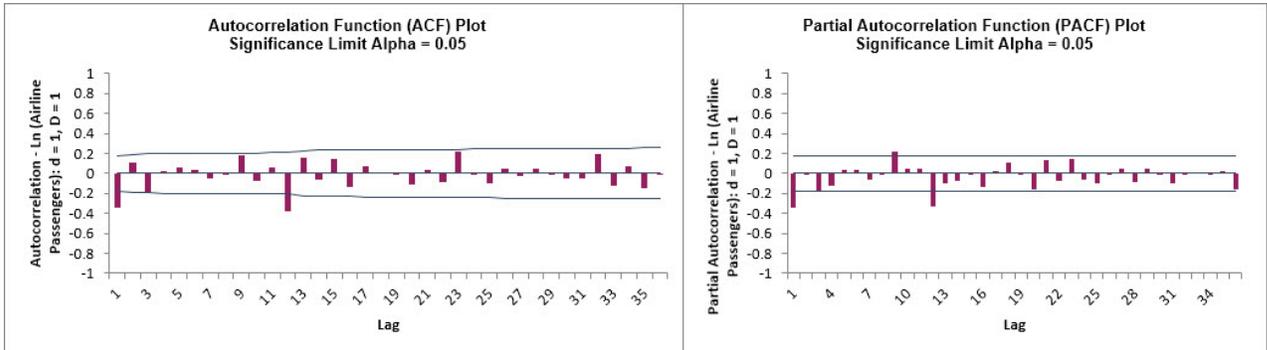


The slow decrease in the ACF as the lags increase is due to the trend, while the “scalloped” shape is due to the seasonality.

- We will now compare to the differenced data for *Ln(Airline Passengers)*. Select the **Difference Data** sheet (or recreate using **SigmaXL > Time Series Forecasting > Utilities > Difference Data**,

with **Nonseasonal Differencing (d) = 1** and **Seasonal Differencing (D) = 1** with **Seasonal Frequency = 12**).

22. Redo the ACF/PACF Plots for the differenced data:



After nonseasonal and seasonal differencing, Lag 1 now shows a negative autocorrelation, but Lag 2 and following have mostly insignificant autocorrelations and partial autocorrelations. Lag 12 is due to the seasonality. This is in agreement with the results in Box and Jenkins (2016, Chapter 9, “Analysis of Seasonal Time Series”, p. 319) who also suggest that, after differencing, the model might be a moving average of order 1 and seasonal moving average of order 1. We will examine this later in ARIMA Forecasting.

Daily Electricity Demand with Predictors – ElecDaily

23. Open **Daily Electricity Demand with Predictors – ElecDaily.xlsx (Sheet 1 tab)**. This is daily electricity demand (GW) for the state of Victoria, Australia, every day during 2014 (Hyndman, fpp2). This data has a seasonal frequency = 7 (observations per week).
24. Click **SigmaXL > Time Series Forecasting > Autocorrelation (ACF/PACF) Plots**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
25. Select *Demand*, click **Numeric Time Series Data (Y) >>**. Select **Automatic Number of Lags**. Specify **Seasonal Frequency = 7** and **Alpha Level = 0.05**.

Autocorrelation (ACF/PACF) Plots

Date

- Temp (C)
- TempSq
- WorkDay

Numeric Time Series Data (Y) >> Demand

<< Remove

Automatic Number of Lags. Seasonal Frequency: 7

Number of Lags:

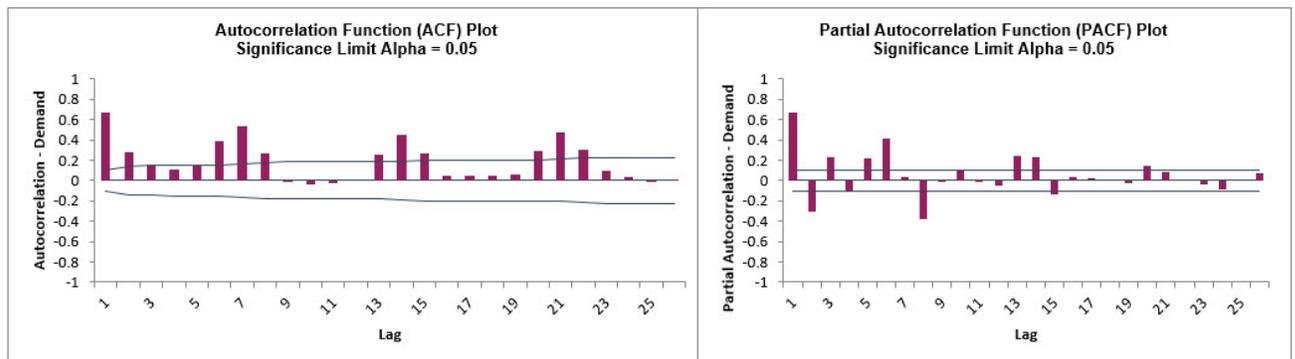
Alpha Level: 0.05

OK >>

Cancel

Help

26. Click **OK**. The ACF and PACF plots are produced.



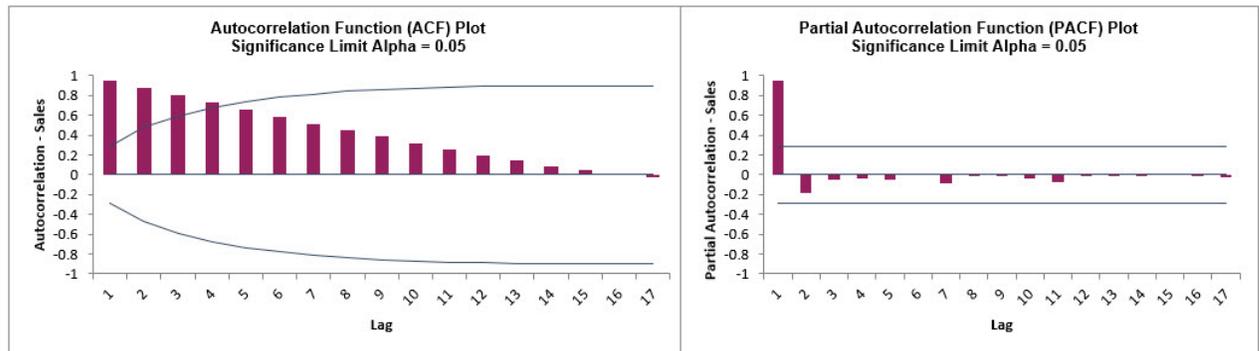
Here we can see the 7-day seasonal pattern. We will not do manual differencing for this example, but the automatic ARIMA algorithm will do a nonseasonal difference (d) = 1 and leave the seasonal differencing (D) = 0.

Sales with Indicator - Modified Series M

27. Open **Sales with Indicator - Modified Series M.xlsx**. (Sheet 1 tab). This is modified Series M data from Box and Jenkins, with 50 quarters of corporate sales values along with a leading indicator. The data is treated as nonseasonal, as done in Box and Jenkins.
28. Click **SigmaXL > Time Series Forecasting > Autocorrelation (ACF/PACF) Plots**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
29. Select *Sales*, click **Numeric Time Series Data (Y) >>**. Select **Automatic Number of Lags**. Specify **Seasonal Frequency = 1** and **Alpha Level = 0.05**.

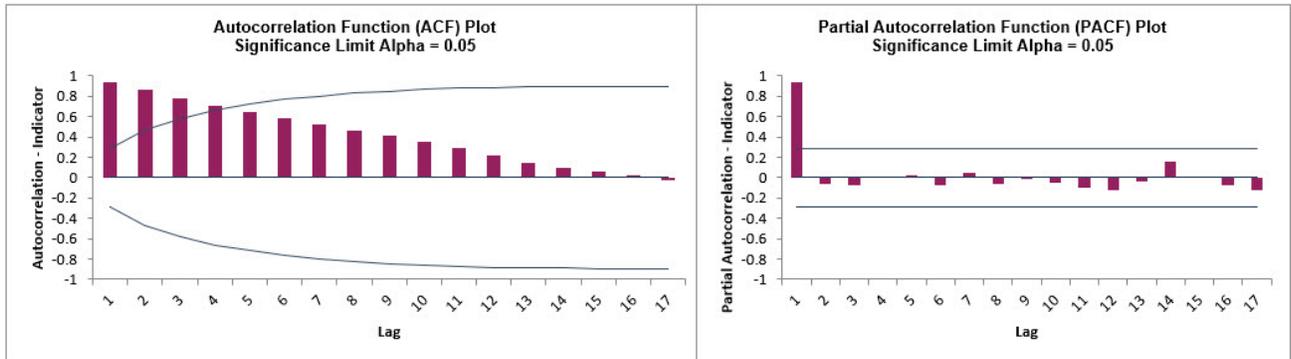


30. Click **OK**. The ACF and PACF plots are produced.



Here we can see the strong autocorrelation with slow decay due to the trend.

31. Click the **Sheet 1** tab. Redo the ACF/PACF Plots for Indicator:



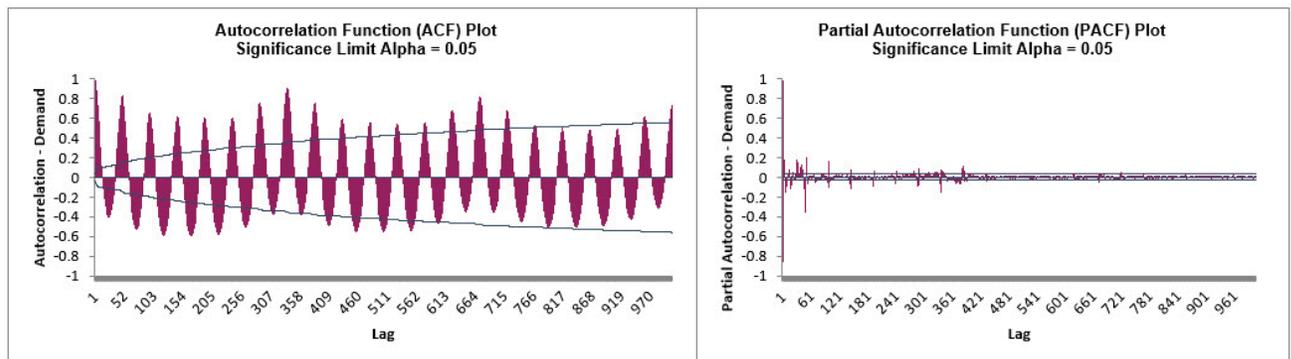
As expected, the autocorrelation for Indicator is very similar to Sales. This will be addressed when Cross Correlation (CCF) Plots are created.

Half-Hourly Multiple Seasonal Electricity Demand – Taylor

32. Open **Half-Hourly Multiple Seasonal Electricity Demand - Taylor.xlsx (Sheet 1 tab)**. This is half-hourly electricity demand (MW) in England and Wales from Monday, June 5, 2000 to Sunday, August 27, 2000 (taylor, R forecast). This data has multiple seasonality with frequency = 48 (observations per day) and 336 (observations per week), with a total of 4032 observations.
33. Click **SigmaXL > Time Series Forecasting > Autocorrelation (ACF/PACF) Plots**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
34. Select *Demand*, click **Numeric Time Series Data (Y) >>**. Select **Automatic Number of Lags**. Specify **Seasonal Frequency** = 336 and **Alpha Level** = 0.05.

The automatic number of lags will be (a minimum of) $\text{Seasonal Frequency} * 3 = 1008$ allowing us to see seasonal patterns in the ACF.

35. Click **OK**. The ACF and PACF plots are produced.



Here we can see the half-hourly seasonal patterns of 48 per day and 336 per week.

Cross Correlation (CCF) Plots

Cross Correlation is similar to autocorrelation, but the correlations are computed on two related time series variables, typically a process input and output. A plot of the X data vs. the Y data at lag k may show a positive or negative trend. If the slope is positive, the cross correlation is positive; if there is a negative slope, the cross correlation is negative. This helps to identify important lags (or leads) in the process and is useful for application when there are predictors in an ARIMA model.

If the X input is autocorrelated, the CCF is affected by its time series structure and any “in common” trends the X and Y series may have over time. Pre-whitening solves this problem by removing the autocorrelation and trends.

For further details and formulas, see Appendix: [Autocorrelation \(ACF\), Partial Autocorrelation \(PACF\) and Cross Correlation \(CCF\)](#).

Sales with Indicator - Modified Series M

1. Open **Sales with Indicator - Modified Series M.xlsx**. (Sheet 1 tab). This is modified Series M data from Box and Jenkins, with 50 quarters of corporate sales values along with a leading indicator. The data is treated as nonseasonal, as done in Box and Jenkins.
2. Click **SigmaXL > Time Series Forecasting > Cross Correlation (CCF) Plot**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
3. Select *Indicator*, click **Input Time Series (X) >>**. Select *Sales*, click **Output Time Series (Y) >>**. Use the default **Automatic Number of Lags**. **Seasonal Frequency = 1** and **Alpha Level = 0.05**.

Cross Correlation (CCF) Plot

Qtr-Year

Input Time Series (X) >> Indicator

Output Time Series (Y) >> Sales

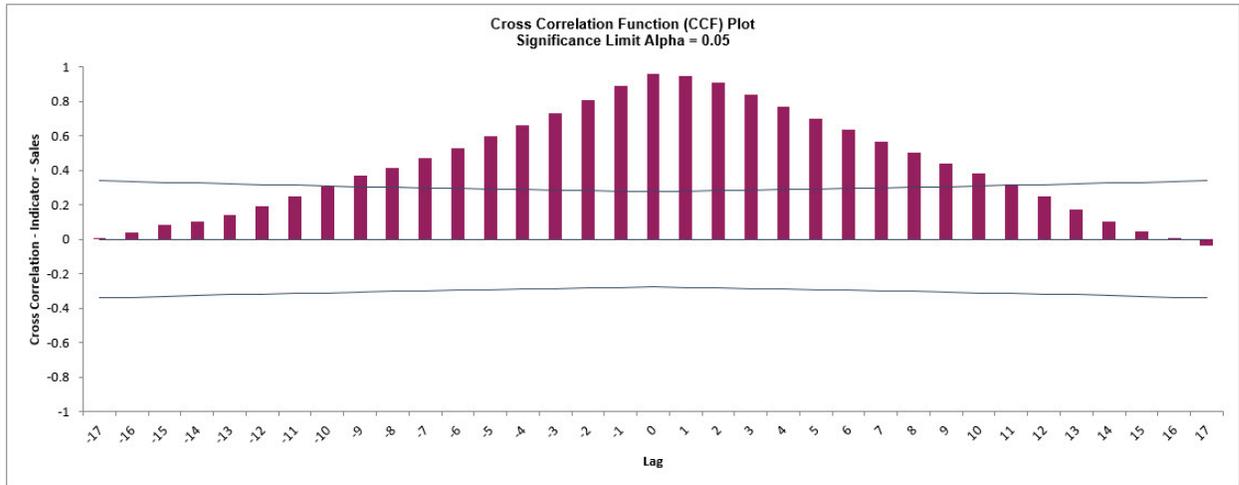
<< Remove

Automatic Number of Lags. Seasonal Frequency: 1
 Number of Lags:

Alpha Level: 0.05 Pre-Whiten Data

OK >> Cancel Help

- Click **OK**. The CCF plot is produced.

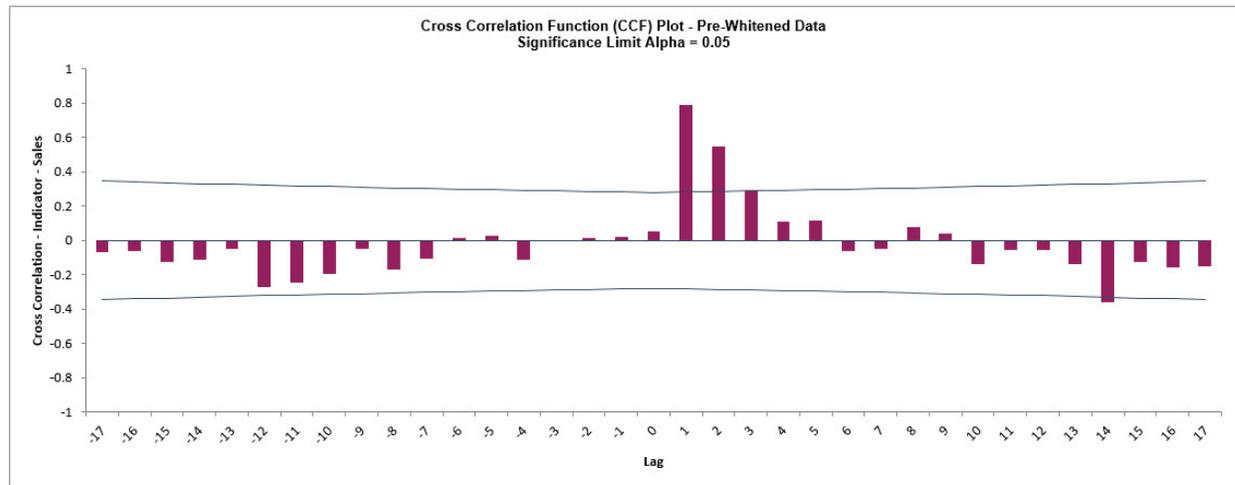


This CCF plot shows significant cross correlation from lag = -9 to +9, with a peak at 0, however the autocorrelation in X and Y data is masking the true nature of the cross correlation.

- We will redo the CCF Plot with the Pre-Whiten Data option. Click **Recall SigmaXL Dialog** menu or press **F3** to recall last dialog. Check **Pre-Whiten Data**.

The figure shows a dialog box titled "Cross Correlation (CCF) Plot" with a close button (X) in the top right corner. On the left side, there is a list box containing "Qtr-Year". To the right of the list box are two buttons: "Input Time Series (X) >>" and "Output Time Series (Y) >>". Below these buttons are two text input fields: "Indicator" and "Sales". Below the input fields is a button labeled "<< Remove". On the right side of the dialog, there are three buttons: "OK >>", "Cancel", and "Help". At the bottom of the dialog, there are two radio button options: "Automatic Number of Lags. Seasonal Frequency: 1" (which is selected) and "Number of Lags:". Below these options is a text input field for "Alpha Level:" containing the value "0.05". To the right of the Alpha Level field is a checked checkbox labeled "Pre-Whiten Data".

6. Click **OK**. The CCF plot is produced.



Pre-Whitening the data has dramatically altered the CCF plot, allowing us to see the underlying cross correlation pattern. Lags 1 and 2 are significantly positive, and Lag 3 is just on the significance line.

Note that while X is called a leading indicator, i.e., X comes before Y in time, the positive lag means that the X variable is lagging the Y variable in terms of correlation structure. SigmaXL uses this convention as given in Box and Jenkins (2016, pp. 437-440).

This CCF plot will be useful later when we model the Sales data using ARIMA Forecast with Predictors.

Spectral Density Plot

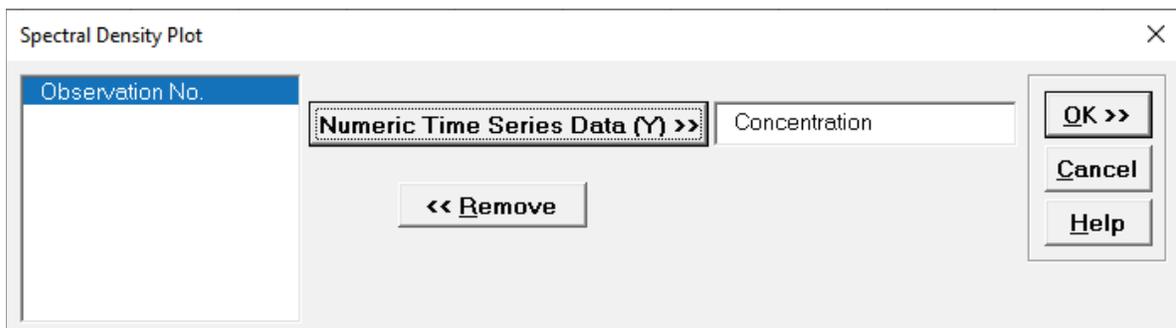
The Spectral Density Plot is used to identify the dominant integer seasonal frequency in time series data using spectral analysis with fast Fourier transforms. The algorithm used here is the same as used in the forecast model option to automatically detect seasonal frequency. If there is multiple seasonality, up to three integer frequencies will be identified. If the peak frequency is not an integer, it is rounded.

The Y axis is Spectral Density, the X axis is Seasonal Frequency. The Spectral Density Plot is also known as a Periodogram. Note that SigmaXL's use of the term "seasonal frequency" is the inverse of what is typically used in Fourier transforms "seasonal period", as discussed in the Introduction.

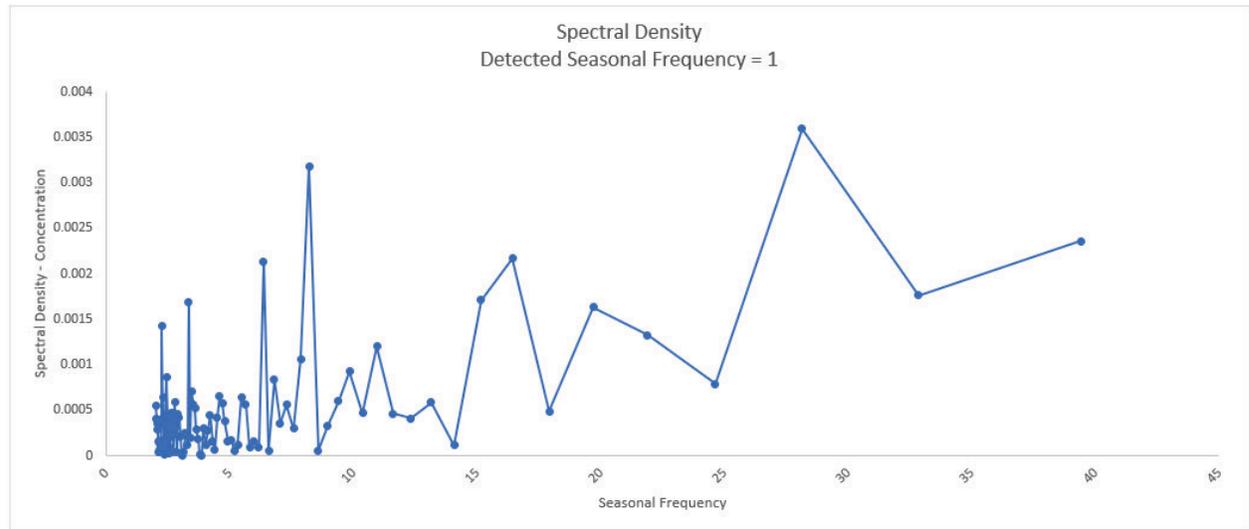
For further details and formulas, see Appendix: [Spectral Density and Automatic Detection of Seasonal Frequency](#).

Chemical Process Concentration - Series A

1. Open **Chemical Process Concentration – Series A.xlsx (Sheet 1 tab)**. This is the Series A data from Box and Jenkins, a set of 197 concentration values from a chemical process taken at two-hour intervals.
2. Click **SigmaXL > Time Series Forecasting > Spectral Density Plot**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
3. Select *Concentration*, click **Numeric Time Series Data (Y) >>**.



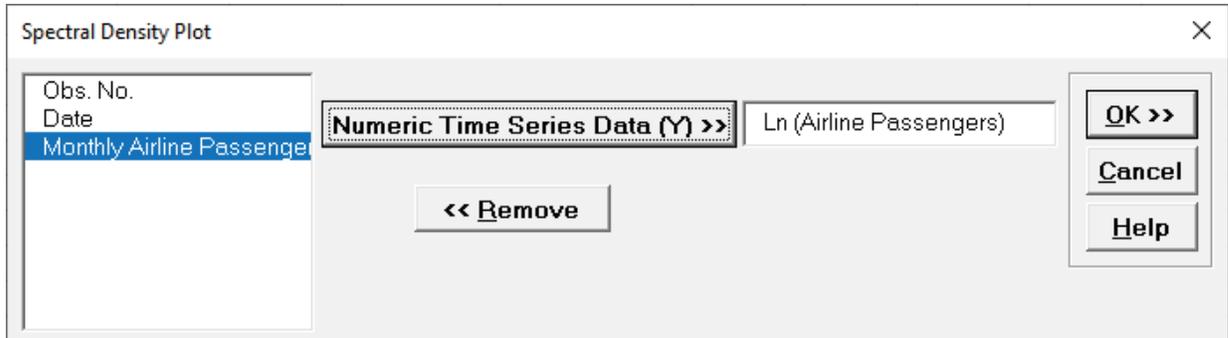
4. Click **OK**. A Spectral Density Plot for Concentration is produced.



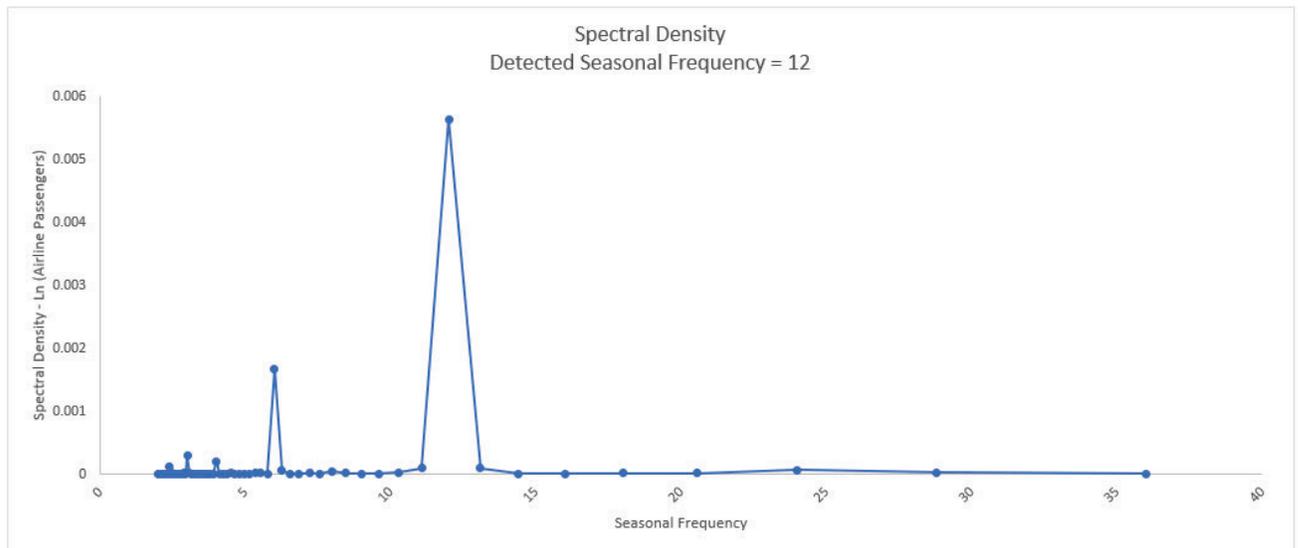
The detected seasonal frequency is 1, which means that it is a nonseasonal process. The peak at 28 does not have enough seasonal “strength” to be considered for use as seasonal frequency in a time series model.

Monthly Airline Passengers - Series G

5. Open **Monthly Airline Passengers - Series G.xlsx (Sheet 1 tab)**. This is the Series G data from Box and Jenkins, monthly total international airline passengers for January 1949 to December 1960.
6. Click **SigmaXL > Time Series Forecasting > Spectral Density Plot**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
7. Select *Ln(Airline Passengers)*, click **Numeric Time Series Data (Y) >>**.



8. Click **OK**. A Spectral Density Plot for Ln(Airline Passengers) is produced.



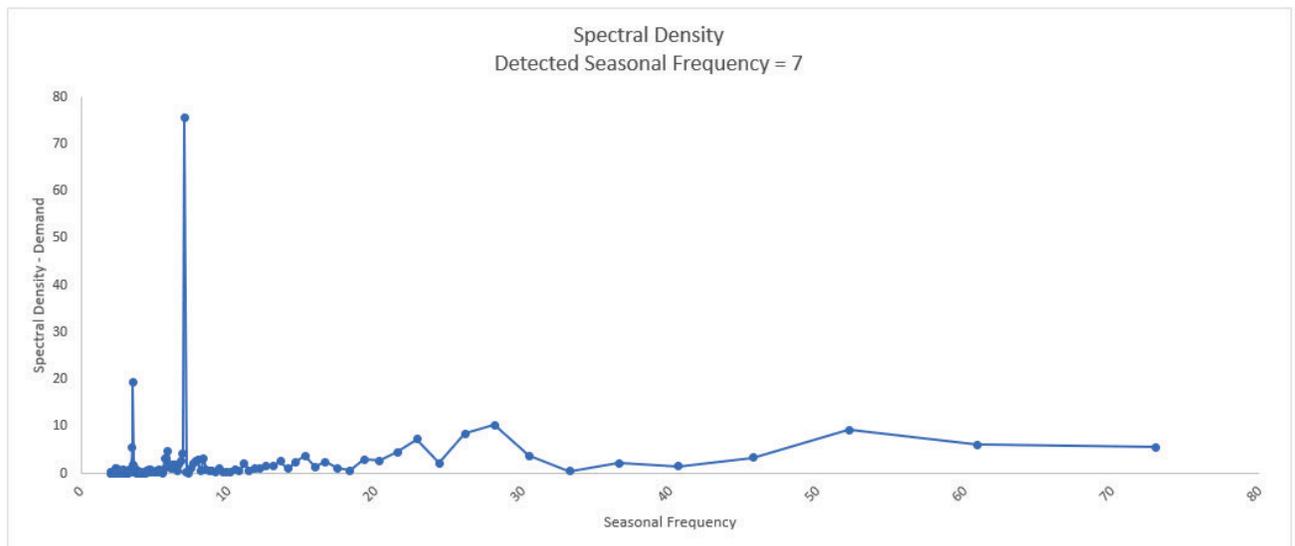
As expected, the detected seasonal frequency for the monthly data is 12. The peak at 6 does not have enough seasonal “strength” to be considered as a second seasonal frequency.

Daily Electricity Demand with Predictors – ElecDaily

9. Open **Daily Electricity Demand with Predictors – ElecDaily.xlsx (Sheet 1 tab)**. This is daily electricity demand (GW) for the state of Victoria, Australia, every day during 2014 (Hyndman, fpp2). This data has a seasonal frequency = 7 (observations per week).
10. Click **SigmaXL > Time Series Forecasting > Spectral Density Plot**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
11. Select *Demand*, click **Numeric Time Series Data (Y) >>**.



12. Click **OK**. A Spectral Density Plot for Demand is produced.



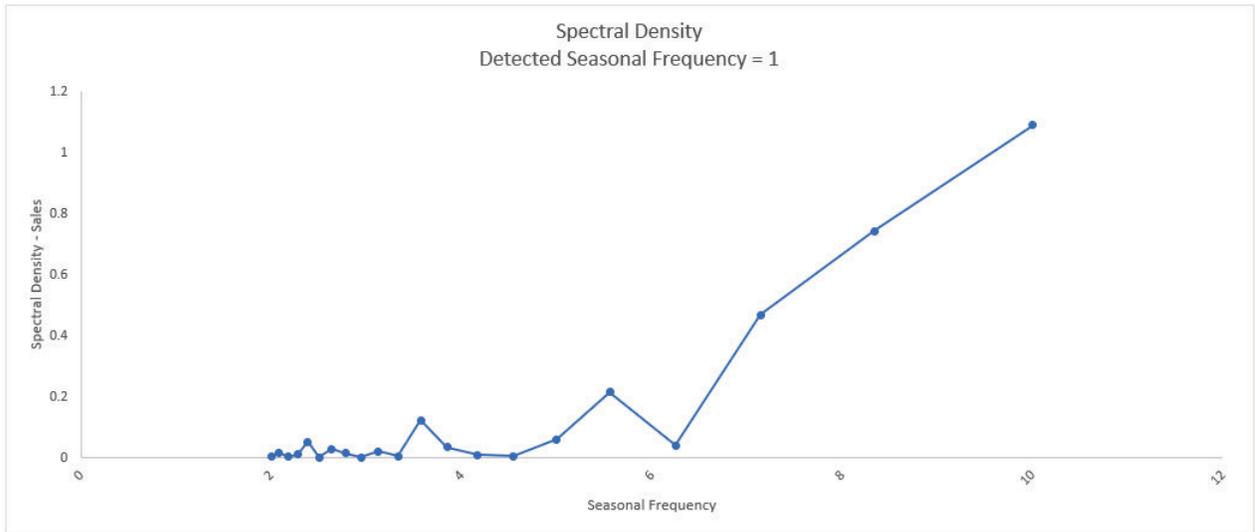
As expected, the detected seasonal frequency for the daily data is 7.

Sales with Indicator - Modified Series M

13. Open **Sales with Indicator - Modified Series M.xlsx**. (Sheet 1 tab). This is modified Series M data from Box and Jenkins, with 50 quarters of corporate sales values along with a leading indicator. The data is treated as nonseasonal, as done in Box and Jenkins.
14. Click **SigmaXL > Time Series Forecasting > Spectral Density Plot**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
15. Select *Sales*, click **Numeric Time Series Data (Y) >>**.



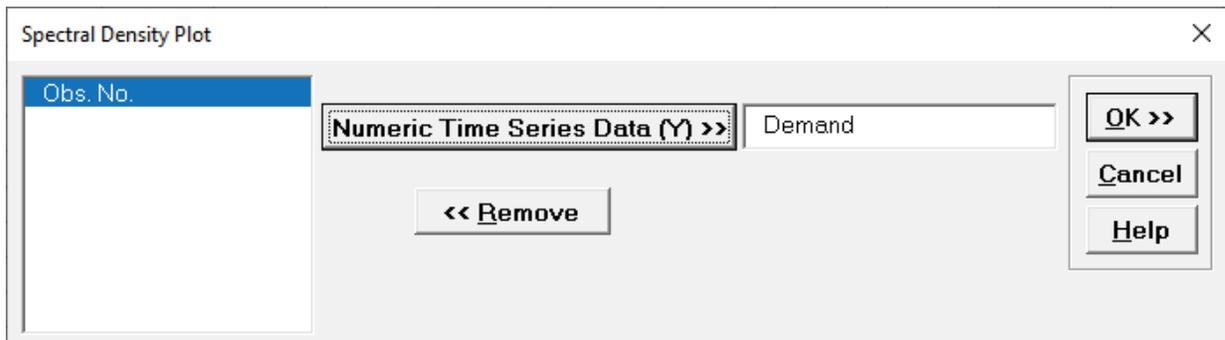
16. Click **OK**. A Spectral Density Plot for Sales is produced.



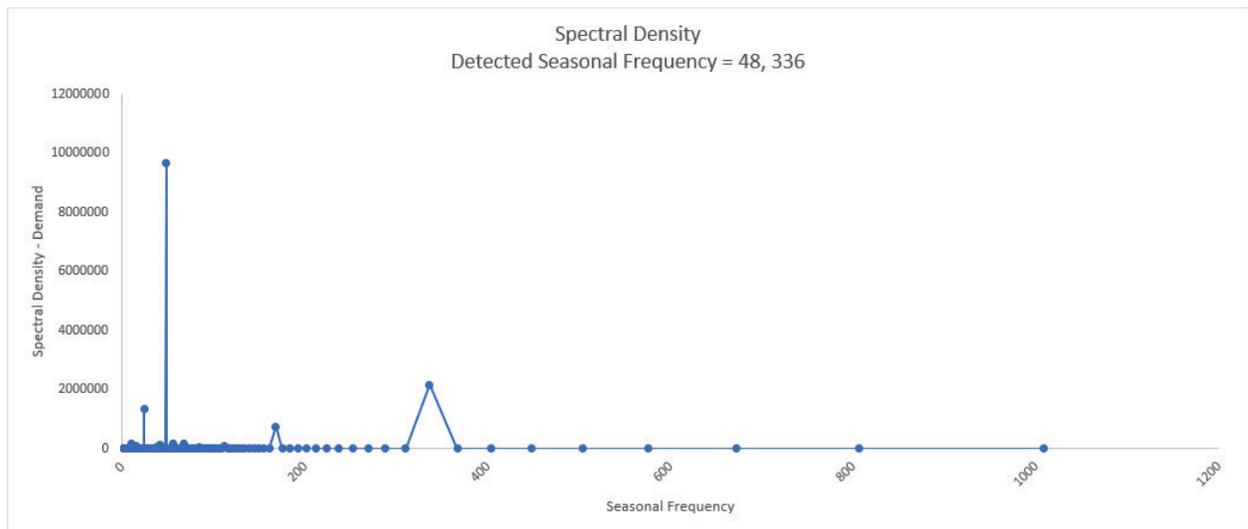
The detected seasonal frequency is 1, which means that it is confirmed to be a nonseasonal process. The peak at 10 does not have enough seasonal “strength” to be considered for use as seasonal frequency in a time series model.

Half-Hourly Multiple Seasonal Electricity Demand – Taylor

17. Open **Half-Hourly Multiple Seasonal Electricity Demand - Taylor.xlsx (Sheet 1 tab)**. This is half-hourly electricity demand (MW) in England and Wales from Monday, June 5, 2000 to Sunday, August 27, 2000 (taylor, R forecast). This data has multiple seasonality with frequency = 48 (observations per day) and 336 (observations per week), with a total of 4032 observations.
18. Click **SigmaXL > Time Series Forecasting > Spectral Density Plot**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
19. Select *Demand*, click **Numeric Time Series Data (Y) >>**.



20. Click **OK**. A Spectral Density Plot for Demand is produced.



The detected multiple seasonal frequency is confirmed as 48, 336.

Seasonal Trend Decomposition Plots

The Seasonal Trend Decomposition Plots are useful to visually distinguish trend and seasonal components in the time series data. If the Seasonal Frequency is unchecked, a Trend Decomposition Plot is produced as the first chart, showing the raw data and the trend. The trend component uses data smoothing, rather than a linear trend so that it may display either a linear trend or cyclical patterns. If a single seasonal frequency is specified, a Seasonal Trend Decomposition Plot is produced, showing the data, smoothed trend and seasonal component. If a multiple seasonal frequency is specified, a Multiple Seasonal Trend Decomposition Plot is produced, showing the data, smoothed trend and multiple seasonal component.

The second chart shows just the smoothed trend; the third chart (if applicable) shows just the seasonal or multiple seasonal component. The final chart is the remainder component.

This is an additive decomposition model, so the sum of the trend value + seasonal value(s) + remainder value gives the original data value. A multiplicative equivalent may be obtained by specifying the Box-Cox Transformation with Lambda = 0, which is a Ln transformation, but the charts will display the transformed data to maintain an additive model. Rounded or Optimal Lambda may also be used, but will only consider the range of values 0 to 1 (this conservative approach is used in time series forecasting, unlike regular Box-Cox in SigmaXL which uses a range of -5 to 5). See Appendix: [Box-Cox Transformation](#).

The decomposition algorithms used here are the same as used in Exponential Smoothing – Multiple Seasonal Decomposition (MSD), and ARIMA – MSD. The seasonal component is first removed through decomposition, a nonseasonal exponential smooth model fitted to the remainder + trend, and then the seasonal component is added back in. This is mainly used for high seasonal frequency and/or multiple seasonal frequency time series. For further details and formulas, see Appendix: [Seasonal Trend Decomposition](#).

Chemical Process Concentration – Series A

1. Open **Chemical Process Concentration – Series A.xlsx (Sheet 1 tab)**. This is the Series A data from Box and Jenkins, a set of 197 concentration values from a chemical process taken at two-hour intervals.
2. Click **SigmaXL > Time Series Forecasting > Seasonal Trend Decomposition Plots**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
3. Select *Concentration*, click **Numeric Time Series Data (Y) >>**. **Seasonal Frequency** and **Box-Cox Transformation** should be unchecked as shown.

Seasonal Trend Decomposition Plots

Observation No.

Numeric Time Series Data (Y) >> Concentration

<< Remove

OK >>
Cancel
Help

Seasonal Frequency

Box-Cox Transformation

Specify 12

Select 4-Quarterly

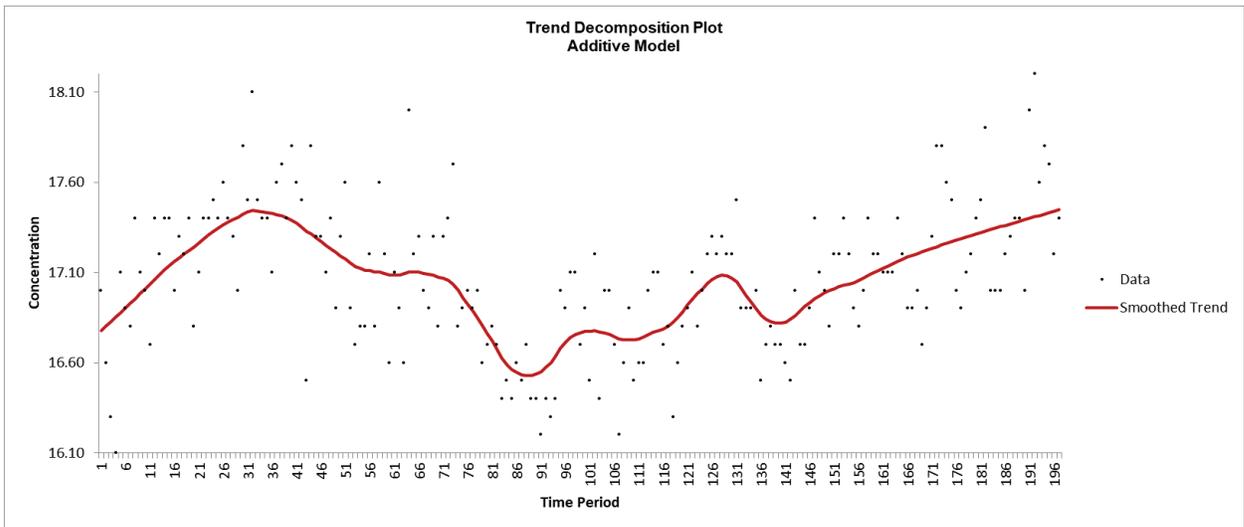
Automatically Detect

Rounded Lambda

Optimal Lambda

Lambda & Threshold (Shift)

- Click **OK**. A Trend Decomposition Plot for Concentration is produced.



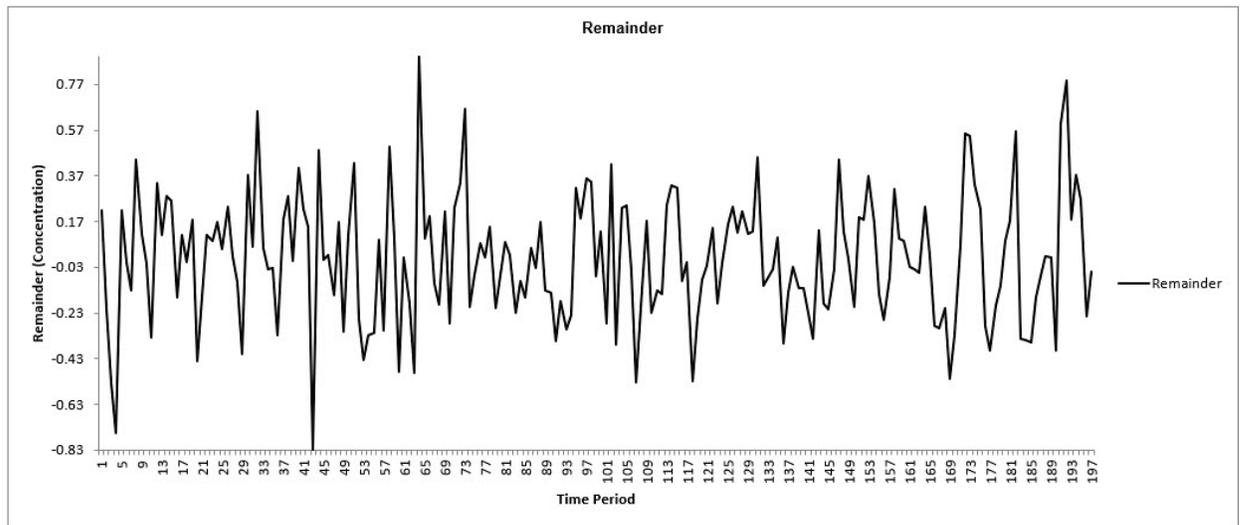
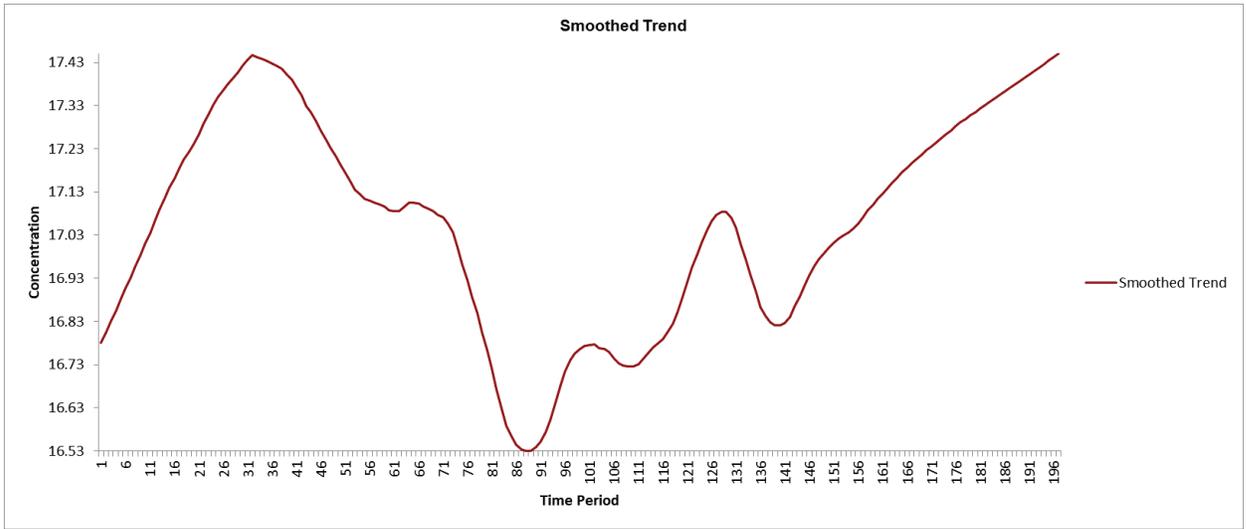
We can clearly see the “wandering mean” in this process. As discussed earlier, this can be modeled with exponential smoothing or with ARIMA after differencing.

- A Decomposition Summary report is also included to the right of the plot, indicating seasonal frequency and Box-Cox parameters (if applicable):

Decomposition Summary	
Seasonal Frequency	1
Box-Cox Transformation	N/A
Lambda	
Threshold	

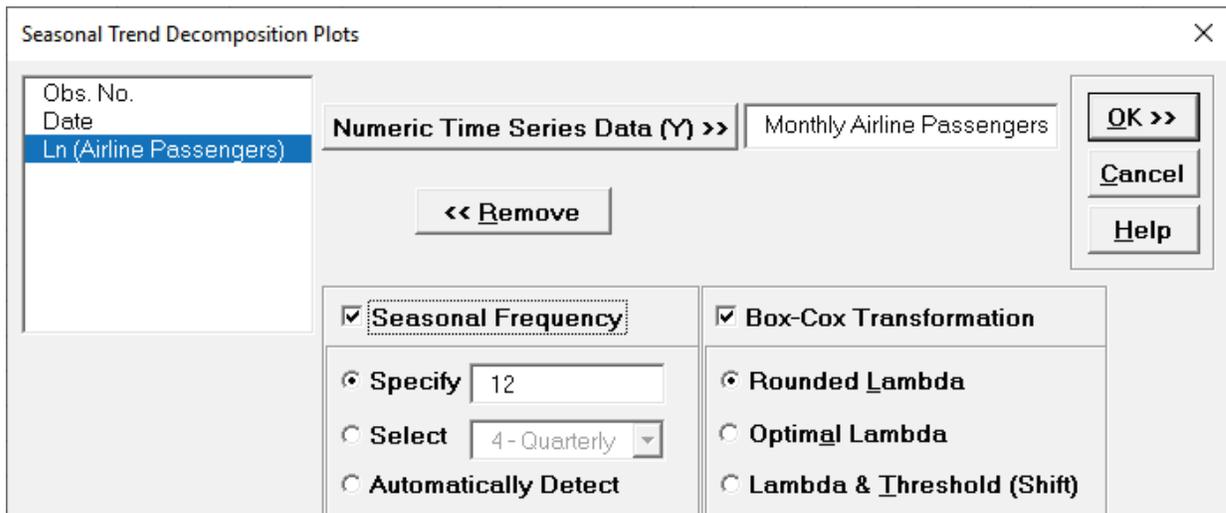
Seasonal Frequency = 1 denotes a nonseasonal process.

6. The Smoothed Trend and Remainder Plots are also produced as shown:

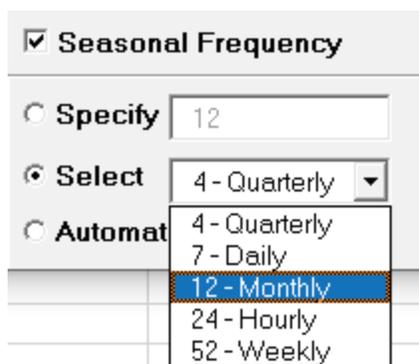


Monthly Airline Passengers - Series G

- Open **Monthly Airline Passengers - Series G.xlsx (Sheet 1 tab)**. This is the Series G data from Box and Jenkins, monthly total international airline passengers for January 1949 to December 1960.
- Click **SigmaXL > Time Series Forecasting > Seasonal Trend Decomposition Plots**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
- Select *Monthly Airline Passengers*, click **Numeric Time Series Data (Y) >>**. Check **Seasonal Frequency** with **Specify** = 12. Check **Box-Cox Transformation** with **Rounded Lambda** (selected because the [Run Chart](#) showed an increase in the seasonal variance over time).



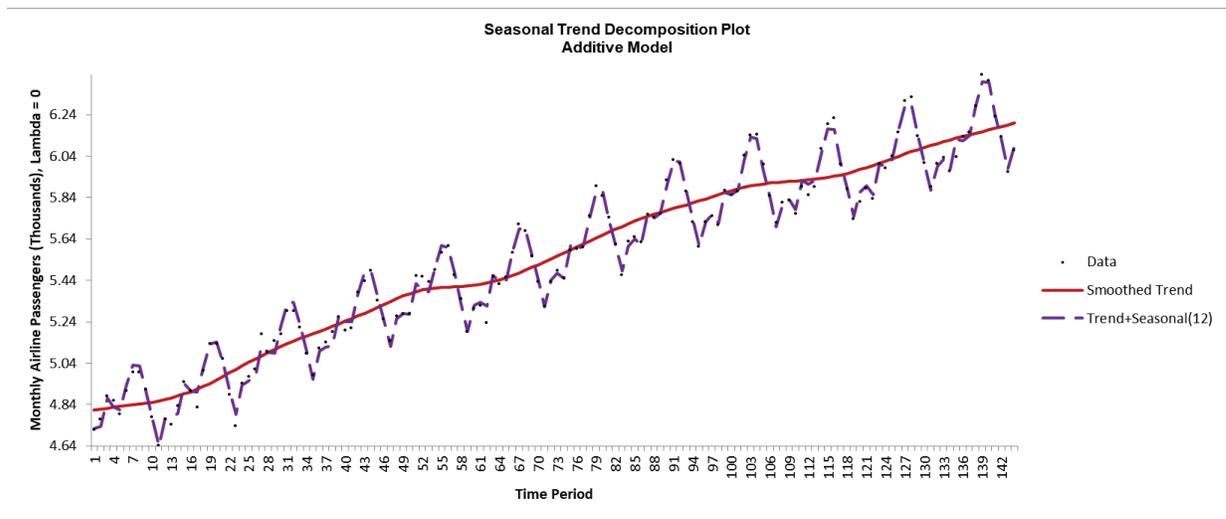
- **Seasonal Frequency Specify** also permits the entry of multiple frequencies
- **Seasonal Frequency Select** gives a drop-down list of commonly used seasonal frequencies:



- **Seasonal Frequency Automatically Detect** should be used if uncertain what the seasonal frequency value is (or do a Spectral Density Plot prior to the Seasonal Trend Decomposition Plots).

- **Box-Cox Transformation with Rounded Lambda** will select Lambda = 0 (Ln), 0.5 (SQRT) or 1 (Untransformed). Threshold (Shift) is computed automatically if the time series data includes 0 or negative values, otherwise it is 0.
- **Box-Cox Transformation with Optimal Lambda** uses the range of 0 to 1 for Lambda. Threshold is computed automatically if the time series data includes 0 or negative values.
- **Box-Cox Transformation with Lambda & Threshold (Shift)** if left blank, will compute optimal lambda and threshold. The user may also specify Lambda and Threshold. Lambda may be specified outside of the 0 to 1 range, but practically for time series analysis, should be limited to -1 to 2. Threshold is typically 0, but if the time series data includes 0 or negative values, a negative threshold value should be entered that is smaller than the minimum data value. This value will be subtracted from the data resulting in positive time series values.

10. Click **OK**. A Seasonal Trend Decomposition Plot for Monthly Airline Passengers is produced.

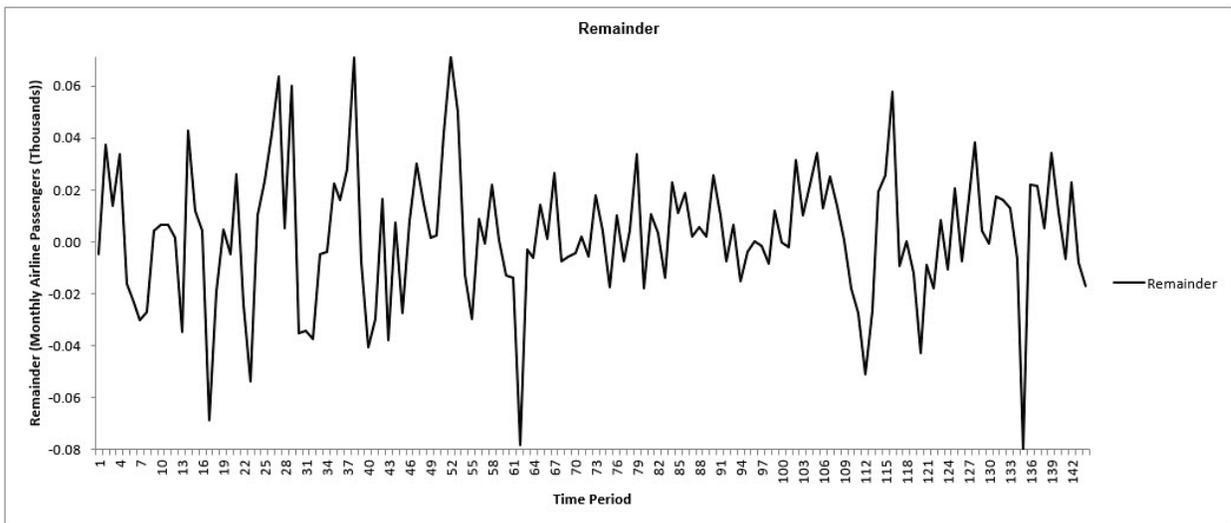
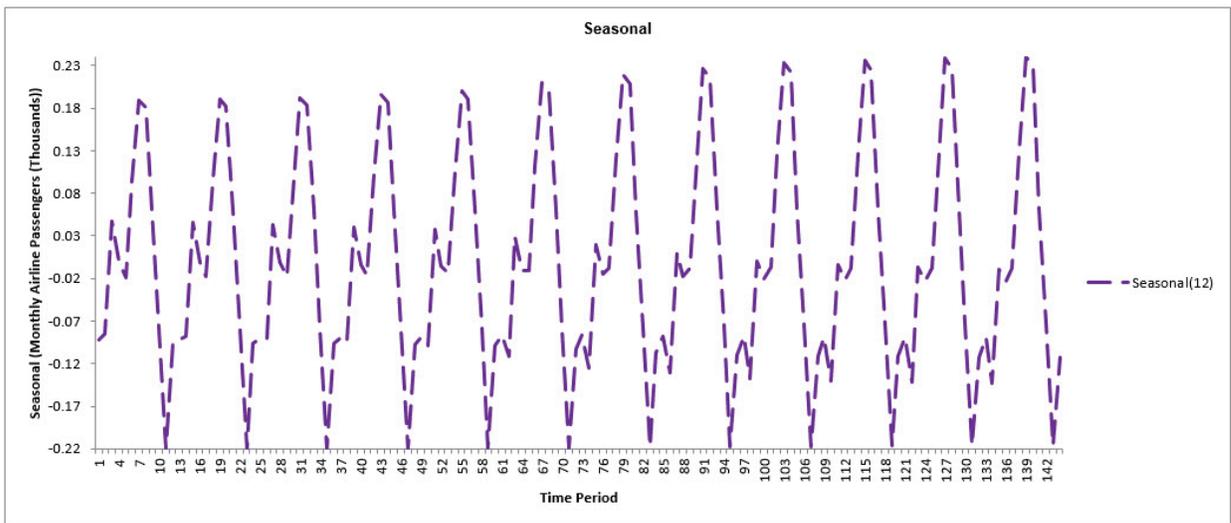
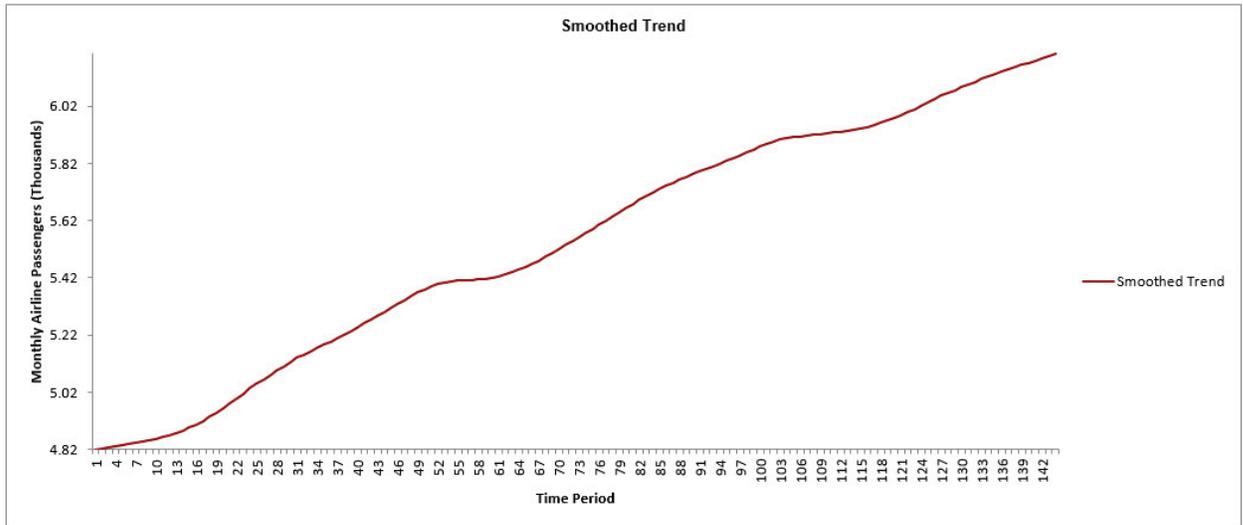


Here we see a strong positive trend as well as the monthly seasonal effect. Note that this is the Lambda = 0 (Ln transformed) data. The Box-Cox transformation information is given in the Decomposition Summary report to the right of the plot:

Decomposition Summary	
Seasonal Frequency	12
Box-Cox Transformation	Rounded Lambda
Lambda	0
Threshold	0

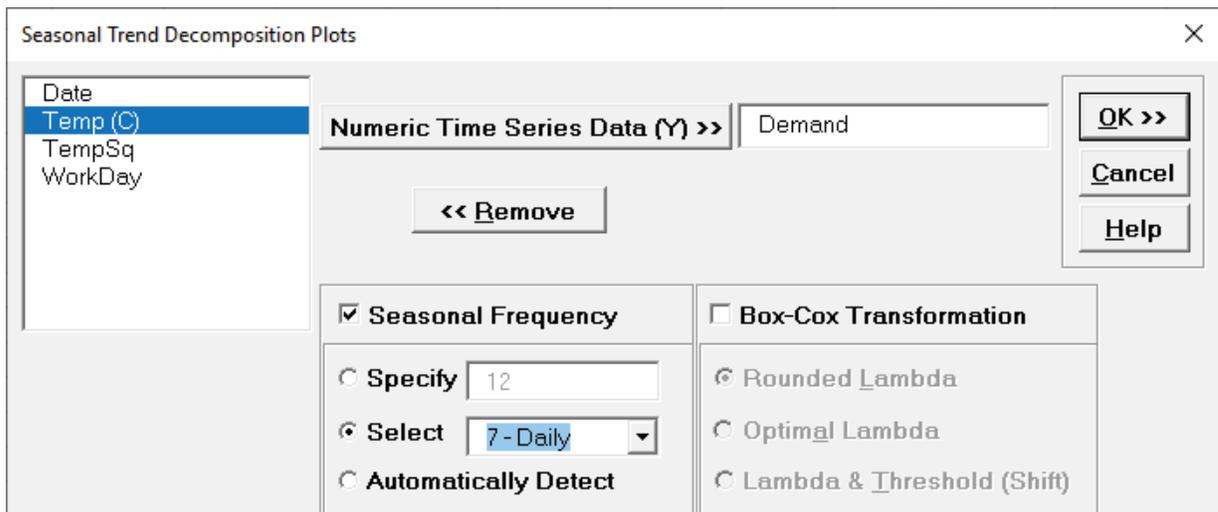
Seasonal Frequency is 12 and Lambda = 0. The Ln transformed data are displayed to maintain an additive model, which is easier to interpret than a multiplicative model.

11. The Smoothed Trend, Seasonal and Remainder Plots are also produced as shown:

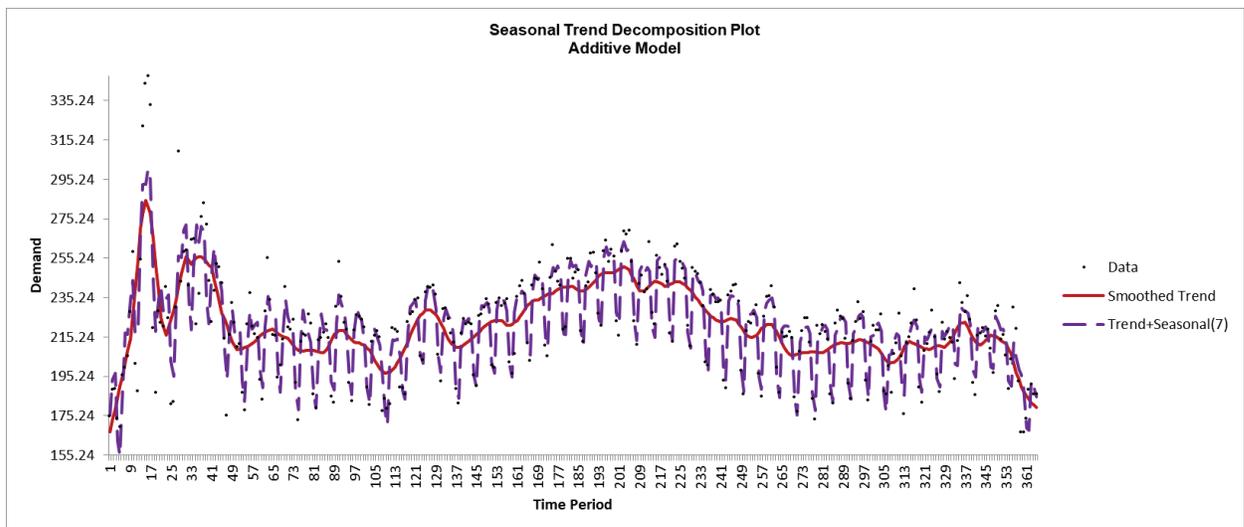


Daily Electricity Demand with Predictors – ElecDaily

12. Open **Daily Electricity Demand with Predictors – ElecDaily.xlsx (Sheet 1 tab)**. This is daily electricity demand (GW) for the state of Victoria, Australia, every day during 2014 (Hyndman, fpp2). This data has a seasonal frequency = 7 (observations per week).
13. Click **SigmaXL > Time Series Forecasting > Seasonal Trend Decomposition Plots**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
14. Select *Demand*, click **Numeric Time Series Data (Y) >>**. Check **Seasonal Frequency** with **Select = 7-Daily** selected from the drop-down list. Uncheck **Box-Cox Transformation**.

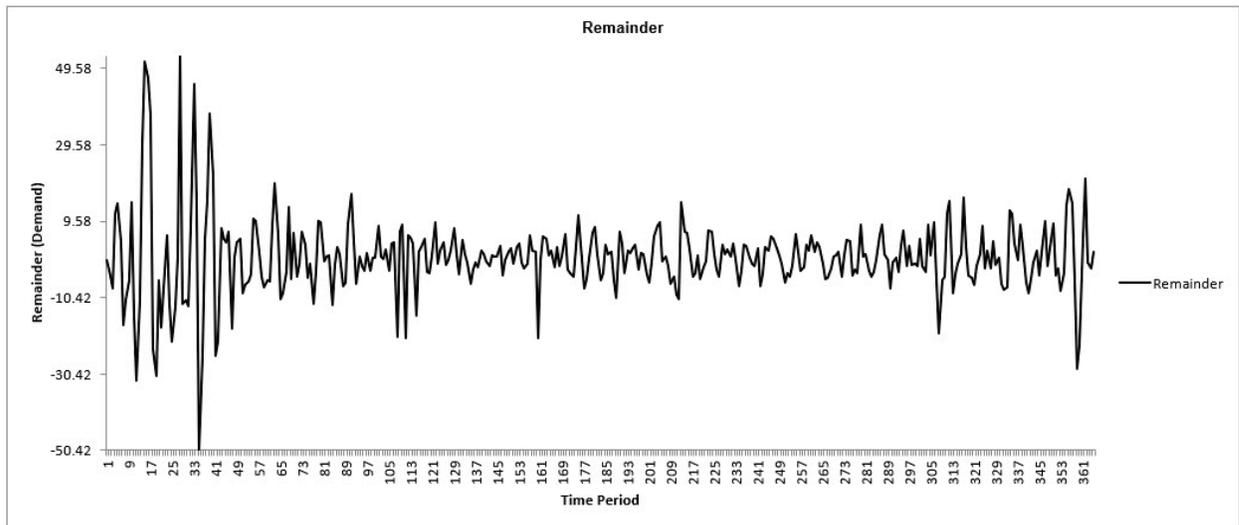
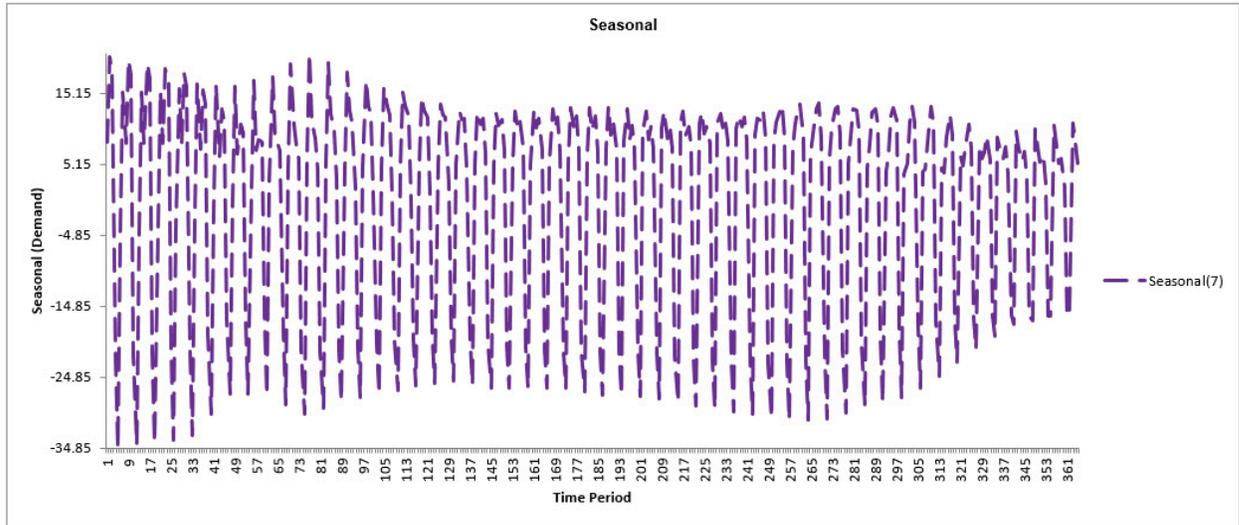
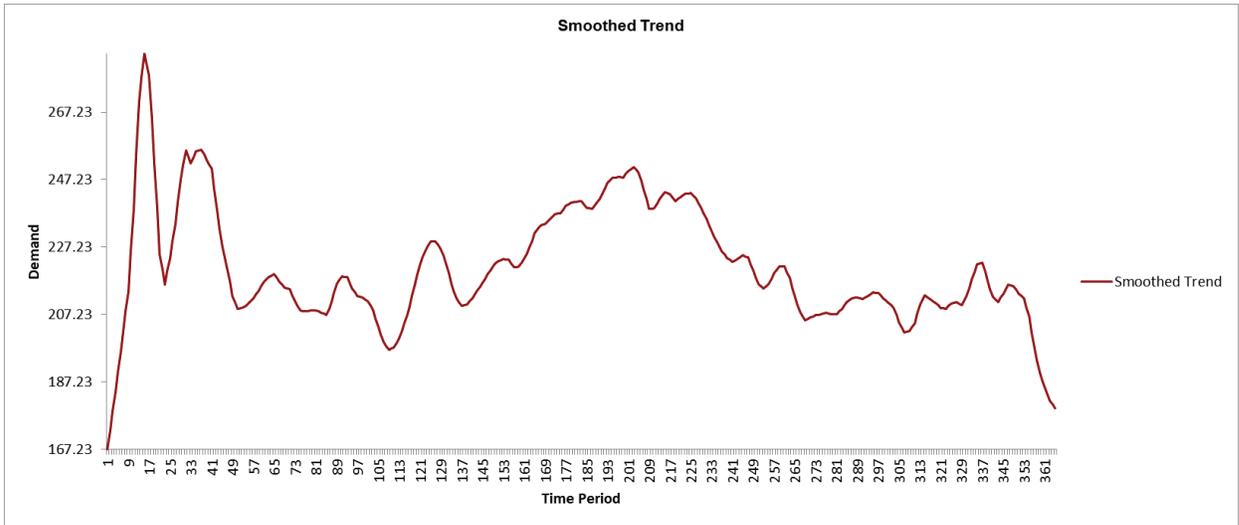


15. Click **OK**. A Seasonal Trend Decomposition Plot for Demand is produced.



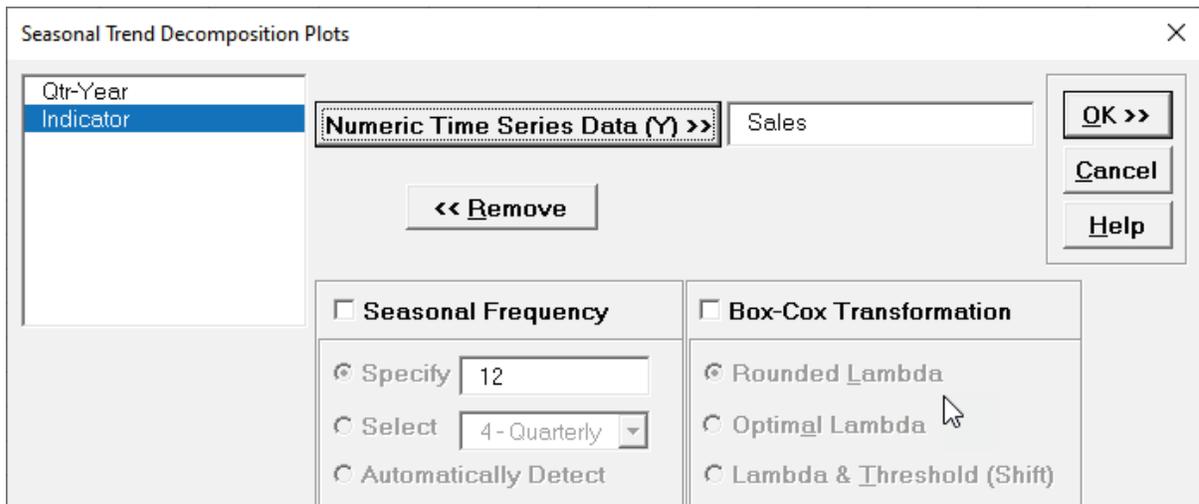
Here we see the daily seasonal effect. As discussed earlier, some of the trend patterns can be explained by the predictors Temp (C), TempSq and WorkDay.

16. The Smoothed Trend, Seasonal and Remainder Plots are also produced as shown:

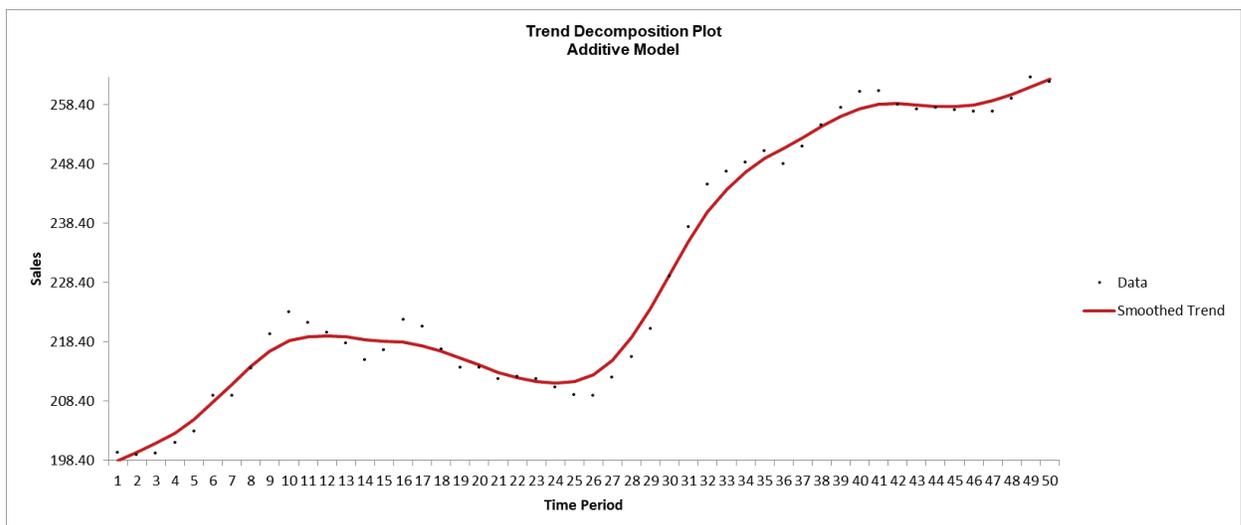


Sales with Indicator - Modified Series M

17. Open **Sales with Indicator - Modified Series M.xlsx**. (Sheet 1 tab). This is modified Series M data from Box and Jenkins, with 50 quarters of corporate sales values along with a leading indicator. The data is treated as nonseasonal, as done in Box and Jenkins.
18. Click **SigmaXL > Time Series Forecasting > Seasonal Trend Decomposition Plots**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
19. Select *Sales*, click **Numeric Time Series Data (Y) >>**. **Seasonal Frequency** and **Box-Cox Transformation** should be unchecked as shown.

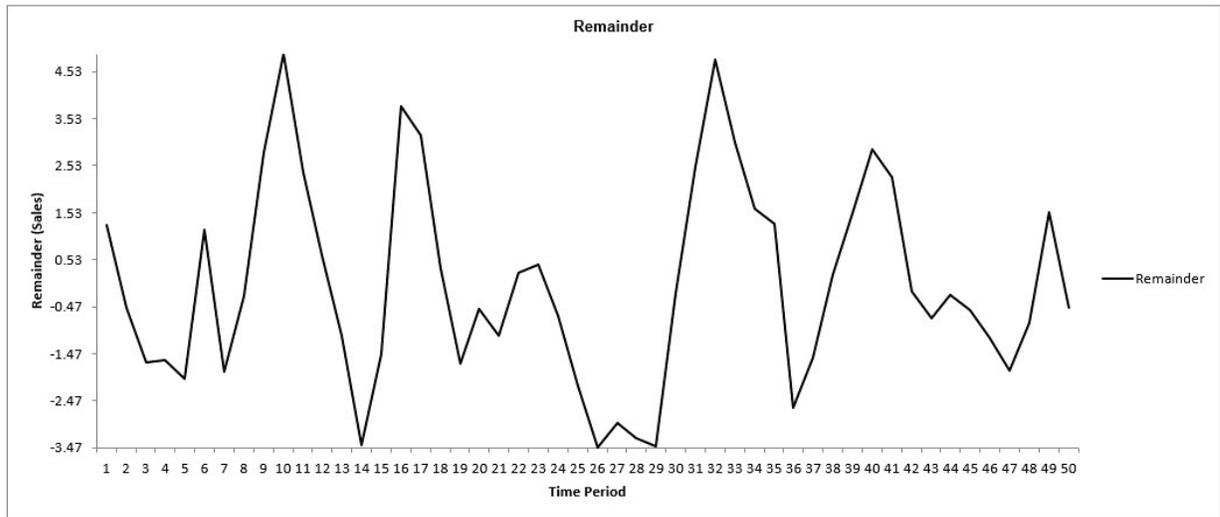
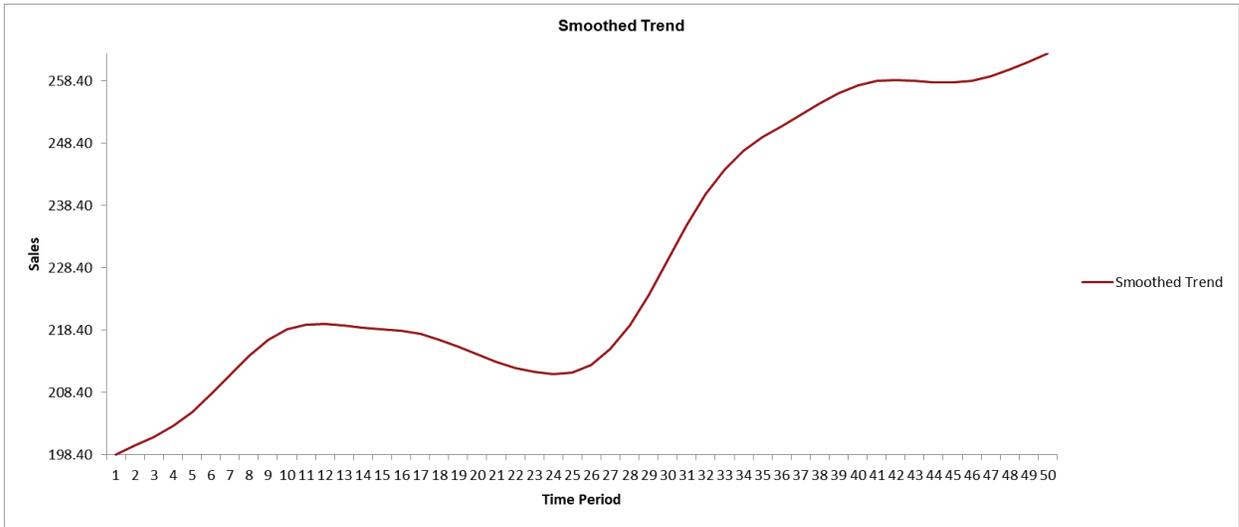


20. Click **OK**. A Trend Decomposition Plot for Concentration is produced.



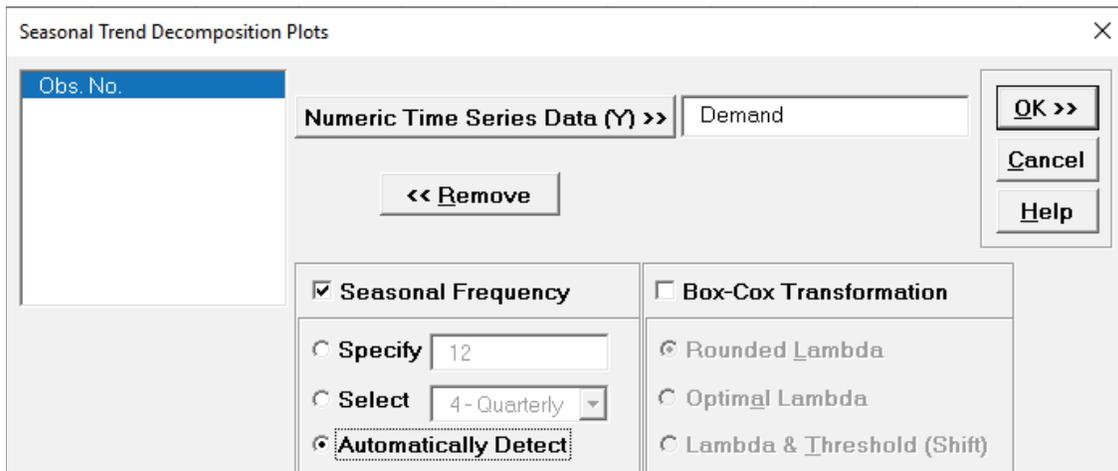
Here we see an overall positive trend over the 50 quarters.

21. The Smoothed Trend and Remainder Plots are also produced as shown:

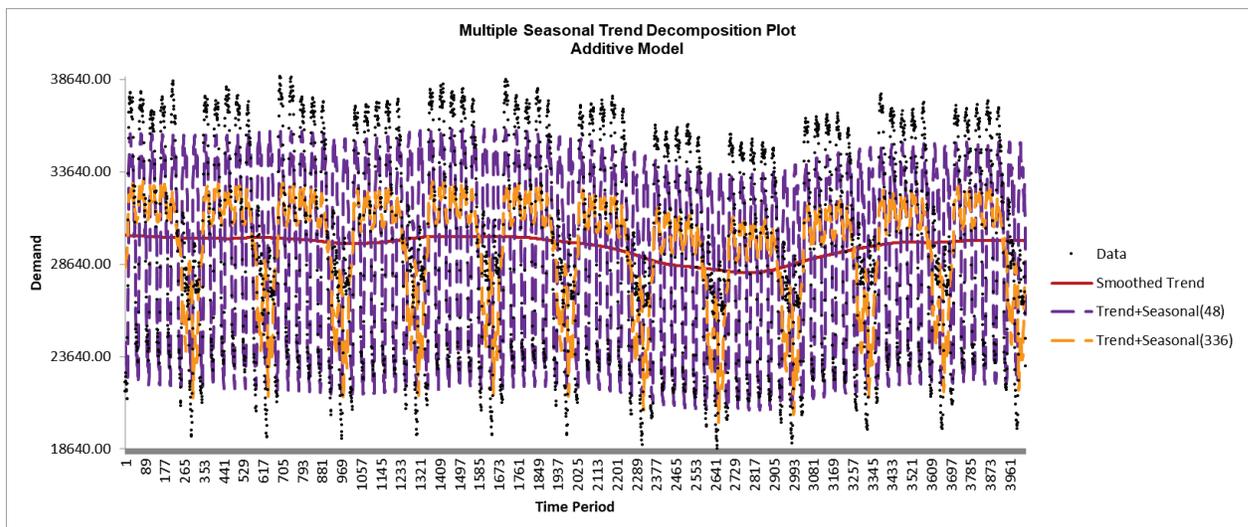


Half-Hourly Multiple Seasonal Electricity Demand – Taylor

22. Open **Half-Hourly Multiple Seasonal Electricity Demand - Taylor.xlsx (Sheet 1 tab)**. This is half-hourly electricity demand (MW) in England and Wales from Monday, June 5, 2000 to Sunday, August 27, 2000 (taylor, R forecast). This data has multiple seasonality with frequency = 48 (observations per day) and 336 (observations per week), with a total of 4032 observations.
23. Click **SigmaXL > Time Series Forecasting > Seasonal Trend Decomposition Plots**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
24. Select *Demand*, click **Numeric Time Series Data (Y) >>**. Check **Seasonal Frequency** with **Automatically Detect**. Uncheck **Box-Cox Transformation**.

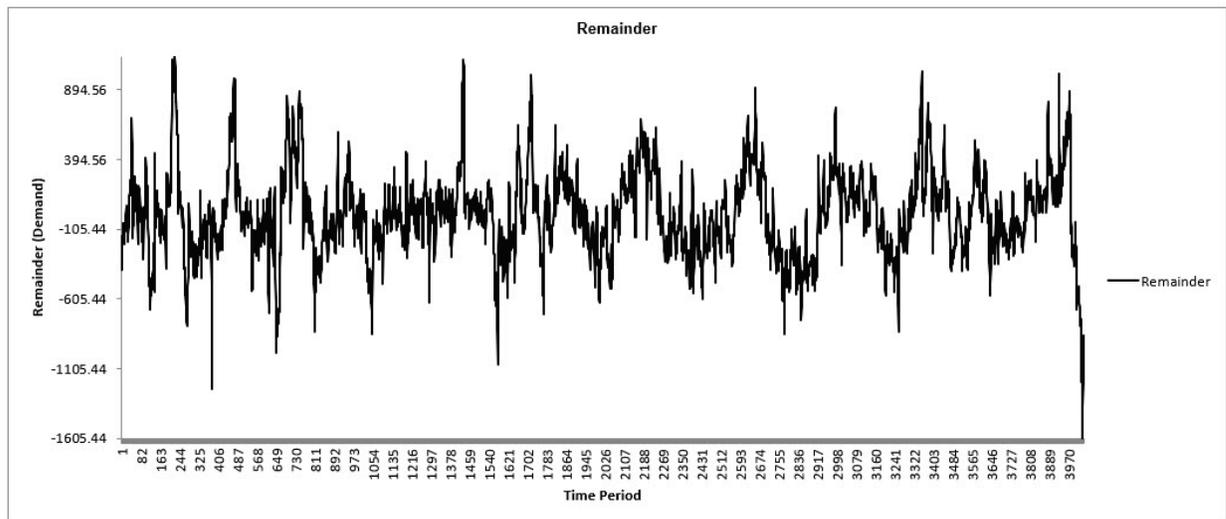
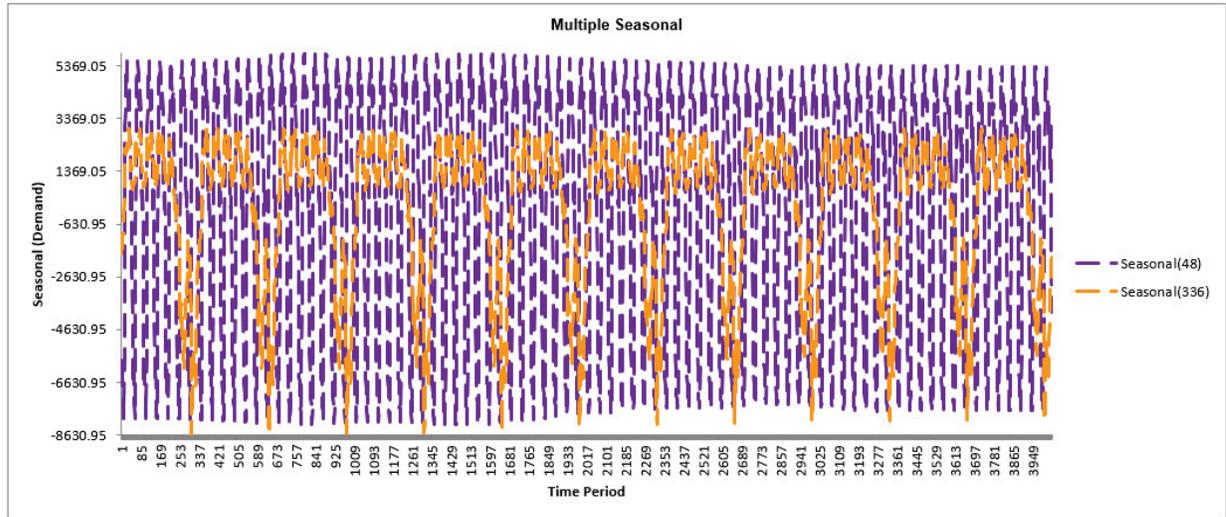
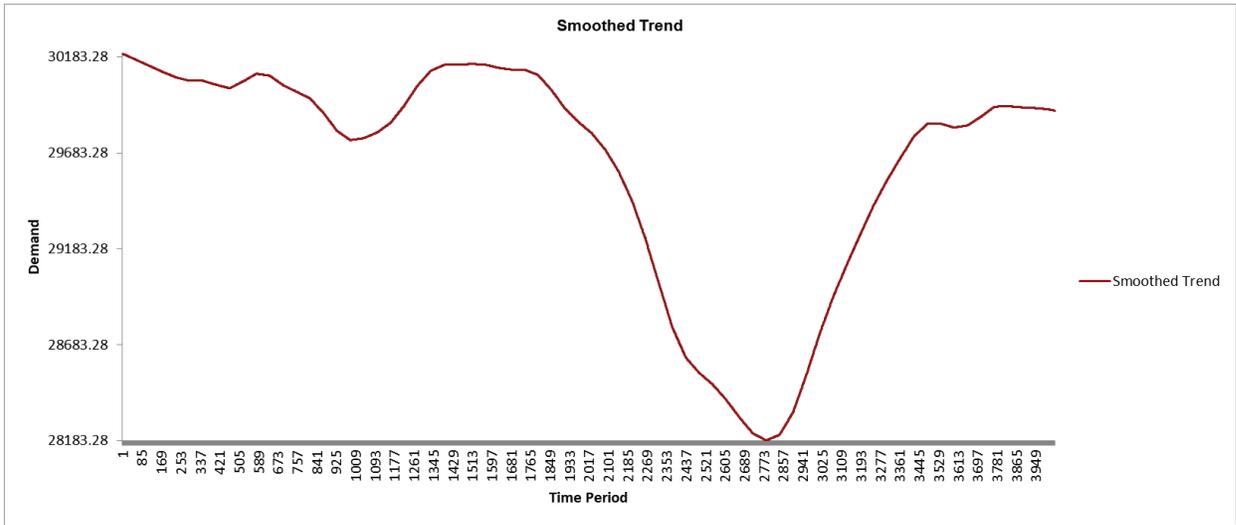


25. Click **OK**. A Seasonal Trend Decomposition Plot for Demand is produced.



Automatic detection of seasonal frequency gives 48, 336, as obtained earlier with the Spectral Density Plot. Here we see the half-hourly multiple seasonal effect with 48 observations per day and 336 observations per week.

26. The Smoothed Trend, Seasonal and Remainder Plots are also produced as shown:



Seasonal Interaction Plots

Bisgaard and Kulahchi (2011) give a novel use of two-way interaction plots to view trends and seasonal effects in data. We will use Two-Way ANOVA to reproduce the interaction charts given in the book.

Note, in order to produce these charts, the data must be balanced, e.g., every year must have 12 months of data.

Reference:

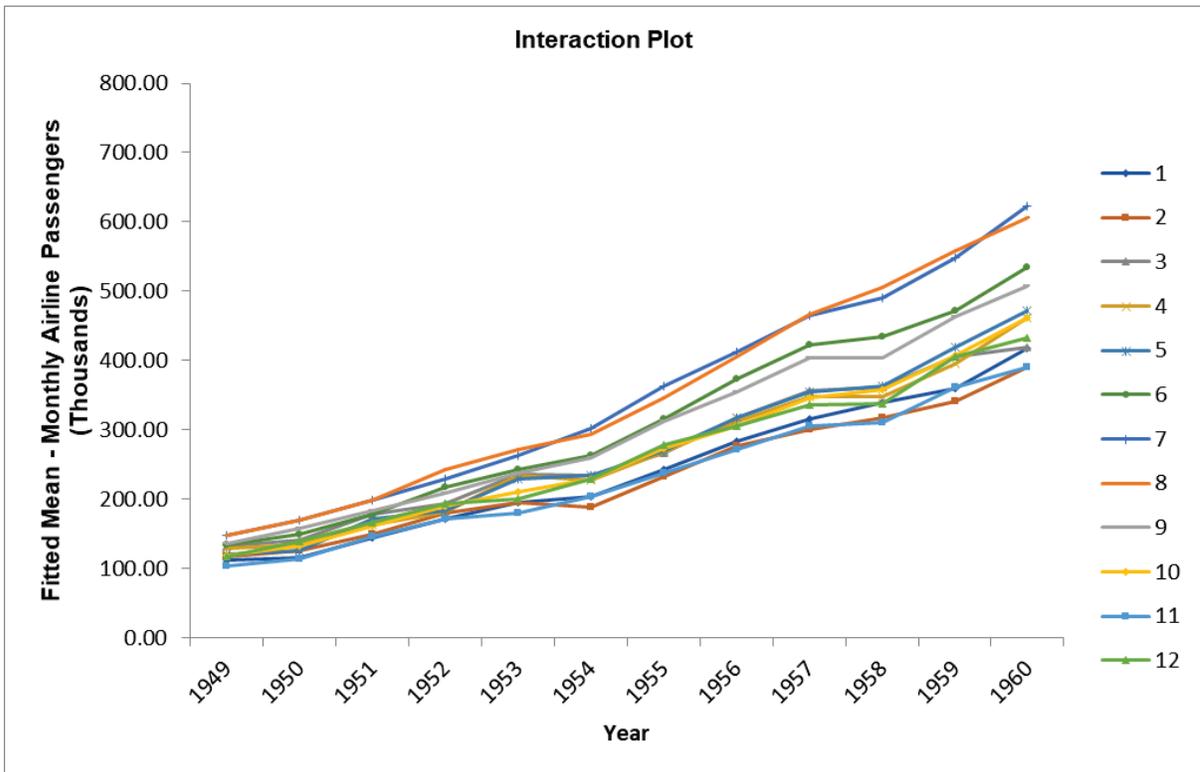
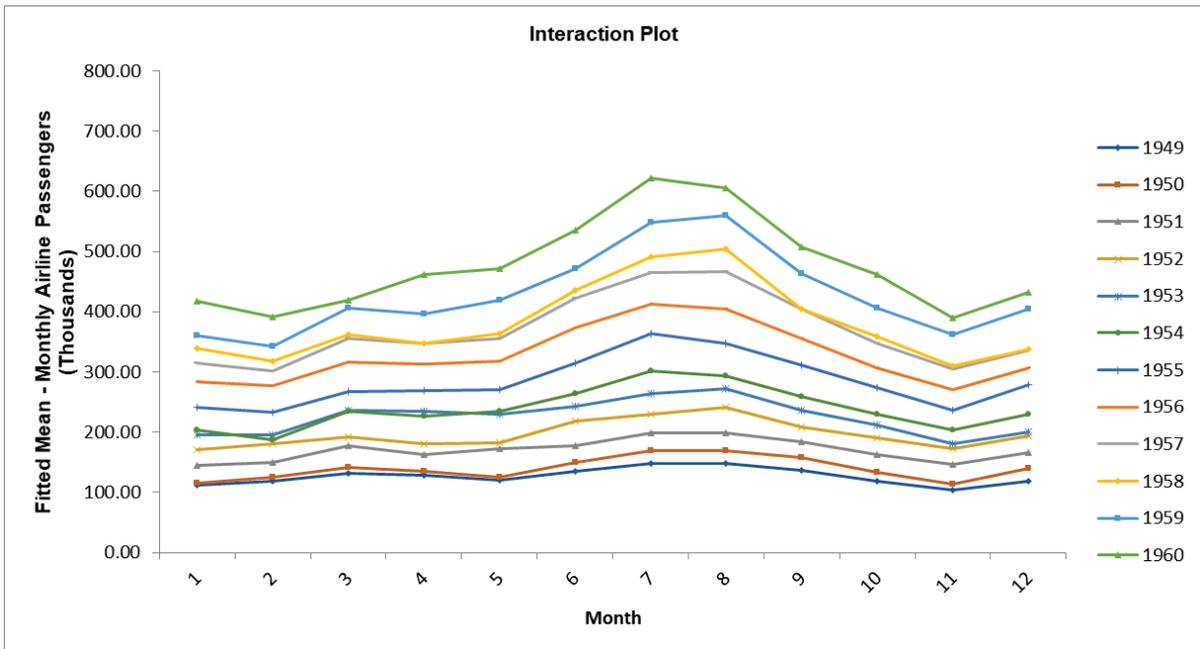
Bisgaard, S. and Kulahchi, M. (2011), *Time Series Analysis and Forecasting by Example*, Wiley, pp. 111-115.

Monthly Airline Passengers - Series G

1. Open **Monthly Airline Passengers - Series G.xlsx**. Click the **Month Year for Interaction Plot** tab. This is the Series G data from Box and Jenkins, monthly total international airline passengers for January 1949 to December 1960. Month and Year columns have been added and calculated using the Excel date functions =MONTH() and =YEAR().
2. Click **SigmaXL > Statistical Tools > Two-Way ANOVA**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
3. Select *Monthly Airline Passengers*, click **Numeric Data Variable (Y) >>**. Select *Month* for **Group Category Factor (X1) >>** and *Year* for **Group Category Factor (X2) >>**. Uncheck all options: **Remove Interaction (Fit Additive Model)**, **Display Residual Plots**, **Display ANOM Normal Two-Way Chart** and **Adjust chart alpha for family-wise error rate**. Use the default **Confidence Level = 95%**.

The screenshot shows the 'Two-Way ANOVA' dialog box. On the left, a list of variables includes 'Obs. No.', 'Date', and 'Ln (Airline Passengers)'. The 'Numeric Data Variable (Y) >>' field contains 'Monthly Airline Passengers'. The 'Group Category Factor (X1) >>' field contains 'Month', and the 'Group Category Factor (X2) >>' field contains 'Year'. Below these fields is a '<< Remove' button. On the right side, there are four unchecked checkboxes: 'Remove Interaction (Fit Additive Model)', 'Display Residual Plots', 'Display ANOM Normal Two-Way Chart', and 'Adjust chart alpha for family-wise error rate'. At the bottom, the 'Confidence Level:' is set to '95.0'. On the far right, there are 'OK >>', 'Cancel', and 'Help' buttons.

- Click **OK**. We will not use the ANOVA report. Scroll down to the Interaction Plots. Resize to view the full legend, double click on each Y axis and set Minimum to 0.



In the first interaction plot (with Month on the X axis) we can see the monthly seasonal effect and how it gets stronger by year. The second interaction plot (with Year on the X axis) shows the same increasing seasonal effect but we can also clearly see the strong positive trend by year.

- Now we will create Seasonal Interaction Plots for Ln (Airline Passengers). Click **Recall SigmaXL Dialog** menu or press **F3** to recall last dialog. Select *Ln (Airline Passengers)* and click **Numeric Data Variable (Y) >>**.

Two-Way ANOVA

Obs. No.
Date
Monthly Airline Pass

Numeric Data Variable (Y) >> Ln (Airline Passengers)

Group Category Factor (X1) >> Month

Group Category Factor (X2) >> Year

<< Remove

Remove Interaction (Fit Additive Model)

Display Residual Plots

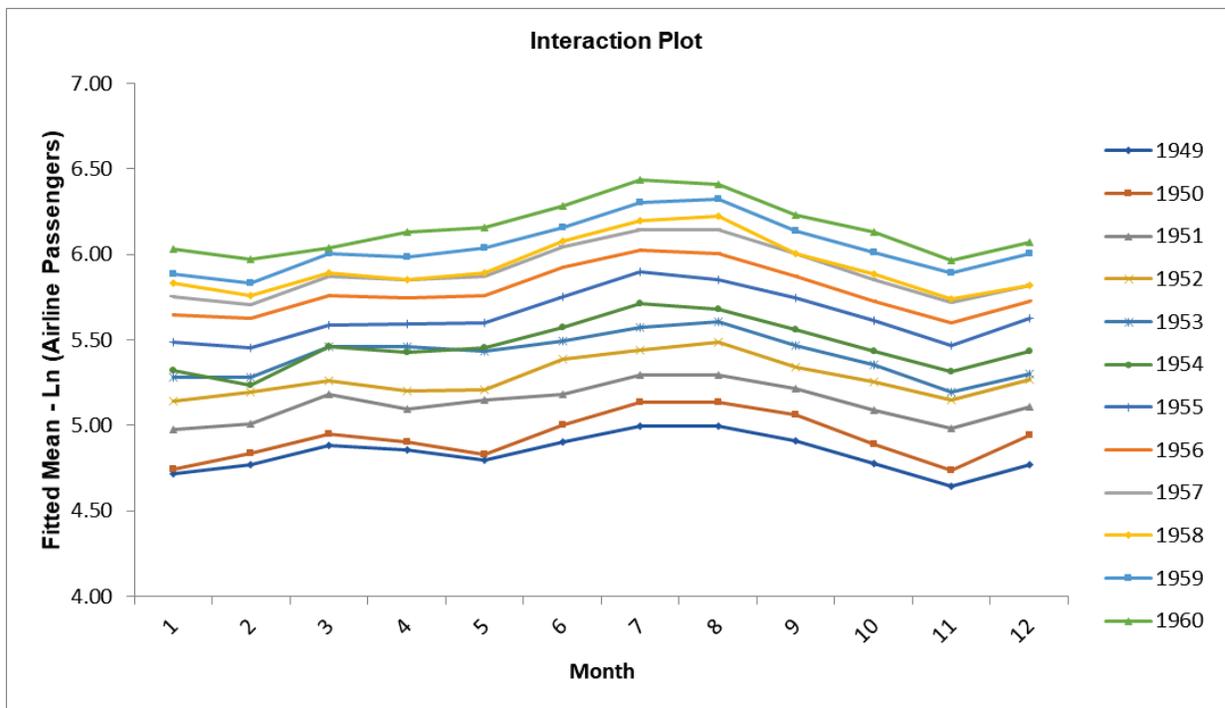
Display ANOM Normal Two-Way Chart

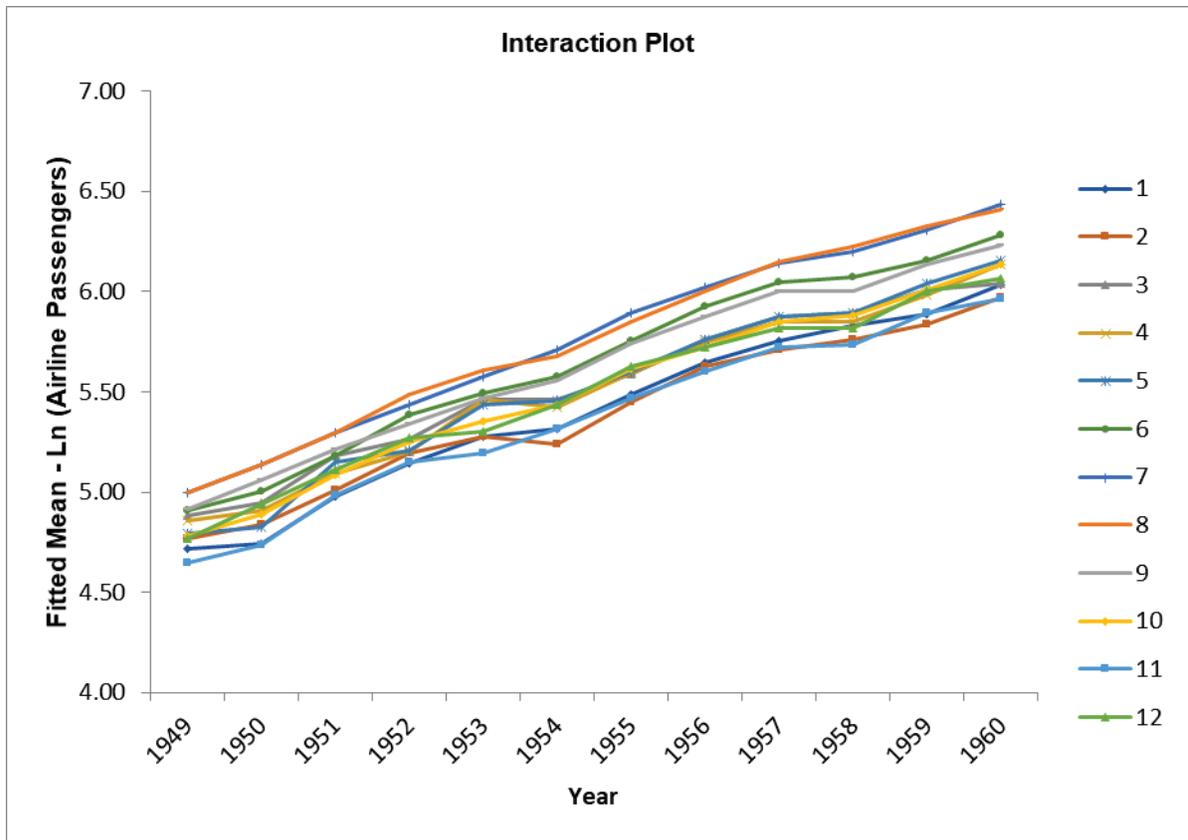
Adjust chart alpha for family-wise error rate

Confidence Level: 95.0

OK >>
Cancel
Help

- Click **OK**. Scroll down to the Interaction Plots. Resize to view the full legend, double click on each Y axis and set Minimum to 4 and Maximum to 7.





In these interaction plots with Ln (Airline Passengers), we can see that the variability due to monthly seasonal effect is consistent and the yearly trend is more linear. Bisgaard and Kulahchi point out that using a traditional interpretation of interaction plots, the similar slopes indicate that the Ln transformation has effectively removed the month by year interaction, so the month and year effect is now additive.

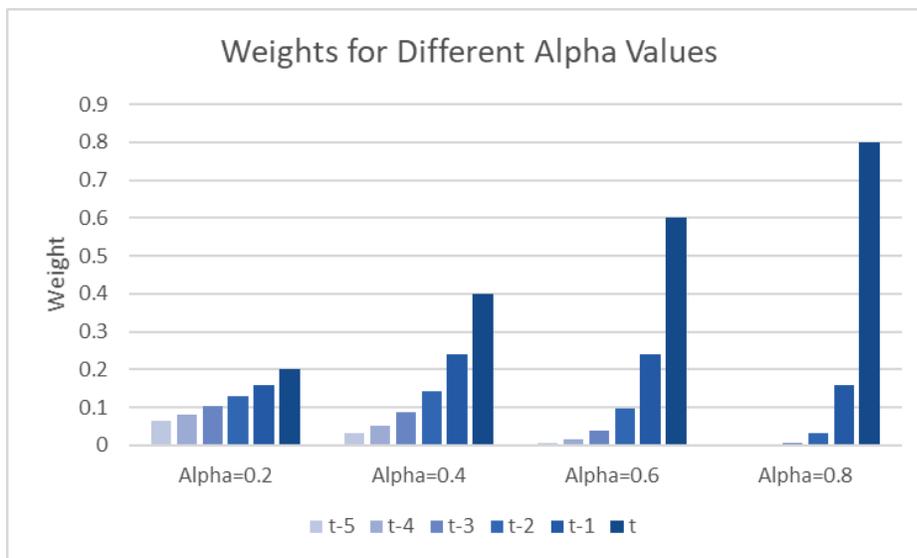
Exponential Smoothing Forecast

Simple Exponential Smoothing

Simple Exponential Smoothing forecasts are calculated using weighted averages, where the weights decrease exponentially as observations come from further in the past with the smallest weights associated with the oldest observations:

$$\hat{y}_{t+1} = \alpha y_t + \alpha(1 - \alpha) y_{t-1} + \alpha(1 - \alpha)^2 y_{t-2} + \dots$$

where $0 \leq \alpha \leq 1$ is the smoothing parameter. This is also known as the Exponentially Weighted Moving Average (EWMA). Weights for different α values are shown graphically:



An equivalent formulation for simple exponential smoothing is:

$$\hat{y}_{t+1} = \alpha y_t + (1 - \alpha) \hat{y}_t$$

with the starting forecast value (initial level) \hat{y}_1 to be estimated and denoted as level l_0 .

The smoothing parameter α and initial level l_0 are determined by minimizing the sum-of-square forecast errors (residuals):

$$SSE = \sum_{t=1}^T (y_t - \hat{y}_t)^2 = \sum_{t=1}^T e_t^2.$$

Error, Trend, Seasonal (ETS) Models

Error, Trend, Seasonal (ETS) models expand on simple exponential smoothing to accommodate trend and seasonal components as well as additive or multiplicative errors. Simple Exponential Smoothing is an Error Model. Error, Trend model is Holt's Linear, also known as double exponential smoothing. Error, Trend, Seasonal model is Holt-Winters, also known as triple exponential smoothing.

Rob Hyndman has developed a complete taxonomy that describes all of the combinations of exponential smooth models in a consistent manner (see fpp2):

- Error:
 - Additive or Multiplicative. Multiplicative error is a relative error: $e_t = \frac{y_t - \hat{y}_t}{\hat{y}_t}$
 - The point forecasts produced by the models are identical if they use the same smoothing parameter values. Multiplicative will, however, generate different prediction intervals to accommodate change in variance.
 - An alternative to multiplicative is to use the Ln transformation (Box-Cox transformation with Lambda = 0).
 - Error models include the smoothing parameter α and initial level value.
- Trend:
 - None, Additive, Additive Damped
 - Multiplicative Trend is not recommended as it tends to produce poor forecasts, so is not included in SigmaXL
 - Trend models add a smoothing parameter β and initial trend value.
 - Damped trend models add a smoothing parameter ϕ that "dampens" the trend to a flat line some time in the future.
- Seasonal:
 - None, Additive, Multiplicative
 - The additive method is preferred when the seasonal variations are roughly constant through the series, while the multiplicative method is preferred when the seasonal variations are changing proportional to the level of the series.
 - Seasonal models add a smoothing parameter γ and initial seasonal values.
 - Maximum seasonal frequency is 24. For higher frequencies use Exponential Smoothing – Multiple Seasonal Decomposition (MSD)
 - # of initial values estimated = seasonal frequency – 1
 - constrained to sum to 0 for additive or seasonal frequency m for multiplicative

For further details, see Appendix: [Exponential Smoothing - ETS](#).

Summary of ETS Models in SigmaXL:

Short hand (Error, Trend, Seasonal)	Method
(A, N, N)	Simple Exponential Smoothing with Additive Errors - Exponentially Weighted Moving Average (EWMA)
(M, N, N)	Simple Exponential Smoothing with Multiplicative Errors
(A, A, N)	Additive Trend Method with Additive Errors (Holt's Linear)
(M, A, N)	Additive Trend Method with Multiplicative Errors (Holt's Linear)
(A, A, A)	Additive Trend, Additive Seasonal Method with Additive Errors (Holt-Winters)
(M, A, A)	Additive Trend, Additive Seasonal Method with Multiplicative Errors (Holt-Winters)
(A, N, A)	Additive Seasonal Method with Additive Errors
(M, N, A)	Additive Seasonal Method with Multiplicative Errors
(A, Ad, N)	Additive Damped Trend Method with Additive Errors
(M, Ad, N)	Additive Damped Trend Method with Multiplicative Errors
(A, Ad, A)	Additive Damped Trend, Additive Seasonal Method with Additive Errors
(M, Ad, A)	Additive Damped Trend, Additive Seasonal Method with Multiplicative Errors
(M, A, M)	Additive Trend, Multiplicative Seasonal Method with Multiplicative Errors (Holt-Winters)
(M, N, M)	Multiplicative Seasonal Method with Multiplicative Errors
(M, Ad, M)	Additive Damped Trend, Multiplicative Seasonal Method with Multiplicative Errors (Holt-Winters)

Exponential Smoothing Parameter Estimation, Model Statistics and Information Criteria for Model Comparison

Model parameters are solved using nonlinear maximization of the Log-Likelihood function. Log-Likelihood is related to $-\ln(\text{Sum-of-Squares Error})$. Information Criteria AICc, AIC and BIC are calculated using $-2 \times \text{Log-Likelihood}$ and incorporate a penalty for the number of terms in the model, so smaller is better. These are used in automatic model selection. AICc is the default Information Criterion, based on forecast error performance with competition data. For further details, see Appendix: [Information Criteria for Model Comparison](#).

Missing Values

Exponential Smoothing handles missing values with seasonally adjusted linear interpolation. While there is robustness against some missing values, if the number of missing values is large then model estimation and forecast accuracy will be degraded. Upon selection of a time series with missing values, you will see a pop-up “Warning: Missing values detected. Seasonally adjusted linear interpolation will be used.” See Appendix: [Seasonally Adjusted Linear Interpolation of Missing Values](#).

Forecast Accuracy Metrics

SigmaXL uses the following forecast accuracy metrics:

$$\text{Root mean squared error: RMSE} = \sqrt{\text{mean}(e_t^2)}$$

$$\text{Mean absolute error: MAE} = \text{mean}(|e_t|)$$

$$\text{Mean absolute percentage error: MAPE} = \text{mean}\left(\left|\frac{100e_t}{y_t}\right|\right)$$

$$\text{Mean absolute scaled error: MASE} = \text{mean}(|e_t|)/\text{scale}$$

where *scale* is the MAE of the in-sample naïve or seasonal naïve forecast (set all forecasts to be the value of the last observation/period). Note that scale for nonseasonal is identical to MR-bar of an Individuals Moving Range chart. A scaled error is less than one if it arises from a better forecast than the average naïve/seasonal naïve forecast. Conversely, it is greater than one if the forecast is worse than the average naïve forecast. MASE is recommended over the popular MAPE, because MAPE becomes infinite if any $y_t=0$. For further details, see Appendix: [Forecast Accuracy](#).

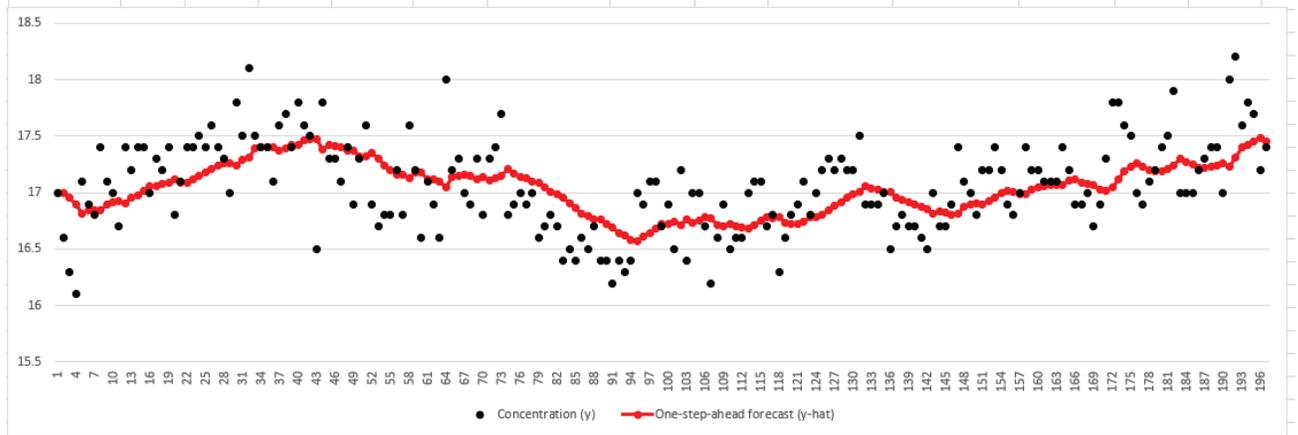
In-Sample is also referred to as the “Train” data. The same metrics may be applied to Out-of-Sample (Withhold), also referred to as the “Test” data and may be One-Step-Ahead or Multi-Step-Ahead. This will be demonstrated later. Out-of-Sample (Withhold) data is not used in the model parameter estimation, so is a much better indicator of true forecast accuracy. In-sample accuracy metrics can be biased due to overfitting. Scale is always computed using the in-sample data.

Demo of Simple Exponential Smoothing – Concentration

1. Open **Demo of Simple Exponential Smoothing – Concentration.xlsx (Sheet 1 tab)**. This is the Series A data from Box and Jenkins, a set of 197 concentration values from a chemical process taken at two-hour intervals in Column B.

	A	B	C	D
1	Obs. No.	Concentration (y)	One-step-ahead forecast (y-hat)	Residuals (e)
2	1	17	17	0
3	2	16.6	17	-0.4
4	3	16.3	16.96	-0.66
5	4	16.1	16.894	-0.794
6	5	17.1	16.8146	0.2854
7	6	16.9	16.84314	0.05686
8	7	16.8	16.848826	-0.048826
9	8	17.4	16.8439434	0.5560566
10	9	17.1	16.89954906	0.20045094
11	10	17	16.91959415	0.080405846

Exponential Smoothing Model Parameters		Exponential Smoothing Model Statistics		In-Sample Forecast Accuracy Metrics	
alpha	0.1	Count	197	RMSE	0.332916231
initial level	17	k terms	2	MAE	0.271930445
Use Solver to obtain parameter values that minimize SSE:		Log-Likelihood	-303.72228	MAPE	1.592107249
SSE	21.8341	AICc	613.56892	Scale (MR-bar)	0.275510204
		AIC	613.44457	MASE	0.9870068
		BIC	623.29418		

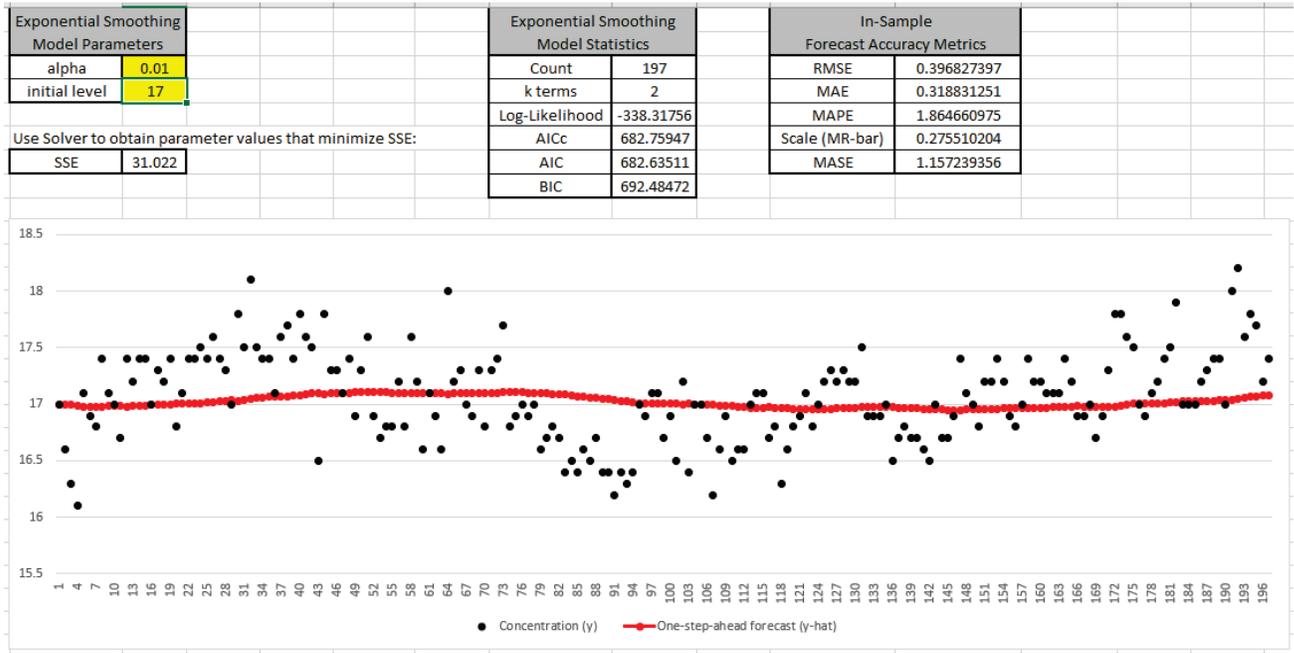


Cell **G3** is the alpha smoothing parameter value, cell **G4** is the initial level. These may be manually entered or solved with Solver. The One-step-ahead forecast is calculated as follows: C2 is the initial level value; $C3 = \alpha * B2 + (1 - \alpha) * C2$; $C4 = \alpha * B3 + (1 - \alpha) * C3$, etc.

Cell **G7** is the Residuals Sum-of-Squares error and is the value to be minimized, in order to produce the most accurate one-step-ahead forecast. The exponential smooth model statistics are also given for later comparison to the SigmaXL report.

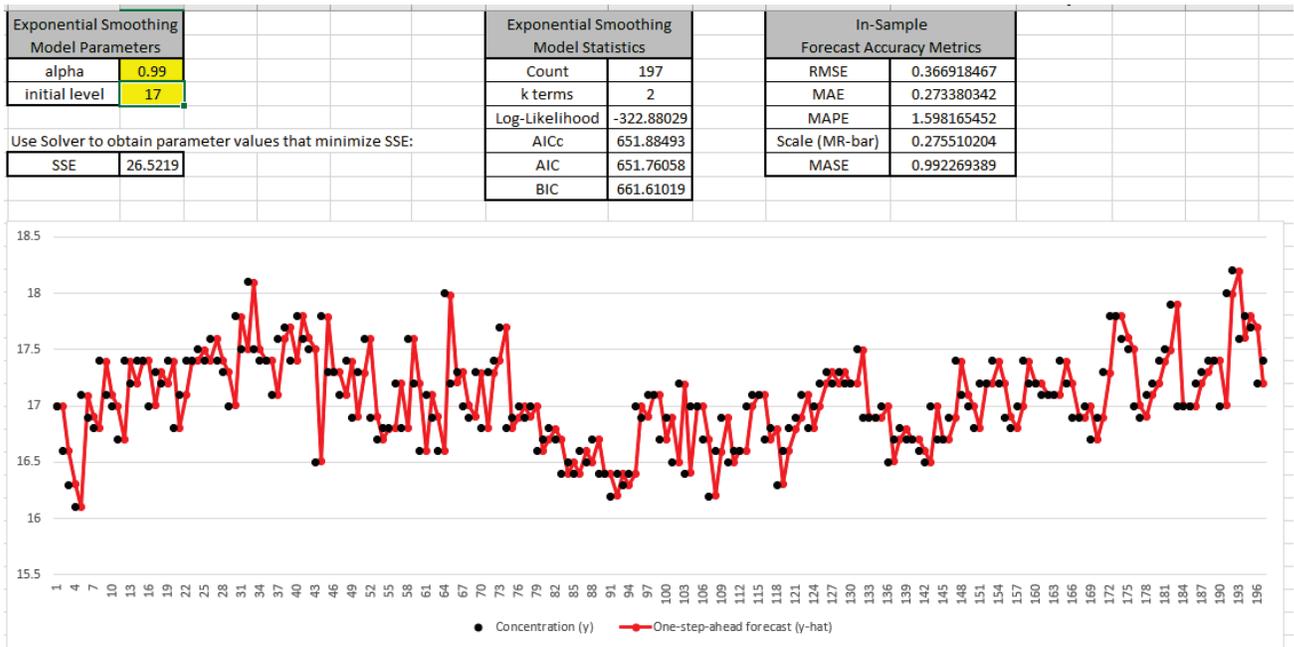
The raw Concentration data are the black dots in the graph; the one-step-ahead forecast values are the red dots. Model Statistics and Forecast Accuracy Metrics will be discussed later.

- Before we use Solver to determine the optimal alpha and initial level values, we will manually enter some values to see how the graph changes. Enter alpha = 0.01 as shown:



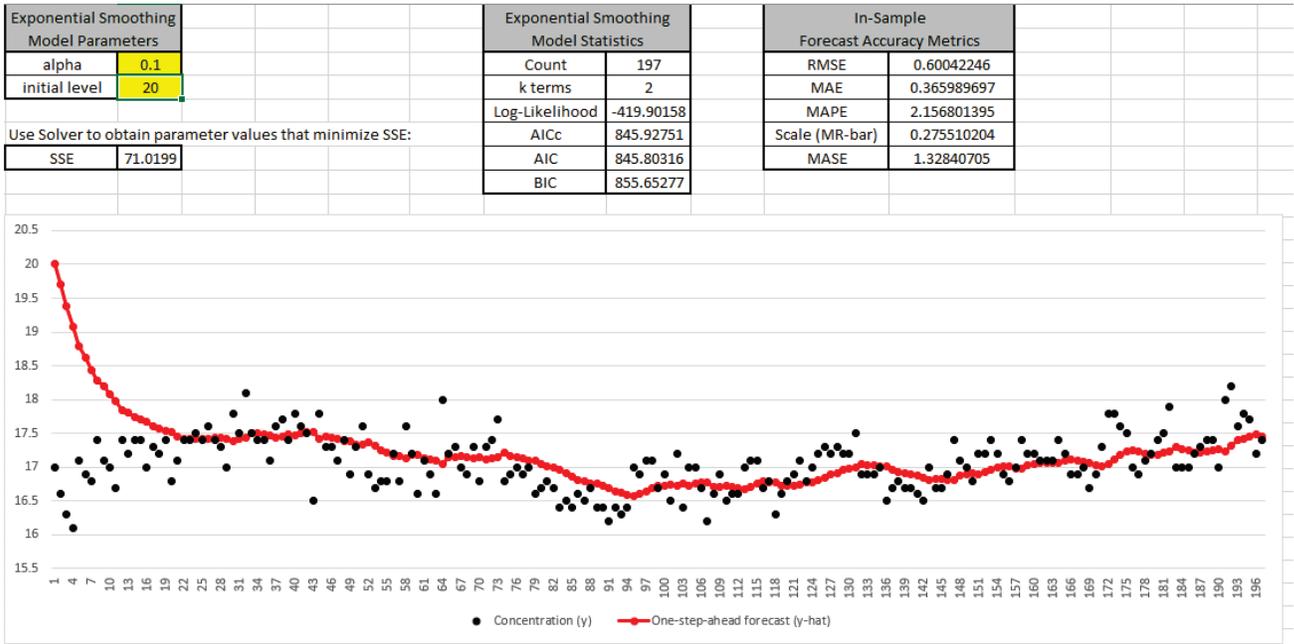
This small value of alpha produces a smooth fit with SSE = 31.02 that is larger than what we started with (SSE = 21.83).

- Now enter alpha = 0.99 as shown:



This large value of alpha results approximately in a naïve forecast, with the one-step-ahead forecast value equal to the previous actual. SSE = 26.52 is larger than what we started with.

4. Enter alpha = 0.1 and initial level = 20 as shown:



This initial level is a poor estimate of the start value and results in SSE = 71.02. After about 30 observations, the influence of this poor start value is negligible. The duration for which the initial level has influence on the one-step-ahead forecast depends on alpha, a smaller alpha would be a longer duration; a larger alpha would be a shorter duration.

5. Now we will use Excel’s Solver to find the optimal values for alpha and initial level. Click **Data > Solver**. (If the Solver menu is not available, click **File > Options > Add-Ins**. Click **Go...** Manage Excel Add-ins, check **Solver Add-in**. Click **OK**. Click **Data > Solver**.)

Solver Parameters [X]

Set Objective: [↑]

To: Max Min Value Of:

By Changing Variable Cells: [↑]

Subject to the Constraints:

alpha <= 1
alpha >= 0

[Add] [Change] [Delete] [Reset All] [Load/Save]

Make Unconstrained Variables Non-Negative

Select a Solving Method: [Options]

Solving Method
Select the GRG Nonlinear engine for Solver Problems that are smooth nonlinear. Select the LP Simplex engine for linear Solver Problems, and select the Evolutionary engine for Solver problems that are non-smooth.

[Help] [Solve] [Close]

The cell addresses, constraints and settings for Solver have been stored in the workbook, so no changes are necessary. Solver will minimize cell **G7** (SSE), by varying cells **G3** (alpha) and **G4** (initial_level). The solving method is GRG Nonlinear; Evolutionary could also be used but it is slower.

- Click **Solve**. The Solver Results dialog shows that a solution has been found.

Solver Results [X]

Solver has converged to the current solution. All Constraints are satisfied.

Keep Solver Solution
 Restore Original Values

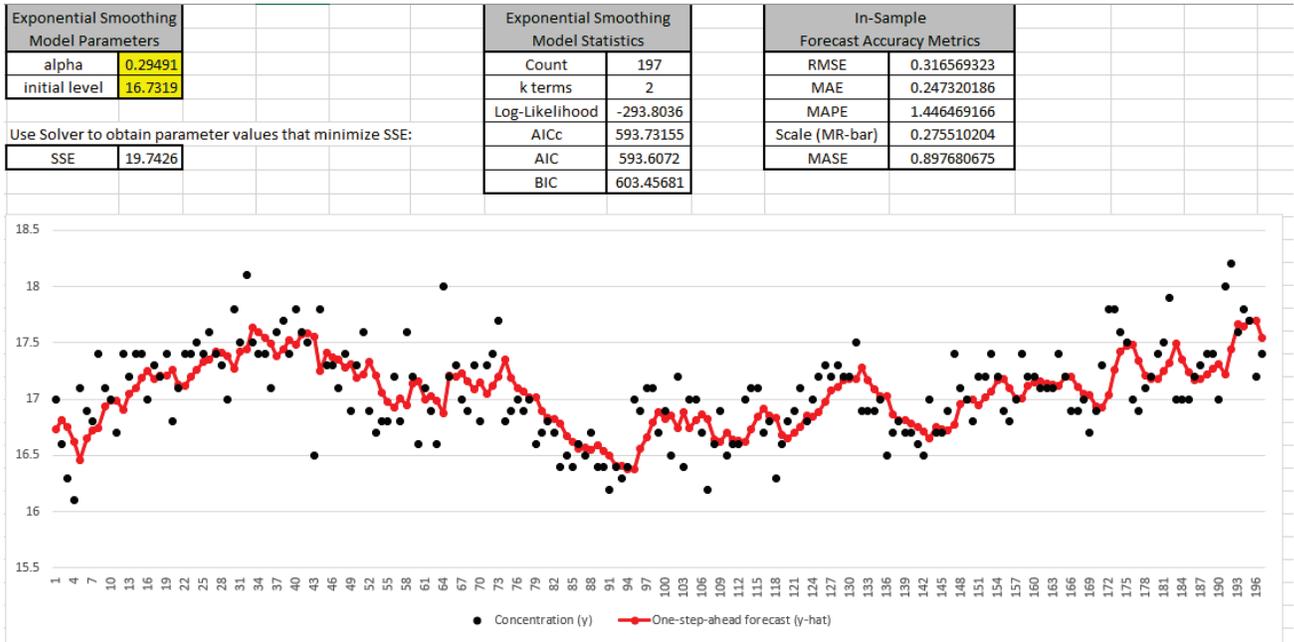
Return to Solver Parameters Dialog Outline Reports

[OK] [Cancel] [Save Scenario...]

Reports
Answer
Sensitivity
Limits

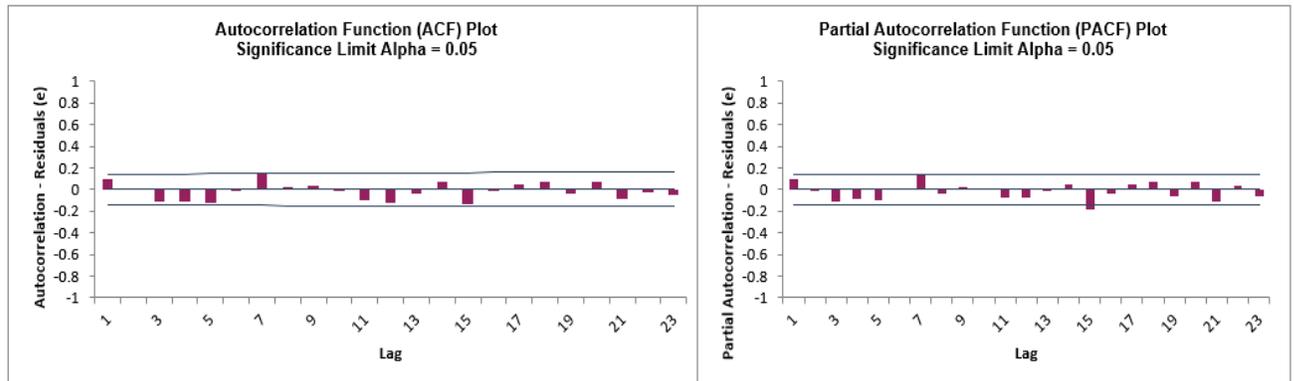
Solver has converged to the current solution. All Constraints are satisfied.
Solver has performed 5 iterations for which the objective did not move significantly. Try a smaller convergence setting, or a different starting point.

7. Click **OK**. The SSE is now 19.74; alpha value = 0.295 and initial level = 16.732.



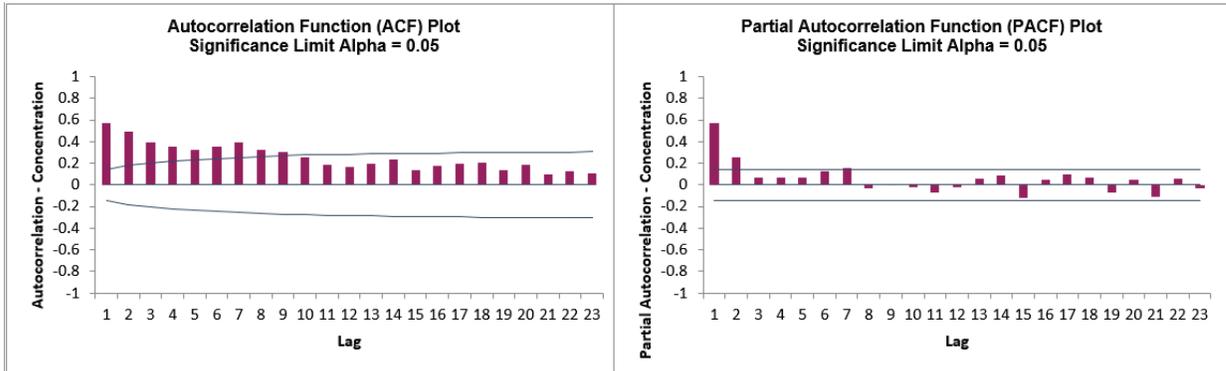
By minimizing the residual sum-of-squares error (SSE), we have our best estimate of the parameters to produce one-step-ahead forecast values.

8. We will now produce ACF/PACF plots for the Residuals. Select cells **D1:D198** and click **SigmaXL > Time Series Forecasting > Autocorrelation (ACF/PACF) Plots**. Click **Next**, Select *Residuals (e)*, Click **OK**.

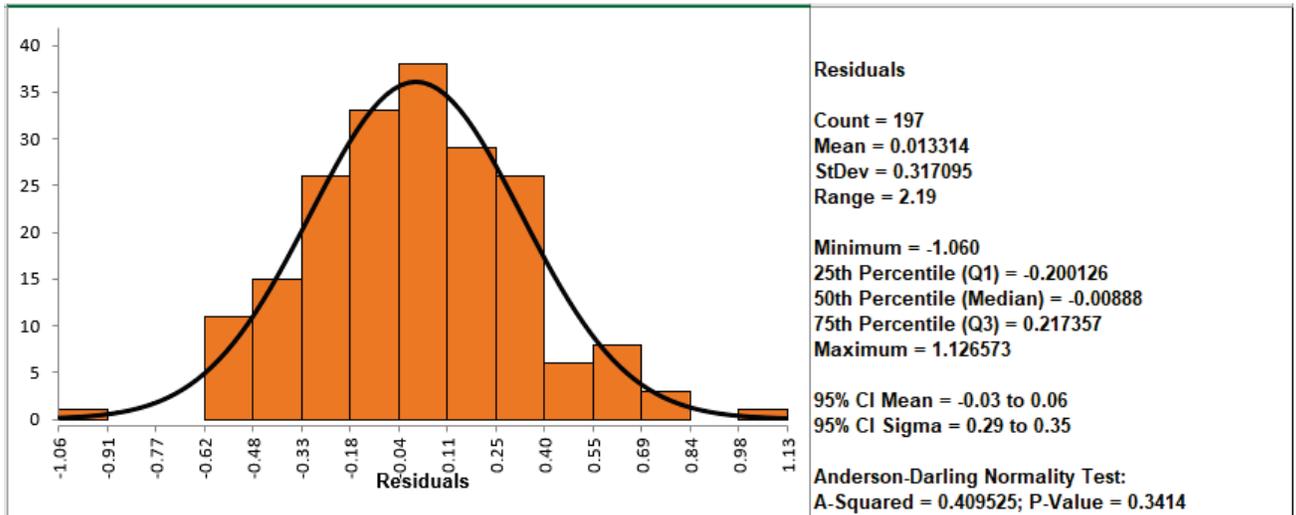


We can see that all of the autocorrelation has been removed by the simple exponential smoothing model (with the exception of lag 15 in the PACF), so this is a good fit to the time series data.

9. By way of comparison, these are the ACF/PACF plots for the original Concentration data that we produced earlier.



10. We will now have a quick look at the Residuals using a Histogram. Select the Residuals data (D1:D198). Click **SigmaXL > Graphical Tools > Histograms and Descriptive Statistics**. Next, Select *Residuals (e)*, Click **OK**.



The residuals are normally distributed with no obvious extreme outliers.

Chemical Process Concentration - Series A

- Open **Chemical Process Concentration – Series A.xlsx (Sheet 1 tab)**. This is the Series A data from Box and Jenkins, a set of 197 concentration values from a chemical process taken at two-hour intervals. See the [Run Chart](#), [ACF/PACF Plots](#), [Spectral Density Plot](#) and [Seasonal Trend Decomposition Plots](#) for this data.
- Click **SigmaXL > Time Series Forecasting > Exponential Smoothing Forecast > Forecast**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
- Select *Concentration*, click **Numeric Time Series Data (Y) >>**. Check **Display ACF/PACF/LB Plots** and **Display Residual Plots**. Leave **Specify Model Periods**, **Seasonal Frequency** and **Box-Cox Transformation** unchecked. We will use the default **No. of Forecast Periods = 24** and **Prediction Interval = 95.0 %**.

Exponential Smoothing Forecast

Observation No.

Numeric Time Series Data (Y) >> Concentration

Optional Time Axis Labels >>

<< Remove

OK >>
Cancel
Help

No. of Forecast Periods 24

Prediction Interval 95.0 %

Specify Model Periods

Start Model at Period 1

Withhold Periods 0

End Model at Period

Model Options

Seasonal Frequency

Specify 12

Select 4-Quarterly

Automatically Detect

Display ACF/PACF/LB Plots

Display Residual Plots

Box-Cox Transformation

Rounded Lambda

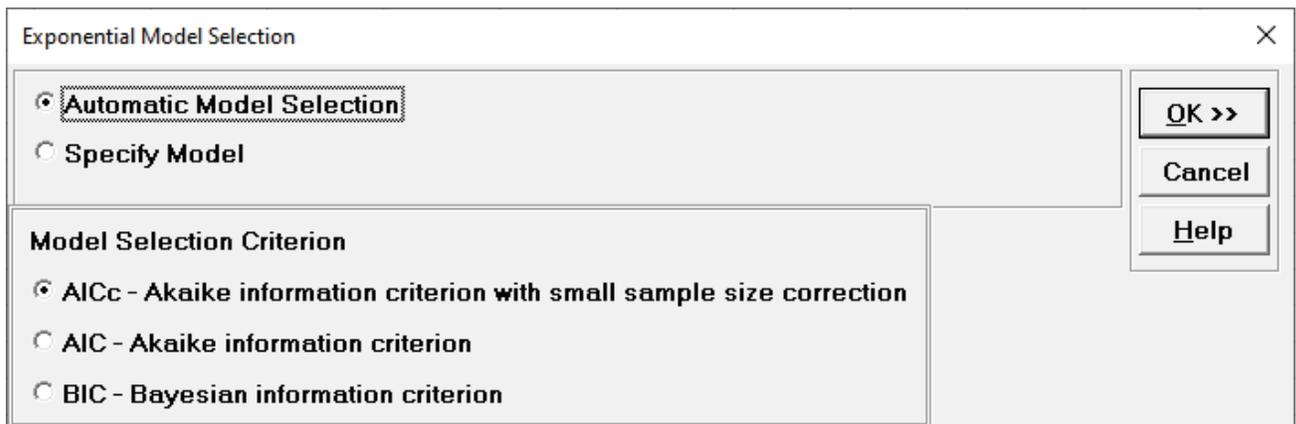
Optimal Lambda

Lambda & Threshold (Shift)

- **Optional Time Axis Labels** will be displayed on the forecast chart time axis. If used, dates for the forecast periods should also be included, otherwise the time axis will be blank for the forecast periods.
- **No. of Forecast Periods** are the number of time series values to be predicted (forecast horizon). The most accurate forecast will be for the first predicted value (one-step-ahead).
- **Prediction Interval %** is the confidence level for the individual predictions. For example, a 95% prediction interval contains a range of values which should include the actual future value with 95% probability. The interval will get larger the further out you predict.

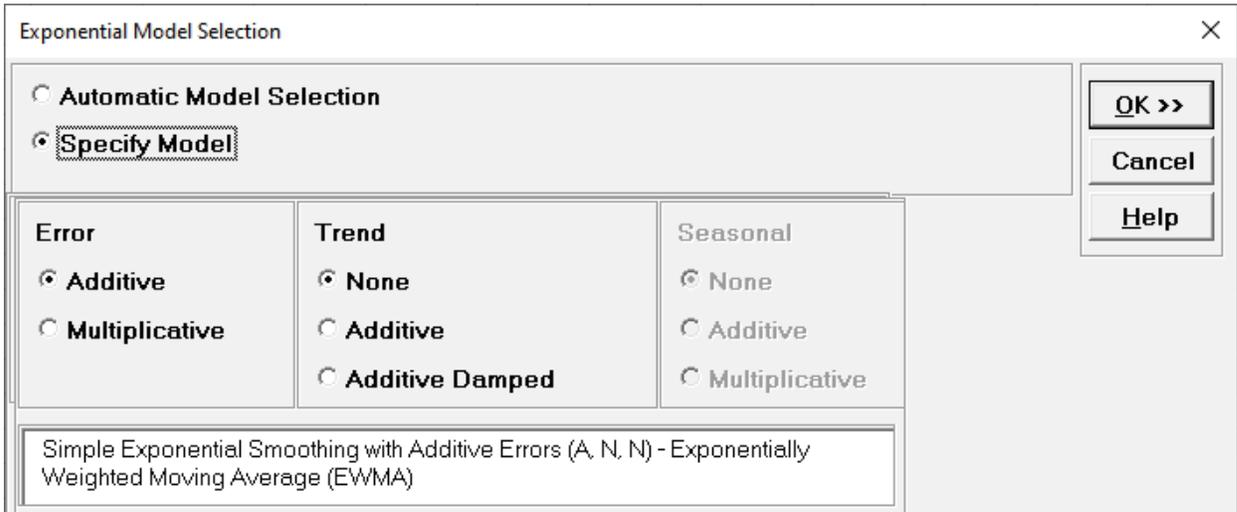
- **Model Options** opens another dialog which allows you to set automatic options or to specify a model.
- **Display ACF/PACF/LB** option will produce ACF and PACF plots for the raw data as well as for the model residuals. The LB plot is a plot of Ljung-Box test P-Values for various lags and is used to determine if a group of autocorrelations are significant, (i.e., the autocorrelations do not come from a white noise series). For further details, see [Ljung-Box Test](#).
- **Display Residual Plots** will produce a table of model residuals and the usual model residual plots: histogram, normal probability plot, residuals versus data order, and residuals vs forecast value. Note that if a Box-Cox transformation is applied, the residuals are transformed so will not be equal to forecast - actual.
- **Specify Model Periods** are used to specify a start period, end period or withhold sample. The withheld data is not used in model estimation, so this is very useful for model validation and comparison. This will be used in a later example.
- **Seasonal Frequency** and **Box-Cox Transformation** will be used in a later example.

14. Click **Model Options**.



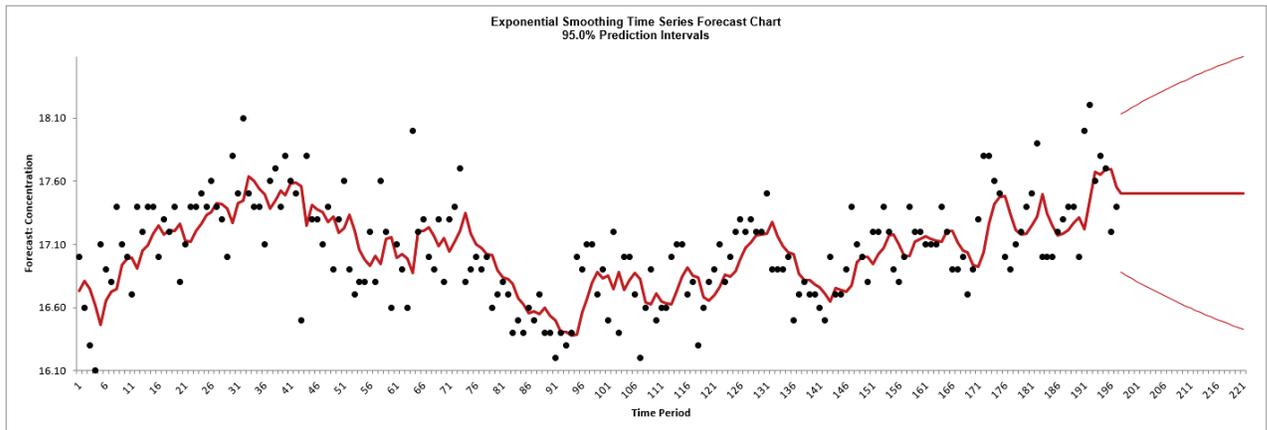
- **Automatic Model Selection** will be used later. It is the default selection.
- **Model Selection Criterion** is the information criterion metric to be used in automatic model selection. AICc is the default selection.
- Clicking **OK** accepts the settings and returns you to the previous dialog. Clicking **Cancel** will cancel any changes and return you to the previous dialog.

15. Select **Specify Model**.



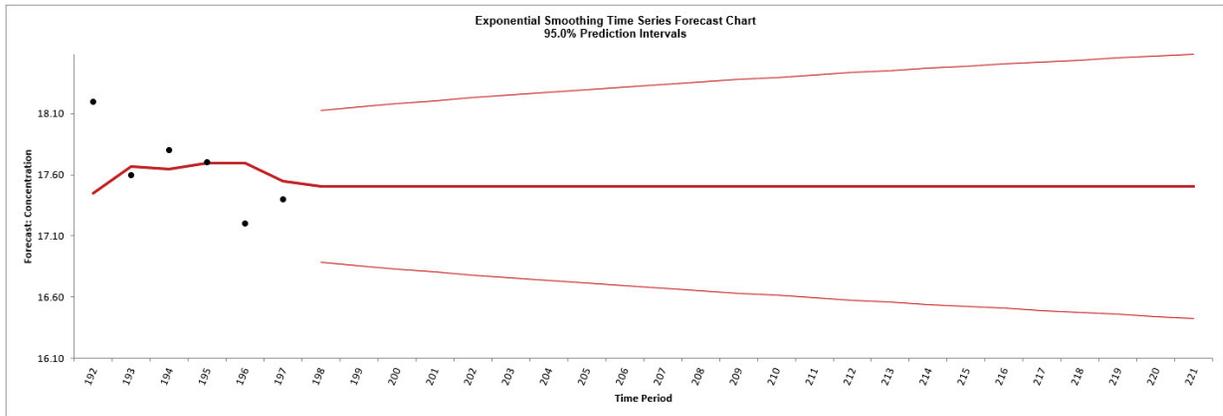
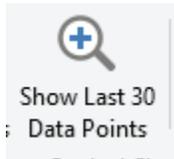
- **Specify Model** allows you to manually specify the Error-Trend-Seasonal model. Seasonal options are greyed out because the **Seasonal Frequency** option in the main dialog is unchecked. A summary description of the model is also given. This will be included in the forecast report.

16. We will use the default **Error: Additive** and **Trend: None**, which is a simple exponential smoothing model, the same as was demonstrated previously. Click **OK** to return to the Exponential Smoothing Forecast dialog. Click **OK**. The exponential smoothing forecast chart is given:

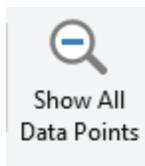


This is very similar to the exponential smooth plot demonstrated above, showing the raw Concentration data (black) and one-step-ahead forecast values (red), but with the addition of a 24-period forecast and the 95% prediction interval.

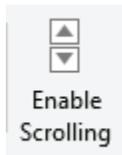
17. You can zoom in to view the last 30 points on the Forecast Chart. Click **SigmaXL Chart Tools > Show Last 30 Data Points**.



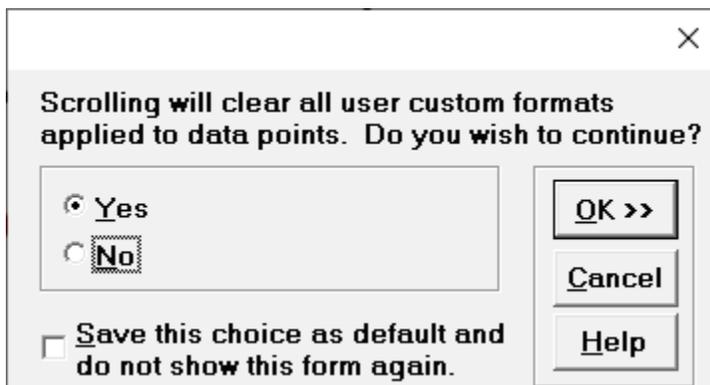
18. To restore the chart, click **SigmaXL Chart Tools > Show All Data Points**.



19. You can also scroll through the data. Click **SigmaXL Chart Tools > Enable Scrolling**

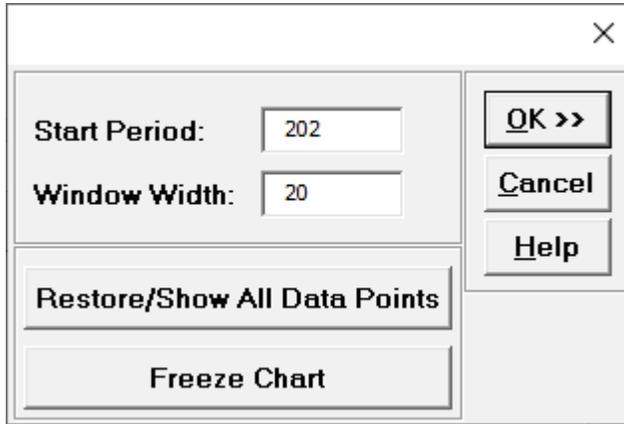


You are prompted with a warning message that custom formatting on the chart will be cleared:



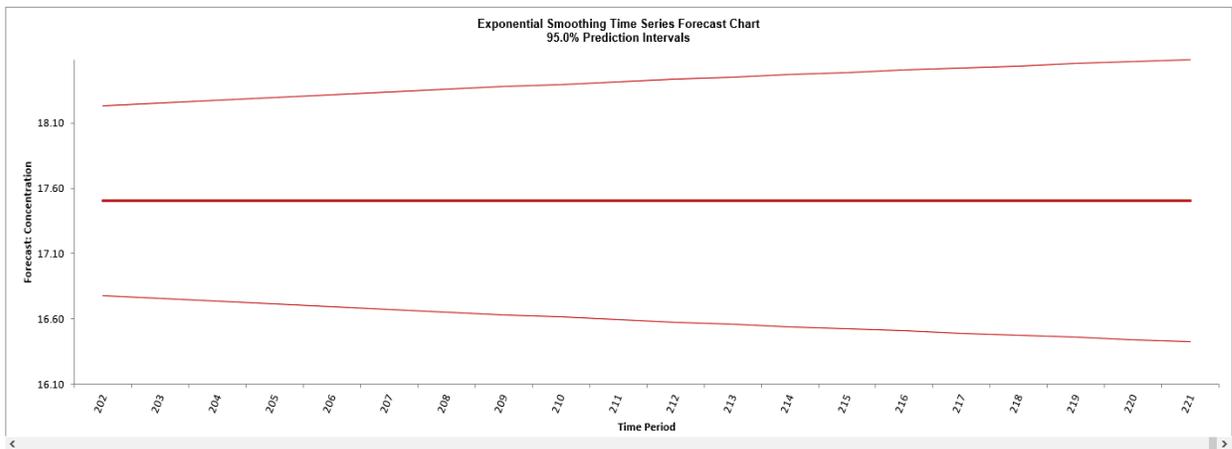
You can avoid seeing this warning by checking **Save this choice as default and do not show this form again**.

20. Click **OK**. The scroll dialog appears allowing you to specify the **Start Period** and **Window Width**:



At any point, you can click **Restore/Show All Data Points** or **Freeze Chart**. Freezing the chart will remove the scroll and unload the dialog. The scroll dialog will also unload if you change worksheets. To restore the dialog, click **SigmaXL Chart Tools > Enable Scrolling**.

21. Click **OK**. A scroll bar appears beneath the forecast chart. You can also change the **Start Subgroup** and **Window Width** and **Update**.



You can scroll through by clicking to the right or left, with the specified window width of 20. Click left once to view the chart as shown.

22. Click **Cancel** to exit the scroll dialog.

23. Scroll down to view the Exponential Smoothing Model header:

Exponential Smoothing Forecast: Simple Exponential Smoothing with Additive Errors (A, N, N) - Exponentially Weighted Moving Average (EWMA) - User Specified Model
Model Periods: All observations are used in the Exponential Smoothing model estimation. No withhold periods available for out-of-sample forecast accuracy evaluation.

The model **Simple Exponential Smoothing with Additive Errors (A, N, N) – Exponentially Weighted Moving Average (EWMA)** is the user specified model. This is the same model information that was displayed in the Model Selection dialog.

If we had checked Specify Model Periods in the main dialog, the start, end or withhold selection would be summarized here as well.

24. The Exponential Smoothing Model Information is given as:

Exponential Smoothing Model Information	
Seasonal Frequency	1
Model Selection Criterion	Specified
Box-Cox Transformation	N/A
Lambda	
Threshold	

This is a summary of model information with Seasonal Frequency = 1 (nonseasonal); Model Selection Criterion = “Specified” because the model was user specified; and Box-Cox Transformation = “N/A” because Box-Cox Transformation was unchecked.

25. The Parameter Estimates are:

Parameter Estimates	
Term	Coefficient
alpha (level smoothing)	0.294785988
I (initial level)	16.73121246

The parameter estimates closely match the values obtained earlier in the demonstration using Solver:

Exponential Smoothing Model Parameters	
alpha	0.29491
initial level	16.7319

The slight differences in parameter results are due to differences in optimization method.

26. The Exponential Smoothing Model Statistics are:

Exponential Smoothing Model Statistics	
No. of Observations	197
DF	195
StDev	0.31819
Variance	0.10124
Log-Likelihood	-293.804
AICc	593.732
AIC	593.607
BIC	603.457

Degrees of freedom (DF) = $n - 2$ (terms in the model). See [Exponential Smoothing Parameter Estimation, Model Statistics and Information Criteria for Model Comparison](#).

The results match those given in the demonstration using Solver:

Exponential Smoothing Model Statistics	
Count	197
k terms	2
Log-Likelihood	-293.8036
AICc	593.73155
AIC	593.6072
BIC	603.45681

The equations may be viewed by clicking on the Demo cells **M5**, **M6**, **M7** or **M8**.

27. The In-Sample Forecast Accuracy metrics are:

Forecast Accuracy		
Metric	In-Sample (Estimation) One-Step-Ahead Forecast	Out-of-Sample (Withhold) One-Step-Ahead Forecast
N	197	
RMSE	0.316569334	
MAE	0.247329038	
MAPE	1.446520183	
MASE	0.897712804	

MASE is less than one, so it is a better forecast than would be obtained from a naïve forecast (set all forecasts to be the value of the last observation). See [Forecast Accuracy Metrics](#). The results closely match those given in the demonstration using Solver:

In-Sample Forecast Accuracy Metrics	
RMSE	0.316569323
MAE	0.247320195
MAPE	1.446469239
Scale (MR-bar)	0.275510204
MASE	0.897680708

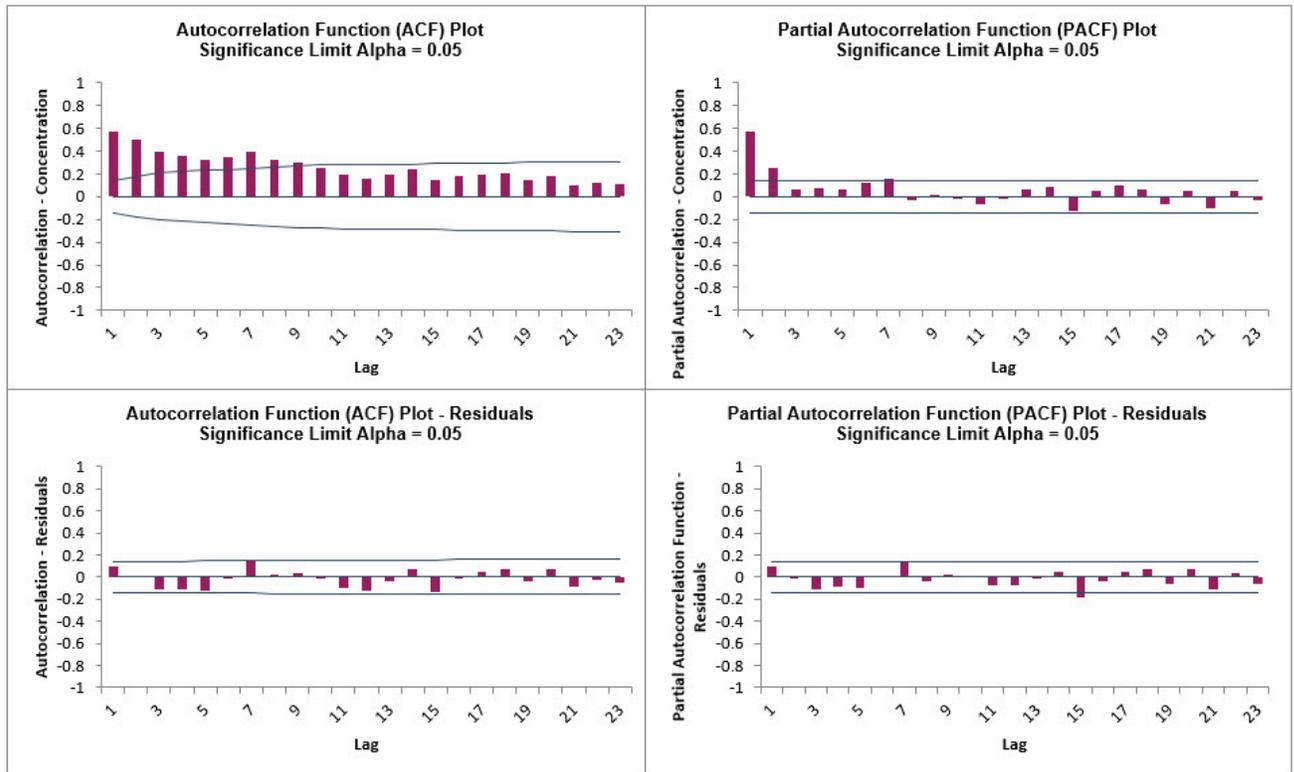
The equations may be viewed by clicking on the Demo cells **P3** to **P7**.

28. The Forecast Table is given as:

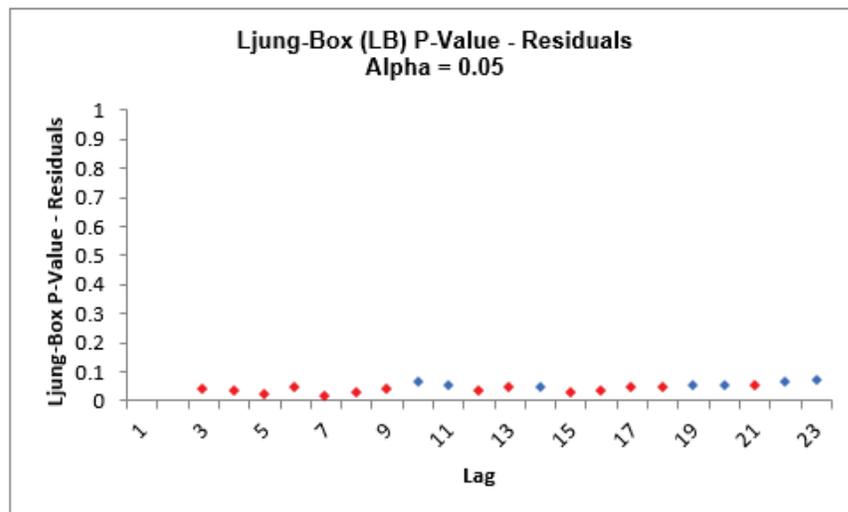
Forecast Table				
Period	Withhold Data	Forecast	Lower 95.0% PI	Upper 95.0% PI
198		17.50540234	16.8817641	18.12904059
199		17.50540234	16.8552318	18.15557289
200		17.50540234	16.82974058	18.1810641
201		17.50540234	16.80517674	18.20562795
202		17.50540234	16.78144587	18.22935881
203		17.50540234	16.75846858	18.25233611
204		17.50540234	16.73617733	18.27462736
205		17.50540234	16.7145141	18.29629058
206		17.50540234	16.69342865	18.31737604
207		17.50540234	16.67287705	18.33792764
208		17.50540234	16.65282071	18.35798398
209		17.50540234	16.63322546	18.37757923

These are the same forecast and prediction interval values displayed in the Forecast Chart but provided for further analysis or charting. If Withhold Periods are specified, the Withhold Data will be displayed as well.

29. Click on the **Exp. Smooth. ACF PACF LB** sheet to view the ACF/PACF/LB Plots:



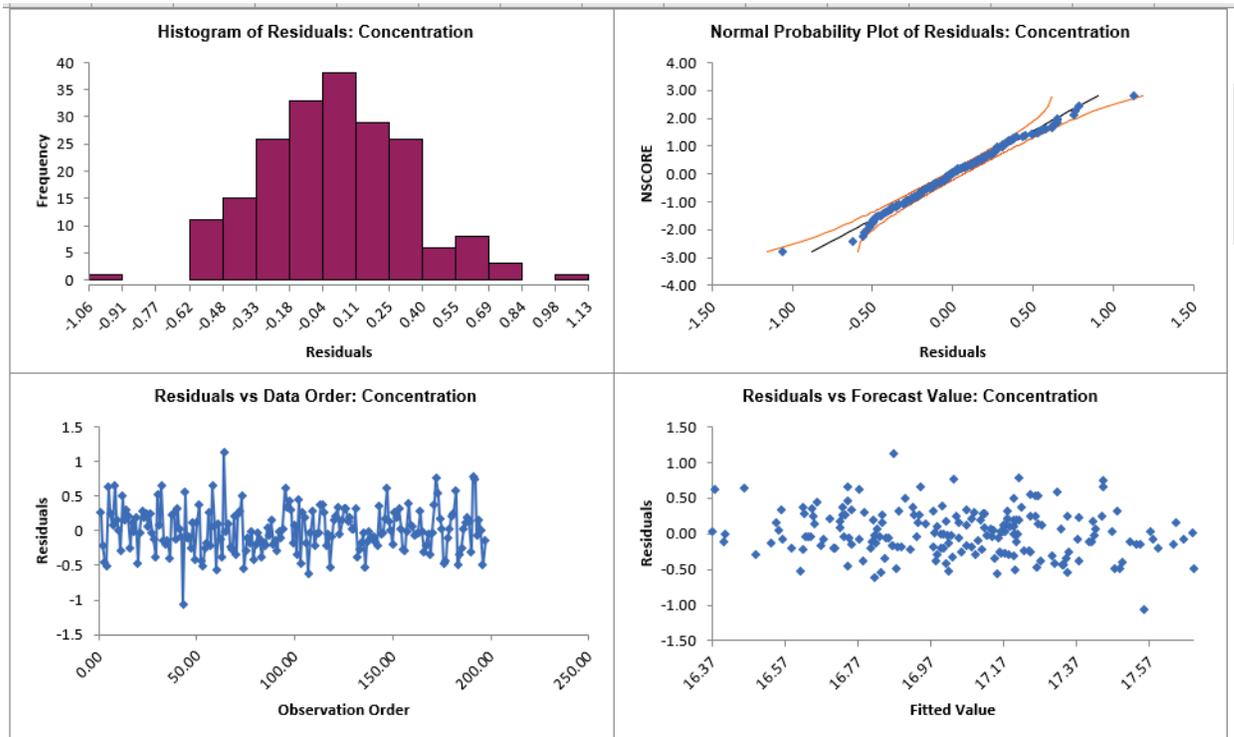
These match the plots that we obtained previously in the [Demo](#). We can see that all of the autocorrelation has been removed by the exponential smoothing model (with the exception of lag 15 in the PACF), so this is a good fit to the time series data.



The LB plot is a plot of Ljung-Box test P-Values for various lags and is used to determine if a group of autocorrelations are significant, (i.e., the autocorrelations do not come from a white noise series). For further details, see [Ljung-Box Test](#).

The red P-Values are significant ($\alpha=.05$) and the blue P-Values are not significant. It is desirable that all P-Values be blue. The ACF/PACF plots indicated that almost all of the correlation has been accounted for in the model, but the Ljung-Box plot shows that some significant autocorrelation still remains - so the model can potentially be improved. This does not mean that the model is a bad model, it can still be very useful for prediction purposes, but the prediction intervals may not provide accurate coverage.

30. Click on the **Exp. Smoothing Residuals** sheet to view the Residual Plots:



These residual plots are the same as used in SigmaXL's Multiple Regression. The histogram matches what we obtained in the Demo (the normal curve is not applied here). The residuals are approximately normally distributed, with a roughly straight line on the normal probability plot. There are no obvious extreme outliers or patterns in the charts. Later, we will apply a control chart to the residuals to formally test for significant outliers or assignable causes.

31. Now we will rerun Exponential Smoothing on the Concentration data but use Automatic Model Selection. Click **Recall SigmaXL Dialog** menu or press **F3** to recall last dialog.

32. Click **Model Options**. Select **Automatic Model Selection**. We will use the default **Model Selection Criterion: AICc – Akaike information criterion with small sample size correction**.

33. Click **OK** to return to the Exponential Smoothing Forecast dialog. Click **OK**. The exponential smoothing forecast report is given.

34. Scroll down to view the Exponential Smoothing Model header:

Exponential Smoothing Forecast: Simple Exponential Smoothing with Multiplicative Errors (M, N, N) - Model Automatically Selected
 Model Periods: All observations are used in the Exponential Smoothing model estimation. No withhold periods available for out-of-sample forecast accuracy evaluation.

The model **Simple Exponential Smoothing with Multiplicative Errors (M, N, N)** was selected as the best fit for the Concentration data based on the AICc criterion. The point forecasts

produced by the Multiplicative and Additive models are identical if they use the same smoothing parameter values. Multiplicative will, however, generate different prediction intervals to accommodate change in variance.

35. The Exponential Smoothing Model Information is given as:

Exponential Smoothing Model Information	
Seasonal Frequency	1
Model Selection Criterion	AICc
Box-Cox Transformation	N/A
Lambda	
Threshold	

This is a summary of model information with Seasonal Frequency = 1 (nonseasonal); Model Selection Criterion = "AICc" and Box-Cox Transformation = "N/A" because Box-Cox Transformation was unchecked.

36. The Parameter Estimates are:

Parameter Estimates	
Term	Coefficient
alpha (level smoothing)	0.295846234
I (initial level)	16.73860868

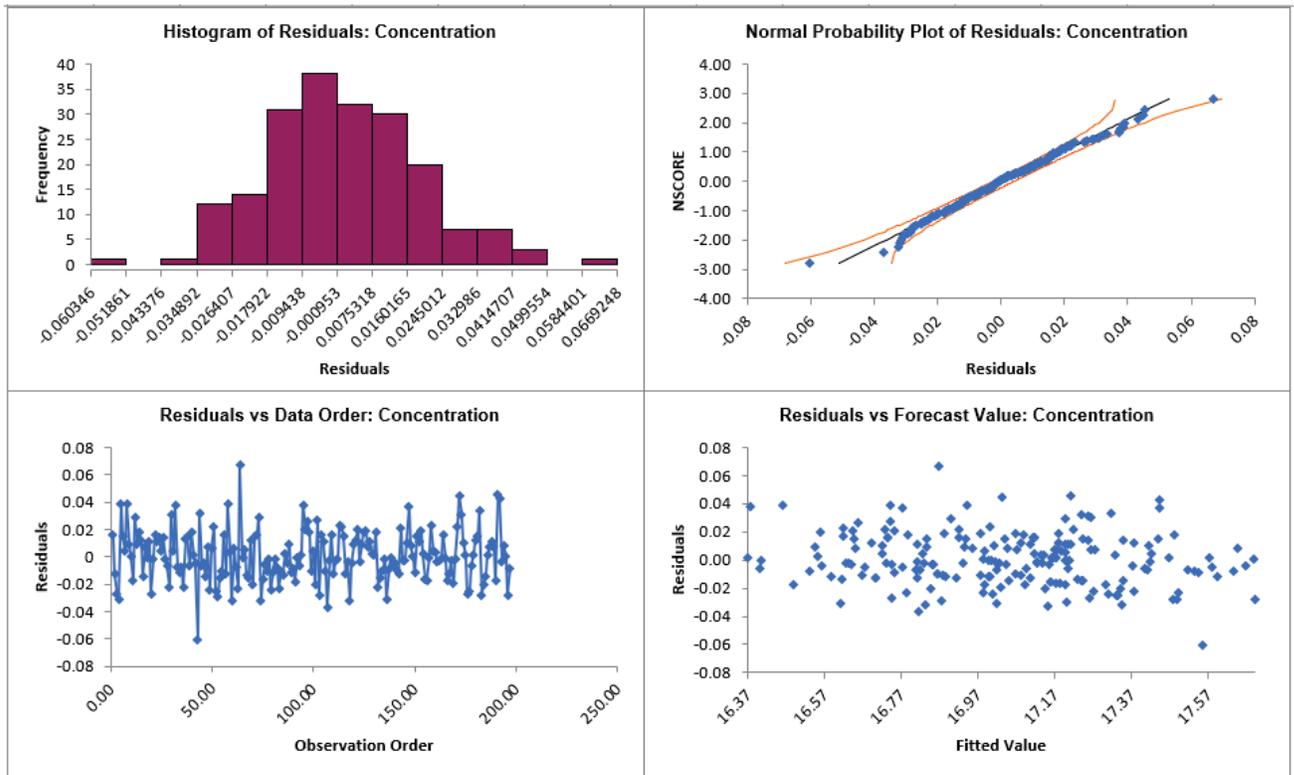
These are fairly close to the parameter estimates obtained [above](#) using the additive model.

37. The Exponential Smoothing Model Statistics are:

Exponential Smoothing Model Statistics	
No. of Observations	197
DF	195
StDev	0.01865
Variance	0.00035
Log-Likelihood	-293.601
AICc	593.326
AIC	593.201
BIC	603.051

The Log-Likelihood, AICc, AIC and BIC are close to the values obtained [above](#) using the additive model, but the Log-Likelihood is slightly higher giving a lower AICc, so multiplicative was selected as the best model. Note however that the StDev and Variance are very different. This difference is due to the multiplicative residuals being relative errors: $e_t = \frac{y_t - \hat{y}_t}{\hat{y}_t}$.

38. The In-Sample Forecast Accuracy metrics, Forecast Table and ACF/PACF/LB Plots for multiplicative are very close to the additive model. The multiplicative residual plots look the same, but note the different scale due to the relative errors.

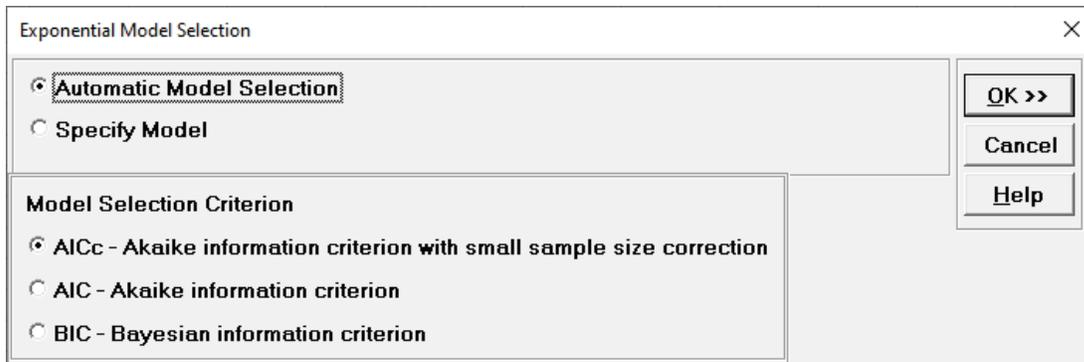


39. Now we will rerun Exponential Smoothing on the Concentration data with Automatic Model Selection, but will use Specify Withhold Periods. Click **Recall SigmaXL Dialog** menu or press **F3** to recall last dialog. Check **Specify Model Periods**. Set **Withhold Periods** = 24 (i.e., we will forecast 24 periods and compare against the withheld actual). We will use the default **Withhold Forecast Type: One-Step-Ahead with Prediction Interval at: Start of Withhold**.

- **Specify Model Periods** option allows you to specify the start and end periods used in automatic model identification and parameter estimation. Typically, **Start Model at Period** is kept = 1 and **Withhold Periods** specifies the number of periods to be withheld for out-of-sample testing. **End Model at Period** specifies the end period, so the withhold sample size would be: total number of observations – end period.
- **Withhold Forecast Type: One-Step-Ahead** will exclude the withhold sample from automatic model identification and parameter estimation, but uses the withhold data to update the predicted one-step ahead forecast. This is useful to assess forecast error when you only care about the short-term one-step ahead prediction.
- **Withhold Forecast Type: One-Step-Ahead with Prediction Interval at: Start of Withhold** will display the prediction interval for the duration of the withhold sample. Note that the length of the prediction interval is determined by the number of withhold periods, so overrides the specified **No. of Forecast Periods**.

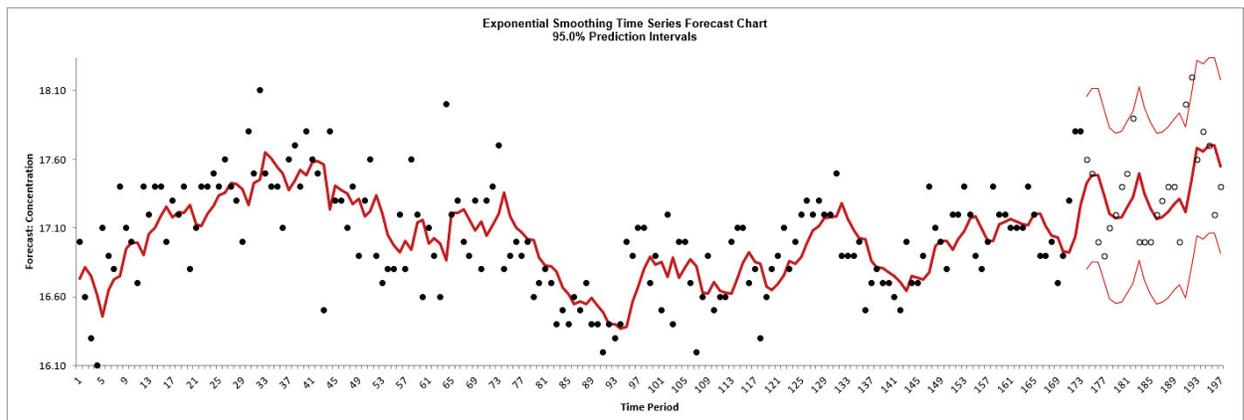
- **Withhold Forecast Type: One-Step-Ahead with Prediction Interval at: End of Withhold** will display the prediction interval at the end of the withhold sample. The length of the prediction interval is determined by the specified **No. of Forecast Periods**.
- **Include in Residuals** will treat the one-step-ahead forecast errors as residuals (even though they were not part of the model estimation process) and will be included in the ACF/PACF/LB Residual Plots along with the Residuals report and graphs. Typically, this is kept unchecked.
- **Withhold Forecast Type: Multi-Step-Ahead with Prediction Interval at Start of Withhold** will exclude the withhold sample from automatic model identification and parameter estimation and does not use the withhold data to update the predicted one-step ahead forecast. This is useful to assess forecast error when you are interested in a long-term forecast window (horizon). The prediction interval will be displayed for the duration of the withhold sample. Note that the length of the prediction interval is determined by the number of withhold periods, so overrides the specified **No. of Forecast Periods**. These forecast errors are not included in ACF/PACF/LB Residual Plots or the Residuals report and graphs.

40. Click **Model Options**. Select **Automatic Model Selection**. We will use the default **Model Selection Criterion: AICc – Akaike information criterion with small sample size correction**.



Tip: When using **Recall SigmaXL Dialog**, and if there are no changes to the **Model Option** settings, the previous settings will be used. It is not necessary to repeat this step.

41. Click **OK** to return to the Exponential Smoothing Forecast dialog. Click **OK**. The exponential smoothing forecast report is given:



The blank dots are the data values in the withhold sample with a one-step-ahead forecast and prediction intervals displayed at the start of the withhold sample.

42. Scroll down to view the Exponential Smoothing Model header:

Exponential Smoothing Forecast: Simple Exponential Smoothing with Multiplicative Errors (M, N, N) - Model Automatically Selected
Model Periods: Model parameter estimates calculated excluding 24 withhold periods.

As with the complete data, the model **Simple Exponential Smoothing with Multiplicative Errors (M, N, N)** was selected as the best fit for the Concentration data based on the AICc criterion. The header also includes the number of specified withhold periods.

43. The Parameter Estimates are:

Parameter Estimates	
Term	Coefficient
alpha (level smoothing)	0.303967286
l (initial level)	16.73554259

These are close to the parameter estimates obtained [above](#) (which used all of the data with the multiplicative model).

44. The Exponential Smoothing Model Statistics are:

Exponential Smoothing Model Statistics	
No. of Observations	173
DF	171
StDev	0.01841
Variance	0.00034
Log-Likelihood	-243.805
AICc	493.752
AIC	493.61
BIC	503.07

These are fairly close to the model statistics obtained [above](#) (which used all of the data with the multiplicative model). Here we are using only 173 of the 197 observations.

45. The Forecast Accuracy metrics are:

Forecast Accuracy		
Metric	In-Sample (Estimation) One-Step-Ahead Forecast	Out-of-Sample (Withhold) One-Step-Ahead Forecast
N	173	24
RMSE	0.31133921	0.35209899
MAE	0.243032918	0.273389284
MAPE	1.425751706	1.567334268
MASE	0.878186174	0.987877246

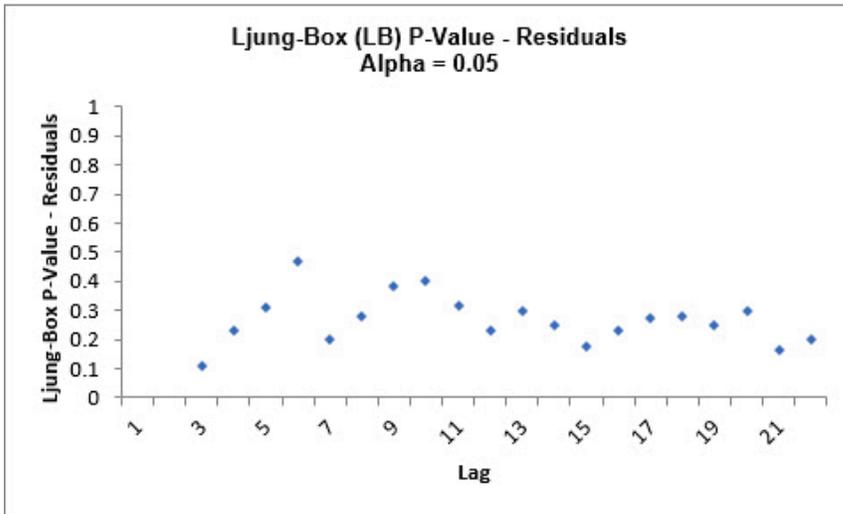
As expected, the **Out-of-Sample (Withhold) One-Step-Ahead Forecast** errors are larger than the **In-Sample (Estimation) One-Step-Ahead Forecast** errors.

46. The Forecast Table is given as:

Forecast Table				
Period	Withhold Data	One-Step-Ahead Forecast	Lower 95.0% PI	Upper 95.0% PI
174	17.6	17.43034504	16.80141557	18.05927452
175	17.5	17.4819146	16.85112437	18.11270483
176	17	17.48741197	16.85642338	18.11840056
177	16.9	17.33925468	16.71361197	17.96489738
178	17.1	17.20573562	16.58491061	17.82656064
179	17.2	17.17359545	16.55393014	17.79326077
180	17.4	17.18162157	16.56166665	17.80157649
181	17.5	17.24800147	16.6256514	17.87035154
182	17.9	17.32460078	16.69948682	17.94971474
183	17	17.49950332	16.86807845	18.13092819
184	17	17.34767065	16.72172427	17.97361703
185	17	17.24199015	16.61985698	17.86412331
186	17.2	17.16843306	16.54895401	17.78791211
187	17.3	17.17802838	16.55820311	17.79785365
188	17.4	17.21510376	16.59394072	17.8362668
189	17.4	17.27130617	16.64811521	17.89449713
190	17	17.31042488	16.68582242	17.93502734
191	18	17.21606587	16.59486812	17.83726363
192	18.2	17.4543562	16.82456035	18.08415206
193	17.6	17.68100752	17.04303354	18.31898151
194	17.8	17.65638389	17.01929838	18.29346939
195	17.7	17.70003849	17.06137782	18.33869916
196	17.2	17.70002679	17.06136654	18.33868704
197	17.4	17.548035	16.91485899	18.18121102

These are the same forecast and prediction interval values displayed in the Forecast Chart, but provided for further analysis or charting. The **Withhold Data** is also displayed.

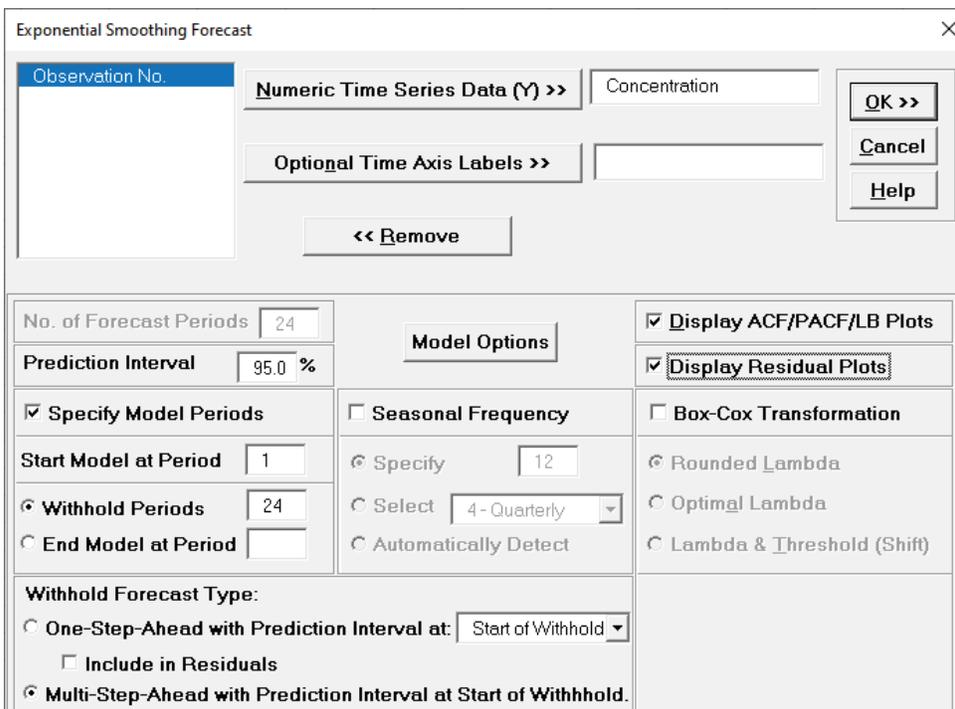
47. The ACF/PACF/LB Residual Plots and Residual Plots are based on the in-sample data. The plots look similar to the complete data [above](#), except for the Ljung-Box P-Values:



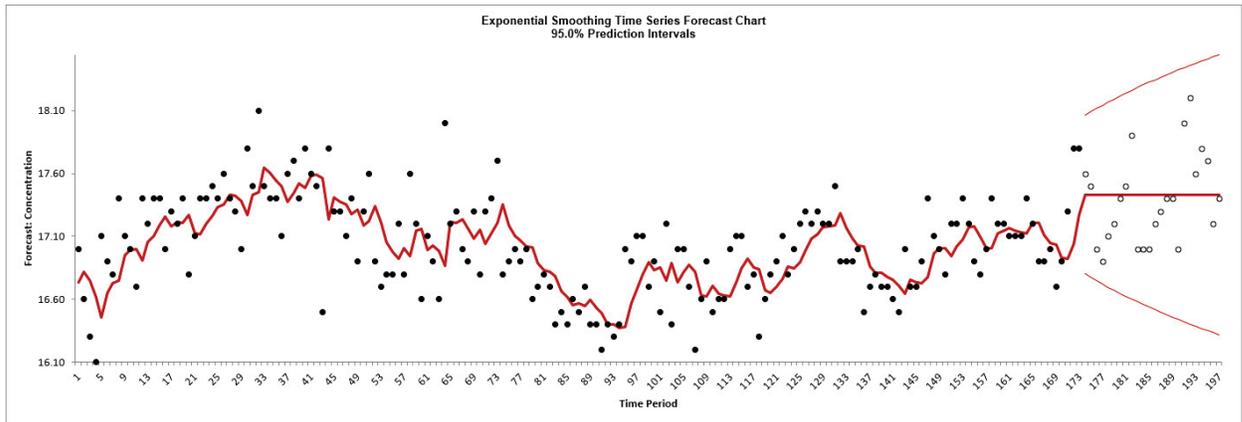
The **Simple Exponential Smoothing with Multiplicative Errors (M, N, N)** model is a better fit to the subset than the complete data, with all P-Values being blue (> .05).

48. If **Include in Residuals** was checked then the residuals would also include the **Out-of-Sample (Withhold) One-Step-Ahead Forecast** errors.

49. Now we will rerun Exponential Smoothing on the Concentration data, but use Multi-Step-Ahead for the Withhold Forecast. Click **Recall SigmaXL Dialog** menu or press **F3** to recall last dialog. Select **Withhold Forecast Type: Multi-Step-Ahead with Prediction Interval at Start of Withhold**.



50. Click **OK**. The exponential smoothing forecast report is given:



The blank dots are the data values in the withhold sample with a multi-step forecast and prediction intervals displayed at the start of the withhold sample.

51. The Forecast Accuracy metrics are:

Forecast Accuracy		
Metric	In-Sample (Estimation) One-Step-Ahead Forecast	Out-of-Sample (Withhold) Multi-Step-Ahead Forecast
N	173	24
RMSE	0.31133921	0.348228059
MAE	0.243032918	0.286752928
MAPE	1.425751706	1.647391419
MASE	0.878186174	1.036166041

As expected, the **Out-of-Sample (Withhold) Multi-Step-Ahead Forecast** errors are larger than the **In-Sample (Estimation) One-Step-Ahead Forecast** errors and the **Out-of-Sample (Withhold) One-Step-Ahead Forecast** errors [above](#).

Monthly Airline Passengers - Series G

52. Open **Monthly Airline Passengers - Series G.xlsx (Sheet 1 tab)**. This is the Series G data from Box and Jenkins, monthly total international airline passengers for January 1949 to December 1960. See the [Run Chart](#), [ACF/PACF Plots](#), [Spectral Density Plot](#) and [Seasonal Trend Decomposition Plots](#) for this data.
53. Click **SigmaXL > Time Series Forecasting > Exponential Smoothing Forecast > Forecast**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
54. Select *Monthly Airline Passengers*, click **Numeric Time Series Data (Y) >>**. Select *Date*, click **Optional Time Axis Labels >>**. Check **Display ACF/PACF/LB Plots** and **Display Residual Plots**. Check **Specify Model Periods**. Set **Withhold Periods** = 24 (i.e., we will forecast 24 months and compare to withheld actual). Select **Withhold Forecast Type: Multi-Step-Ahead with Prediction Interval at Start of Withhold**. Check **Seasonal Frequency** with **Specify** = 12. Check **Box-Cox Transformation** and select **Rounded Lambda** (selected because the [Run Chart](#) showed an increase in the seasonal variance over time). We will use the default **Prediction Interval** = 95.0 %.

Exponential Smoothing Forecast

Obs. No.
Ln (Airline Passengers)

Numeric Time Series Data (Y) >> Monthly Airline Passenger

Optional Time Axis Labels >> Date

<< Remove

OK >>
Cancel
Help

No. of Forecast Periods 24

Prediction Interval 95.0 %

Model Options

Specify Model Periods

Start Model at Period 1

Withhold Periods 24

End Model at Period

Seasonal Frequency

Specify 12

Select 4-Quarterly

Automatically Detect

Display ACF/PACF/LB Plots

Display Residual Plots

Box-Cox Transformation

Rounded Lambda

Optimal Lambda

Lambda & Threshold (Shift)

Withhold Forecast Type:

One-Step-Ahead with Prediction Interval at: Start of Withhold

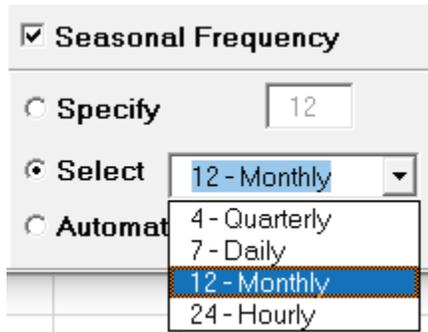
Include in Residuals

Multi-Step-Ahead with Prediction Interval at Start of Withhold.

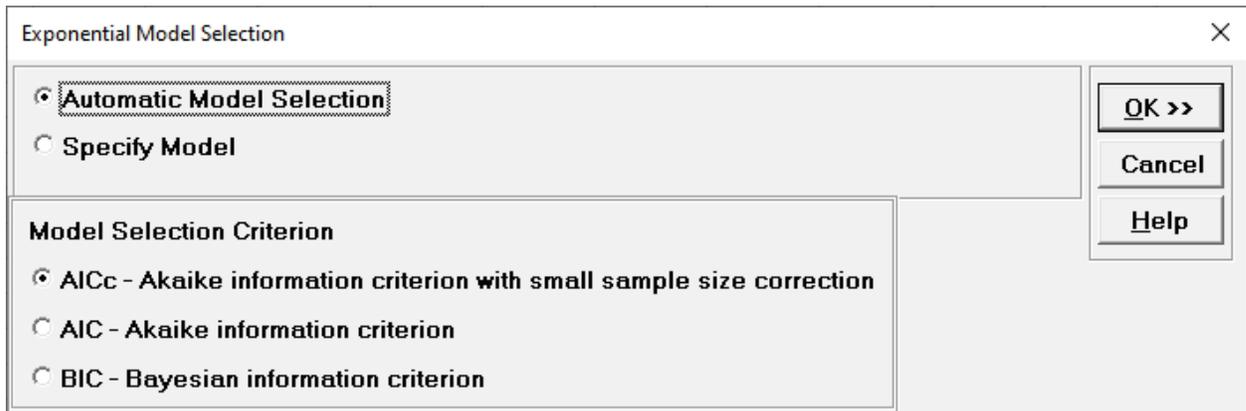
- **Withhold Forecast Type: Multi-Step-Ahead with Prediction Interval at Start of Withhold** will exclude the withhold sample from automatic model identification and parameter estimation and does not use the withhold data to update the predicted one-step ahead forecast. This is

useful to assess forecast error when you are interested in a long-term forecast window (horizon). The prediction interval will be displayed for the duration of the withhold sample. Note that the length of the prediction interval is determined by the number of withhold periods, so overrides the specified **No. of Forecast Periods**.

- **Seasonal Frequency Specify** is used to specify the seasonal frequency. Note that Exponential Smoothing is limited to a maximum seasonal frequency of 24. For higher frequencies use Exponential Smoothing – Multiple Seasonal Decomposition (MSD)
- **Seasonal Frequency Select** gives a drop-down list of commonly used seasonal frequencies:

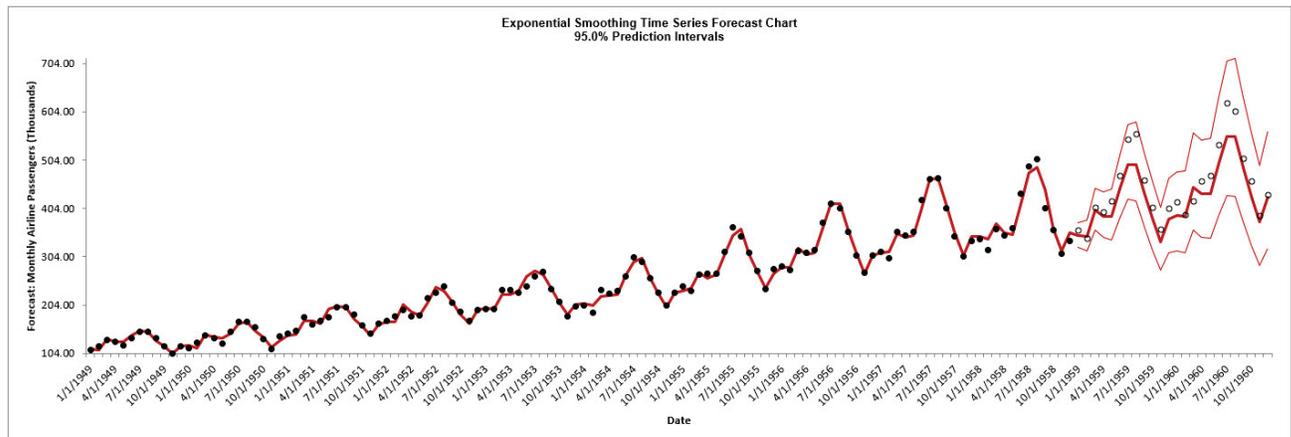


- **Seasonal Frequency Automatically Detect** should be used if uncertain what the seasonal frequency value is (or do a Spectral Density Plot prior to the Seasonal Trend Decomposition Plots). If detected frequency is > 24, 1 is returned. If that occurs, please use Exponential Smoothing – Multiple Seasonal Decomposition (MSD).
- **Box-Cox Transformation with Rounded Lambda** will select Lambda = 0 (Ln), 0.5 (SQRT) or 1 (Untransformed). Threshold (Shift) is computed automatically if the time series data includes 0 or negative values, otherwise it is 0.
- **Box-Cox Transformation with Optimal Lambda** uses the range of 0 to 1 for Lambda. Threshold is computed automatically if the time series data includes 0 or negative values.
- **Box-Cox Transformation with Lambda & Threshold (Shift)** if left blank, will compute optimal lambda and threshold. The user may also specify Lambda and Threshold. Lambda may be specified outside of the 0 to 1 range, but practically for time series analysis, should be limited to -1 to 2. Threshold is typically 0, but if the time series data includes 0 or negative values, a negative threshold value should be entered that is smaller than the minimum data value. This value will be subtracted from the data resulting in positive time series values.

55. Click **Model Options**.


The dialog box titled "Exponential Model Selection" has a close button (X) in the top right corner. It contains two main sections. The first section has two radio buttons: "Automatic Model Selection" (which is selected) and "Specify Model". The second section, titled "Model Selection Criterion", has three radio buttons: "AICc - Akaike information criterion with small sample size correction" (selected), "AIC - Akaike information criterion", and "BIC - Bayesian information criterion". On the right side of the dialog, there are three buttons: "OK >>", "Cancel", and "Help".

56. We will use the default **Automatic Model Selection** with **AICc** as the **Model Selection Criterion**. Click **OK** to return to the Exponential Smoothing Forecast dialog. Click **OK**. The exponential smoothing forecast report is given:



The blank dots are the data values in the withhold sample with a multi-step forecast and prediction intervals displayed at the start of the withhold sample. Note that the data has been transformed using Box-Cox, but the inverse transformation is applied to produce this chart with the original units.

57. Scroll down to view the Exponential Smoothing Model header:

Exponential Smoothing Model: Additive Trend, Additive Seasonal Method with Additive Errors (Holt-Winters) (A, A, A) - Model Automatically Selected
Model Periods: Model parameter estimates calculated excluding 24 withhold periods.

The model **Additive Trend, Additive Seasonal Method with Additive Errors (Holt-Winters) (A, A, A)** was automatically selected as the best fit for the Box-Cox transformed Airline Passenger data based on the AICc criterion. Note that multiplicative models are not considered when **Box-Cox Transformation** is checked. If a multiplicative model was desired, then one could manually transform the data and model $\ln(\text{Airline Passengers})$ with **Box-Cox Transformation** unchecked.

The header also includes the number of specified withhold periods.

58. The Exponential Smoothing Model Information is given as:

Exponential Smoothing Model Information	
Seasonal Frequency	12
Model Selection Criterion	AICc
Box-Cox Transformation	Rounded Lambda
Lambda	0
Threshold	0

This is a summary of model information with Seasonal Frequency = 12; Model Selection Criterion = “AICc” and Box-Cox Transformation = “Rounded Lambda” with Lambda = 0 (i.e., a Ln transformation).

59. The Parameter Estimates are:

Parameter Estimates	
Term	Coefficient
alpha (level smoothing)	0.770415301
beta (trend smoothing)	0.0001
gamma (seasonal smoothing)	0.0001
l (initial level)	4.80925762
b (initial trend)	0.00935261
s1 (initial seasonal)	-0.099989196
s2 (initial seasonal)	-0.217929483
s3 (initial seasonal)	-0.076652841
s4 (initial seasonal)	0.063787774
s5 (initial seasonal)	0.198157807
s6 (initial seasonal)	0.205968013
s7 (initial seasonal)	0.108472035
s8 (initial seasonal)	-0.016634733
s9 (initial seasonal)	-0.008278553
s10 (initial seasonal)	0.033611999
s11 (initial seasonal)	-0.101347496
s12 (initial seasonal)	-0.089165324

- Error includes the smoothing parameter *alpha* and initial level value (*l*). The error is additive, but on the Ln transformed data.
- Trend adds a smoothing parameter (*beta*) and initial trend value (*b*).
- Seasonal adds a smoothing parameter (*gamma*) and initial seasonal values (*s1* to *s12*), constrained to sum to zero for additive.

60. The Exponential Smoothing Model Statistics are:

Exponential Smoothing Model Statistics	
No. of Observations	120
DF	104
StDev	0.037144
Variance	0.00138
Log-Likelihood	116.4916
AICc	-192.983
AIC	-198.983
BIC	-151.596

- The number of observations, $n = 144 - 24$ (*withhold*) = 120
- Degrees of freedom (DF) = 120 (n) – 16 (terms in the model, excluding s_{12}) = 104
- Note that the model statistics are based on the Ln transformed data, not the original data.

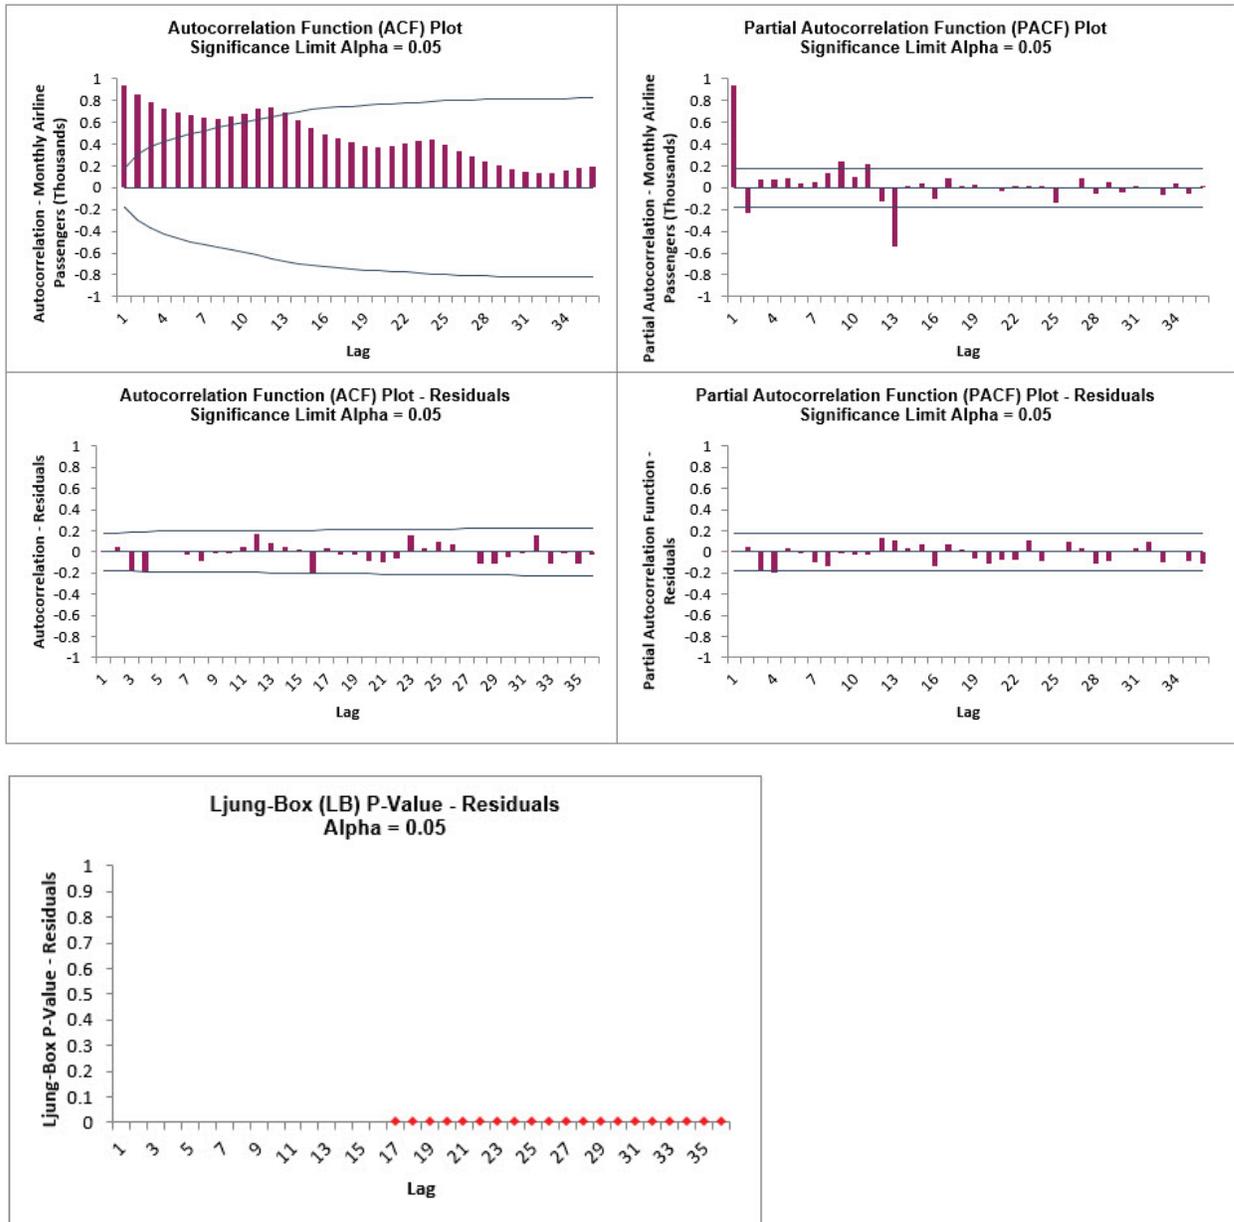
61. The Forecast Accuracy metrics are:

Forecast Accuracy		
Metric	In-Sample (Estimation) One-Step-Ahead Forecast	Out-of-Sample (Withhold) Multi-Step-Ahead Forecast
N	120	24
RMSE	8.67099231	33.06342473
MAE	6.481327147	27.83270693
MAPE	2.730421504	5.80549634
MASE	0.226825448	0.974054552

As expected, the **Out-of-Sample (Withhold) Multi-Step-Ahead Forecast** errors are larger than the **In-Sample (Estimation) One-Step-Ahead Forecast** errors. Note, if we were primarily interested in a short term one-step ahead forecast, then we would have selected **Withhold Forecast Type: One-Step-Ahead** and this table would show **Out-of-Sample (Withhold) One-Step-Ahead Forecast** errors.

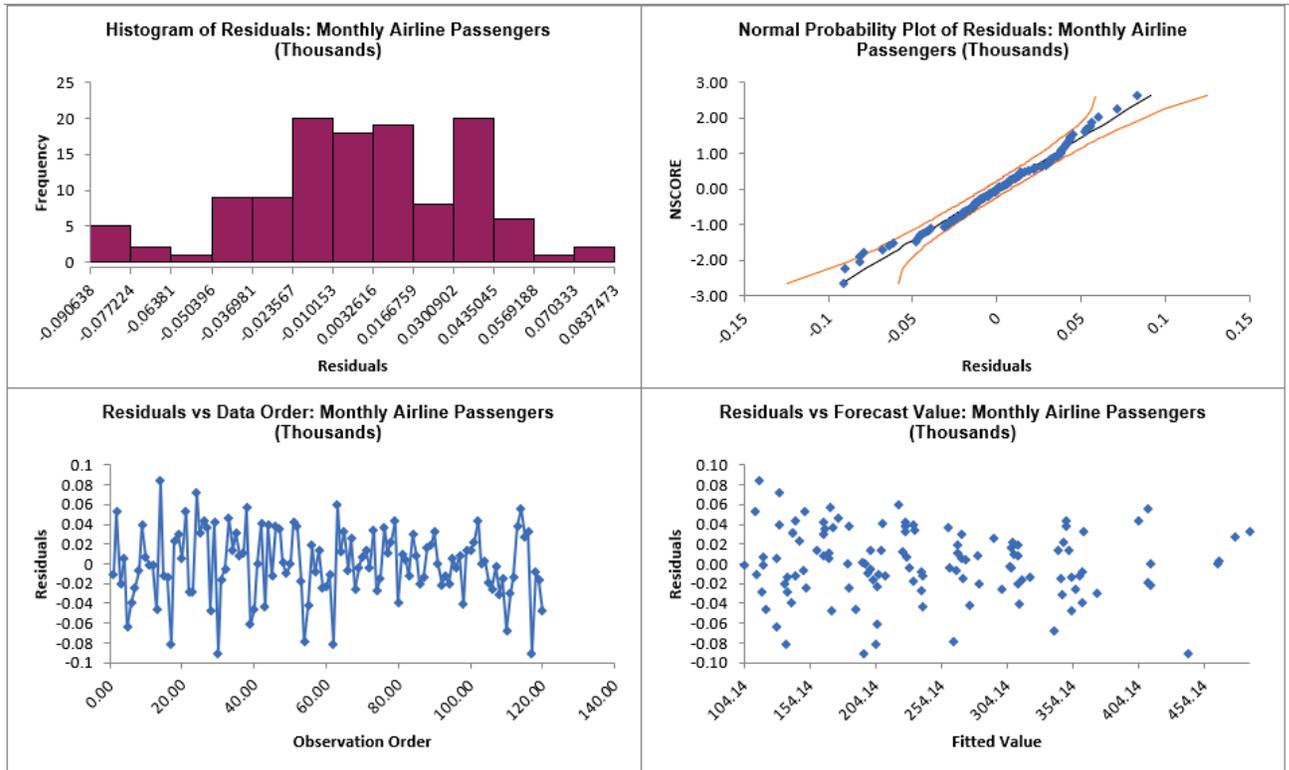
Forecast Accuracy metrics are calculated using the actual raw data versus inverse transformed forecast as displayed in the Forecast Chart and Table, so allows comparison across all model types and transformations.

62. Click on the **Exp. Smooth. ACF PACF LB** sheet to view the ACF/PACF/LB Plots:



The ACF/PACF Residuals Plots indicate that almost all of the autocorrelation has been accounted for in the model, but the Ljung-Box plot shows that significant autocorrelation still remains (the red P-Values are significant at $\alpha=0.05$) - so the model can potentially be improved. This does not mean that the model is a bad model, it can still be very useful for prediction purposes, but the prediction intervals may not provide accurate coverage.

63. Click on the **Exp. Smoothing Residuals** sheet to view the Residual Plots:



The residuals are approximately normally distributed, with a roughly straight line on the normal probability plot. There are no obvious extreme outliers or patterns in the charts.

Note that the residuals are based on the Ln transformed data, not the original data. The **Exponential Smoothing Model Information** to the right of the plots shows the Box-Cox Transformation information.

Exponential Smoothing Model Information	
Seasonal Frequency	12
Model Selection Criterion	AICc
Box-Cox Transformation	Rounded Lambda
Lambda	0
Threshold	0

Exponential Smoothing – Multiple Seasonal Decomposition (MSD) **Forecast**

Exponential Smoothing is limited to a maximum seasonal frequency of 24. For higher frequencies use Exponential Smoothing – Multiple Seasonal Decomposition (MSD). The seasonal component is first removed through decomposition, a nonseasonal exponential smooth model fitted to the remainder (+trend), and then the seasonal component is added back in. For forecasting, a naïve seasonal forecast is used on the seasonal component. Note that the prediction intervals are derived from the exponential smoothing model and do not include uncertainty in the seasonal component.

As the name implies, Multiple Seasonal Decomposition (MSD) also accommodates multiple seasonality, for example the half-hourly data with a seasonal frequency of 48 observations per day and 336 observations per week. When using MSD, it is recommended to limit the forecast period to 2*dominant seasonal frequency.

For further details and formulas, see Appendix: [Seasonal Trend Decomposition](#).

Monthly Airline Passengers - Series G

1. Open **Monthly Airline Passengers - Series G.xlsx (Sheet 1 tab)**. This is the Series G data from Box and Jenkins, monthly total international airline passengers for January 1949 to December 1960. See the [Run Chart](#), [ACF/PACF Plots](#), [Spectral Density Plot](#) and [Seasonal Trend Decomposition Plots](#) for this data. The Multiple Seasonal Decomposition (MSD) option is not necessary for this data, but by way of introduction, we will use this to compare to the previous analysis.
2. Click **SigmaXL > Time Series Forecasting > Exponential Smoothing Forecast > Multiple Seasonal Decomposition Forecast**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
3. Select *Monthly Airline Passengers*, click **Numeric Time Series Data (Y) >>**. Select *Date*, click **Optional Time Axis Labels >>**. Check **Display ACF/PACF/LB Plots** and **Display Residual Plots**. Check **Specify Model Periods**. Set **Withhold Periods = 24**. Select **Withhold Forecast Type: Multi-Step-Ahead with Prediction Interval at Start of Withhold**. Select **Seasonal Frequency Specify** and enter 12. Check **Box-Cox Transformation** and select **Rounded Lambda**. We will use the default **Prediction Interval = 95.0 %**.

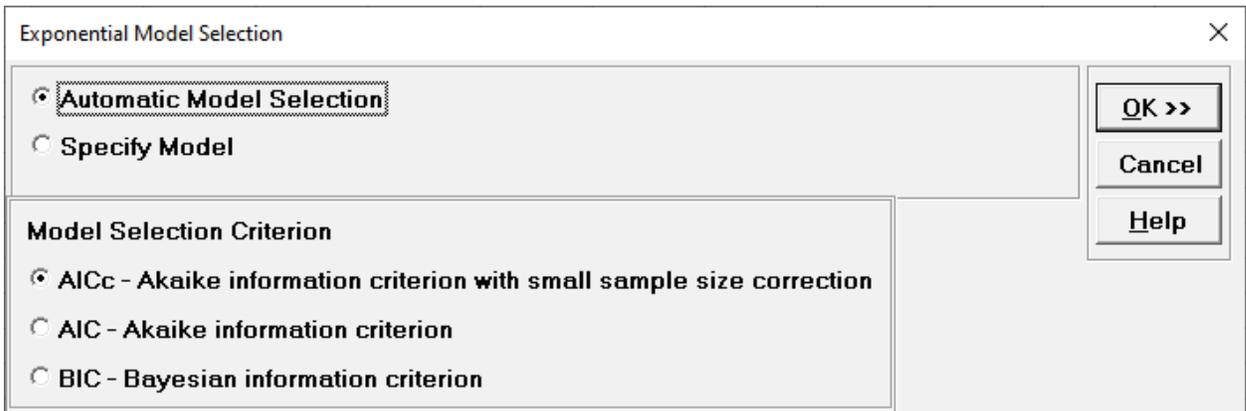
Exponential Smoothing MSD Forecast

Obs. No. Ln (Airline Passengers)	Numeric Time Series Data (Y) >>	Monthly Airline Passenger	OK >>
	Optional Time Axis Labels >>	Date	Cancel
	<< Remove		Help

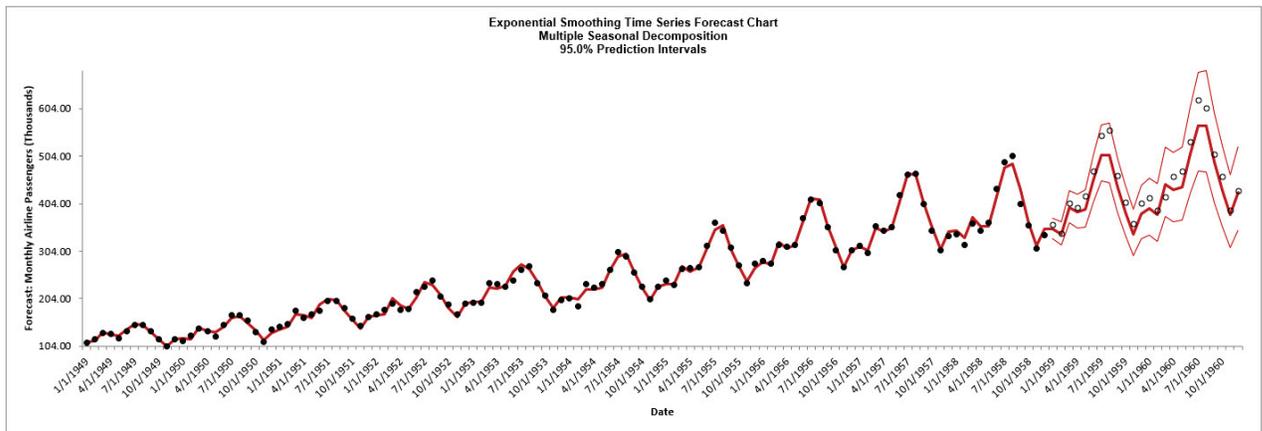
No. of Forecast Periods <input type="text" value="24"/>	Model Options	<input checked="" type="checkbox"/> Display ACF/PACF/LB Plots
Prediction Interval <input type="text" value="95.0 %"/>		<input checked="" type="checkbox"/> Display Residual Plots
<input checked="" type="checkbox"/> Specify Model Periods	Seasonal Frequency	<input checked="" type="checkbox"/> Box-Cox Transformation
Start Model at Period <input type="text" value="1"/>	<input checked="" type="radio"/> Specify <input type="text" value="12"/>	<input checked="" type="radio"/> Rounded Lambda
<input checked="" type="radio"/> Withhold Periods <input type="text" value="24"/>	<input type="radio"/> Select <input type="text" value="4 - Quarterly"/>	<input type="radio"/> Optimal Lambda
<input type="radio"/> End Model at Period <input type="text" value=""/>	<input type="radio"/> Automatically Detect	<input type="radio"/> Lambda & Threshold (Shift)
Withhold Forecast Type:		
<input type="radio"/> One-Step-Ahead with Prediction Interval at: <input type="text" value="Start of Withhold"/>		
<input type="checkbox"/> Include in Residuals		
<input checked="" type="radio"/> Multi-Step-Ahead with Prediction Interval at Start of Withhold.		

- **Seasonal Frequency** can have multiple entries however, we recommend no more than 3 values.

4. Click **Model Options**.



5. We will use the default **Automatic Model Selection** with **AICc** as the **Model Selection Criterion**. Click **OK** to return to the Exponential Smoothing MSD Forecast dialog. Click **OK**. The exponential smoothing forecast report is given:



6. Scroll down to view the Exponential Smoothing Model header:

Exponential Smoothing Model (Multiple Seasonal Decomposition): Additive Trend Method with Additive Errors (Holt's Linear) (A, A, N) - Model Automatically Selected
Model Periods: Model parameter estimates calculated excluding 24 withhold periods.

After **Multiple Seasonal Decomposition**, the model is **Additive Trend Method with Additive Errors (Holt's Linear) (A, A, N)**.

(A, A, N) was automatically selected as the best fit for the deseasonalized Airline Passenger data based on the AICc criterion.

The header also includes the number of specified withhold periods.

7. The Exponential Smoothing Model Information is given as:

Exponential Smoothing Model Information	
Seasonal Frequency	12 Decomposition
Model Selection Criterion	AICc
Box-Cox Transformation	Rounded Lambda
Lambda	0
Threshold	0

This is a summary of model information with Seasonal Frequency = 12 using Decomposition and Model Selection Criterion = "AICc". The Box-Cox Transformation is "Rounded Lambda" with Lambda = 0 (Ln transformation).

8. The Parameter Estimates for the deseasonalized Airline Passenger data are:

Parameter Estimates	
Term	Coefficient
alpha (level smoothing)	0.674923407
beta (trend smoothing)	0.0001
l (initial level)	4.806932238
b (initial trend)	0.009425139

- Error includes the smoothing parameter *alpha* and initial level value (*l*). The error is additive, but on the Ln transformed data.
- Trend adds a smoothing parameter (*beta*) and initial trend value (*b*).
- Seasonal smoothing parameter (*gamma*) and initial seasonal values are not computed because the data has been deseasonalized.

9. The Exponential Smoothing Model Statistics are:

Exponential Smoothing Model Statistics	
No. of Observations	120
DF	116
StDev	0.030549511
Variance	0.000933273
Log-Likelihood	133.3933791
AICc	-256.2604423
AIC	-256.7867581
BIC	-242.8492994

- The number of observations, $n = 144 - 24$ (*withhold*) = 120
- Degrees of freedom (DF) = 120 (n) - 4 (terms in the model) = 116
- Note that the model statistics are based on the Ln transformed data, not the original data.

- The model statistics are better than the [previous analysis](#): lower StDev and Variance, higher Log-Likelihood and lower AICc, AIC and BIC. This is to be expected because the data has been deseasonalized so the seasonal error component is not included.

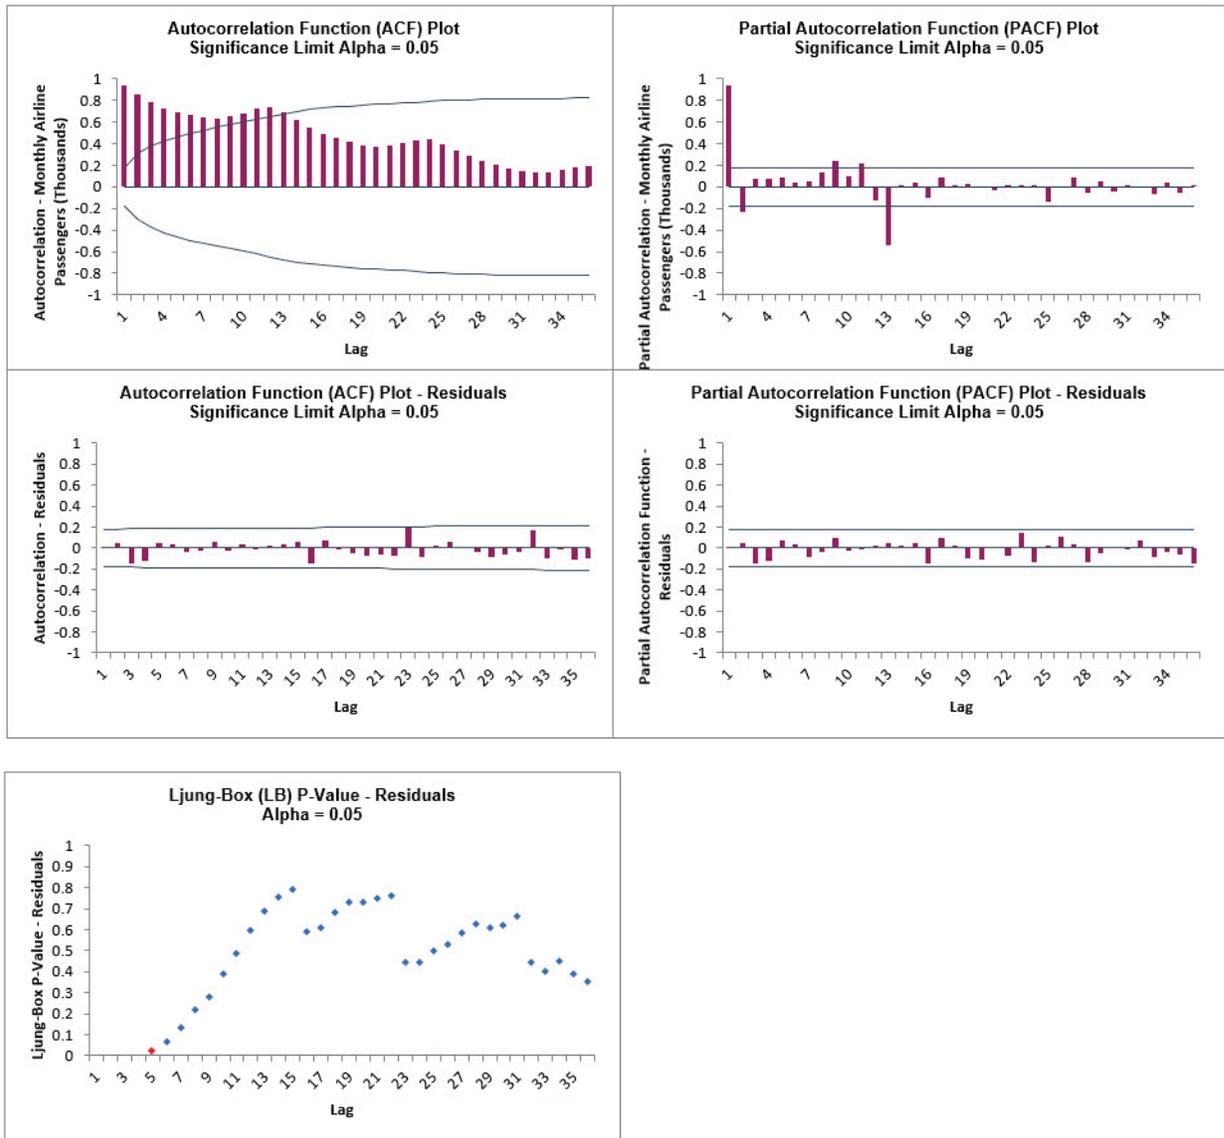
10. The Forecast Accuracy metrics are:

Forecast Accuracy		
Metric	In-Sample (Estimation) One-Step-Ahead Forecast	Out-of-Sample (Withhold) Multi-Step-Ahead Forecast
N	120	24
RMSE	7.115702971	26.70412899
MAE	5.455681009	22.99181598
MAPE	2.349324345	4.842245097
MASE	0.190931156	0.804639056

Comparing to our [earlier analysis](#), both the **In-Sample (Estimation) One-Step-Ahead Forecast** errors and **Out-of Sample (Withhold) Multi-Step-Ahead Forecast** errors are slightly smaller. (This was not expected, typically a Seasonal Exponential Smoothing or ARIMA model would give a more accurate forecast, based on comparison of methods using forecast competition data). However, given that we are forecasting out for two years, both models look very good.

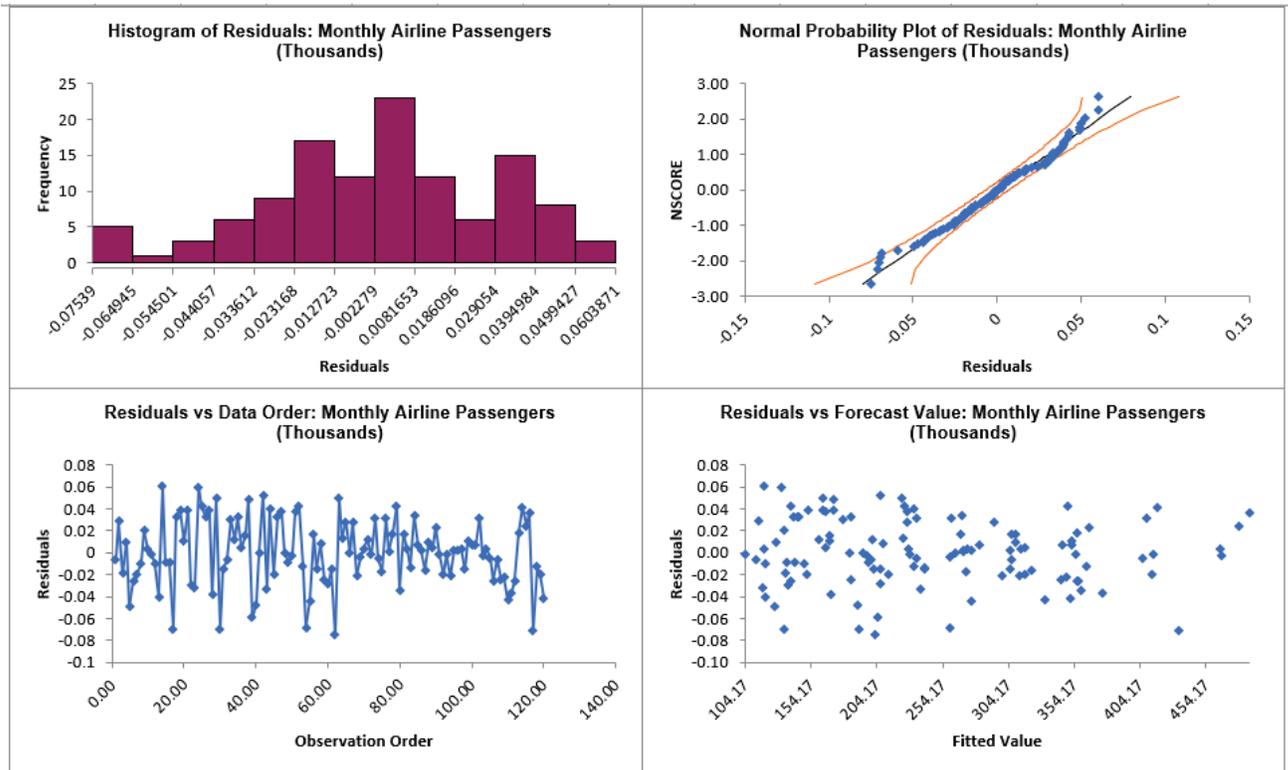
Forecast Accuracy metrics are calculated using the actual raw data versus inverse transformed forecast as displayed in the Forecast Chart and Table, so allow comparison across all model types and transformations.

11. Click on the **Exp Smooth MSD ACF PACF LB** sheet to view the ACF/PACF/LB Plots:



The ACF/PACF Residuals Plots are similar to the [previous analysis](#) and indicate that almost all of the autocorrelation has been accounted for in the model, however the Ljung-Box plot confirms that this is a better fit, with most P-Values being blue ($> .05$).

12. Click on the **Exp Smooth MSD Residuals** sheet to view the Residual Plots:



The residuals are approximately normally distributed, with a roughly straight line on the normal probability plot. There are no obvious extreme outliers or patterns in the charts.

Note that Residuals for MSD are the final observed -predicted values, so there are no scaling differences if the model uses additive or multiplicative error. Since a Box-Cox transformation was used, the residuals are in Ln transformed units.

Half-Hourly Multiple Seasonal Electricity Demand – Taylor

13. Open **Half-Hourly Multiple Seasonal Electricity Demand - Taylor.xlsx (Sheet 1 tab)**. This is half-hourly electricity demand (MW) in England and Wales from Monday, June 5, 2000 to Sunday, August 27, 2000 (taylor, R forecast). This data has multiple seasonality with frequency = 48 (observations per day) and 336 (observations per week), with a total of 4032 observations. See the [Run Chart](#), [ACF/PACF Plots](#), [Spectral Density Plot](#) and [Seasonal Trend Decomposition Plots](#) for this data.
14. Click **SigmaXL > Time Series Forecasting > Exponential Smoothing Forecast > Multiple Seasonal Decomposition Forecast**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
15. Select *Demand*, click **Numeric Time Series Data (Y) >>**. Check **Display ACF/PACF/LB Plots** and **Display Residual Plots**. Check **Specify Model Periods**. Set **Withhold Periods** = 96. Select **Withhold Forecast Type: Multi-Step-Ahead with Prediction Interval at Start of Withhold**. Check **Seasonal Frequency** with **Specify** = 48 336. Leave **Box-Cox Transformation** unchecked. We will use the default **Prediction Interval** = 95.0 %.

Exponential MSD Forecast

Obs. No.

Numeric Time Series Data (Y) >> Demand

Optional Time Axis Labels >>

<< Remove

OK >>

Cancel

Help

No. of Forecast Periods 24

Prediction Interval 95.0 %

Specify Model Periods

Start Model at Period 1

Withhold Periods 96

End Model at Period

Model Options

Seasonal Frequency

Specify 48 336

Select 4-Quarterly

Automatically Detect

Display ACF/PACF/LB Plots

Display Residual Plots

Box-Cox Transformation

Rounded Lambda

Optimal Lambda

Lambda & Threshold (Shift)

Withhold Forecast Type:

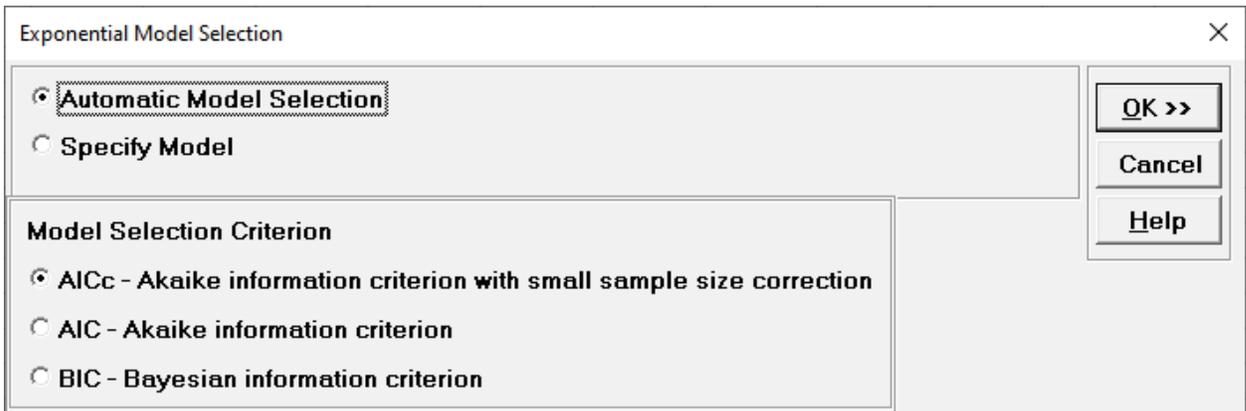
One-Step-Ahead with Prediction Interval at: Start of Withhold

Include in Residuals

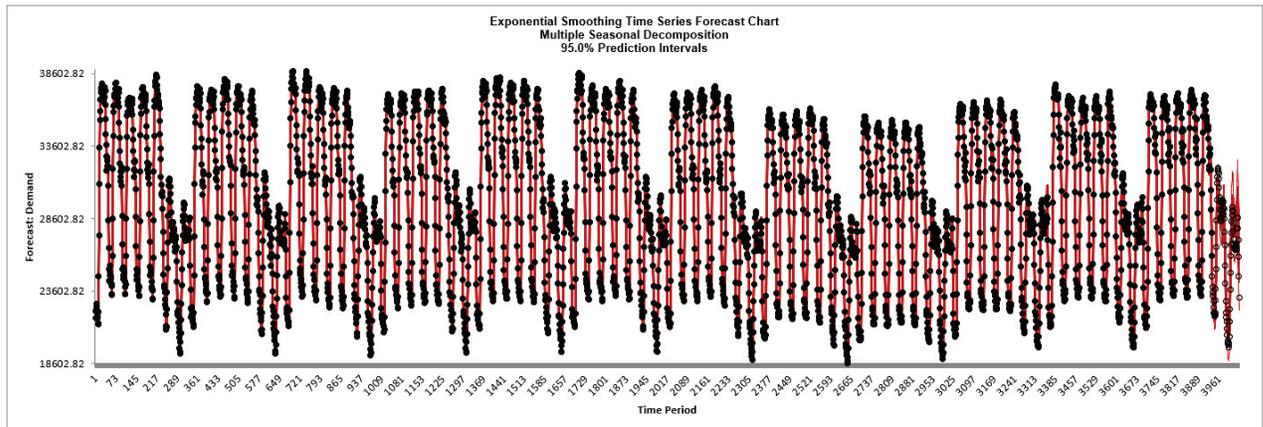
Multi-Step-Ahead with Prediction Interval at Start of Withhold.

Withhold Periods is 2*dominant seasonal frequency (48). Dominant frequency is obtained from the Spectral Density Plot. Start Model at Period = 1 is always greyed out for MSD.

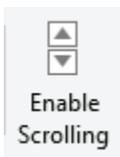
16. Click **Model Options**.



17. We will use the default **Automatic Model Selection** with **AICc** as the **Model Selection Criterion**. Click **OK** to return to the Exponential Smoothing MSD Forecast dialog. Click **OK**. The exponential smoothing forecast report is given:

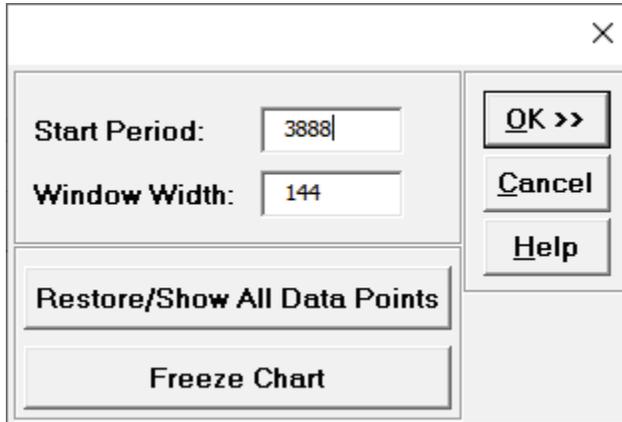


18. We will want to zoom in on the last 3 days, i.e., 144 half-hourly time periods, using chart scrolling. Click **SigmaXL Chart Tools > Enable Scrolling**



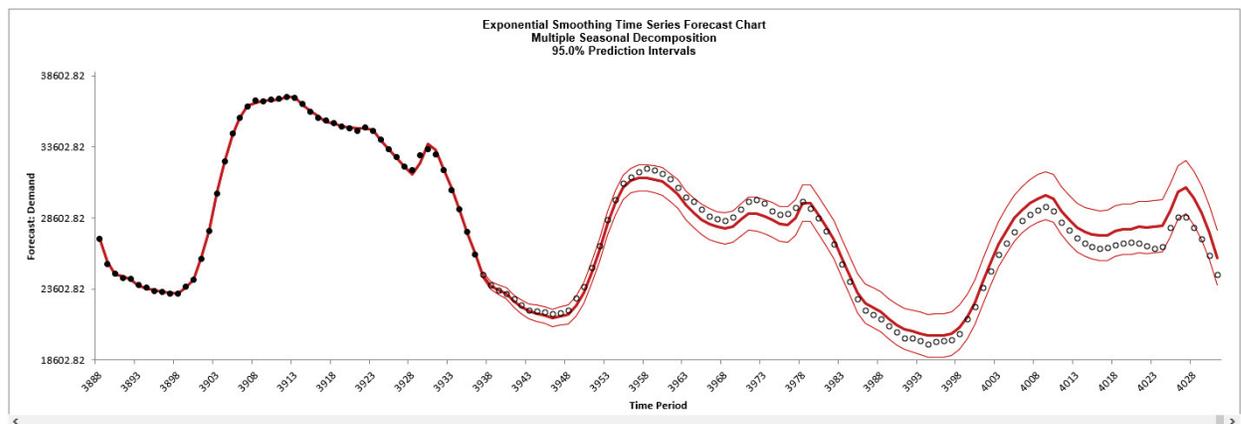
You may be prompted with a warning message that custom formatting on the chart will be cleared. You can avoid seeing this warning by checking **Save this choice as default and do not show this form again**.

19. Click **OK**. The scroll dialog appears allowing you to specify the **Start Period** and **Window Width**. Enter **Start Period** = 3888 and **Window Width** = 144:



At any point, you can click **Restore/Show All Data Points** or **Freeze Chart**. Freezing the chart will remove the scroll and unload the dialog. The scroll dialog will also unload if you change worksheets. To restore the dialog, click **SigmaXL Chart Tools > Enable Scrolling**.

20. Click **OK**. A scroll bar appears beneath the forecast chart. You can also change the **Start Subgroup** and **Window Width** and **Update**.



You can scroll through by clicking to the right or left, with the specified window width of 144.

The blank dots are the data values in the withhold sample with a multi-step forecast and prediction intervals displayed at the start of the withhold sample. The model does quite well at predicting the withhold 96 half-hour demand values. Note that the prediction error increases the further out we predict.

Click **Cancel** to exit the scroll dialog

21. Scroll down to view the Exponential Smoothing Model header:

Exponential Smoothing Model (Multiple Seasonal Decomposition): Simple Exponential Smoothing with Multiplicative Errors (M, N, N) - Model Automatically Selected
Model Periods: Model parameter estimates calculated excluding 96 withhold periods.

After **Multiple Seasonal Decomposition**, the model is **Simple Exponential Smoothing with Multiplicative Errors (M, N, N)**.

(M, N, N) was automatically selected as the best fit for the deseasonalized Demand data based on the AICc criterion.

The header also includes the number of specified withhold periods.

22. The Exponential Smoothing Model Information is given as:

Exponential Smoothing Model Information	
Seasonal Frequency	48, 336 Decomposition
Model Selection Criterion	AICc
Box-Cox Transformation	N/A
Lambda	
Threshold	

This is a summary of model information with Seasonal Frequency = 48, 336 using decomposition and Model Selection Criterion = "AICc".

23. The Parameter Estimates for the deseasonalized Demand data are:

Parameter Estimates	
Term	Coefficient
alpha (level smoothing)	0.9999
I (initial level)	29798.08498

- Error includes the smoothing parameter *alpha* and initial level value (*I*). With *alpha* = 0.9999, this is effectively a naïve forecast, with multiplicative errors.
- Seasonal smoothing parameter (*gamma*) is not computed because the data has been deseasonalized.

24. The Exponential Smoothing Model Statistics are:

Exponential Smoothing Model Statistics	
No. of Observations	3936
DF	3934
StDev	0.003336007
Variance	1.11289E-05
Log-Likelihood	-34366.34649
AICc	68738.69908
AIC	68738.69298
BIC	68757.52674

- The number of observations, $n = 4032 - 96$ (*withhold*) = 3936
- Degrees of freedom (DF) = 3936 (n) - 2 (terms in the model) = 3934
- Note that the model statistics are calculated using the deseasonalized data and residuals are multiplicative (relative) errors: $e_t = \frac{y_t - \hat{y}_t}{\hat{y}_t}$.

25. The Forecast Accuracy metrics are:

Forecast Accuracy		
Metric	In-Sample (Estimation) One-Step-Ahead Forecast	Out-of-Sample (Withhold) Multi-Step-Ahead Forecast
N	3936	96
RMSE	101.5029509	776.7232991
MAE	73.73426784	649.5237662
MAPE	0.253138484	2.45436482
MASE	0.040157098	0.353743116

As expected, the **Out-of Sample (Withhold) Multi-Step-Ahead Forecast** errors are larger than the **In-Sample (Estimation) One-Step-Ahead Forecast** errors. Note, if we were primarily interested in a short term one-step ahead forecast, then we would have selected **Withhold Forecast Type: One-Step-Ahead** and the above table would show **Out-of Sample (Withhold) One-Step-Ahead Forecast** errors.

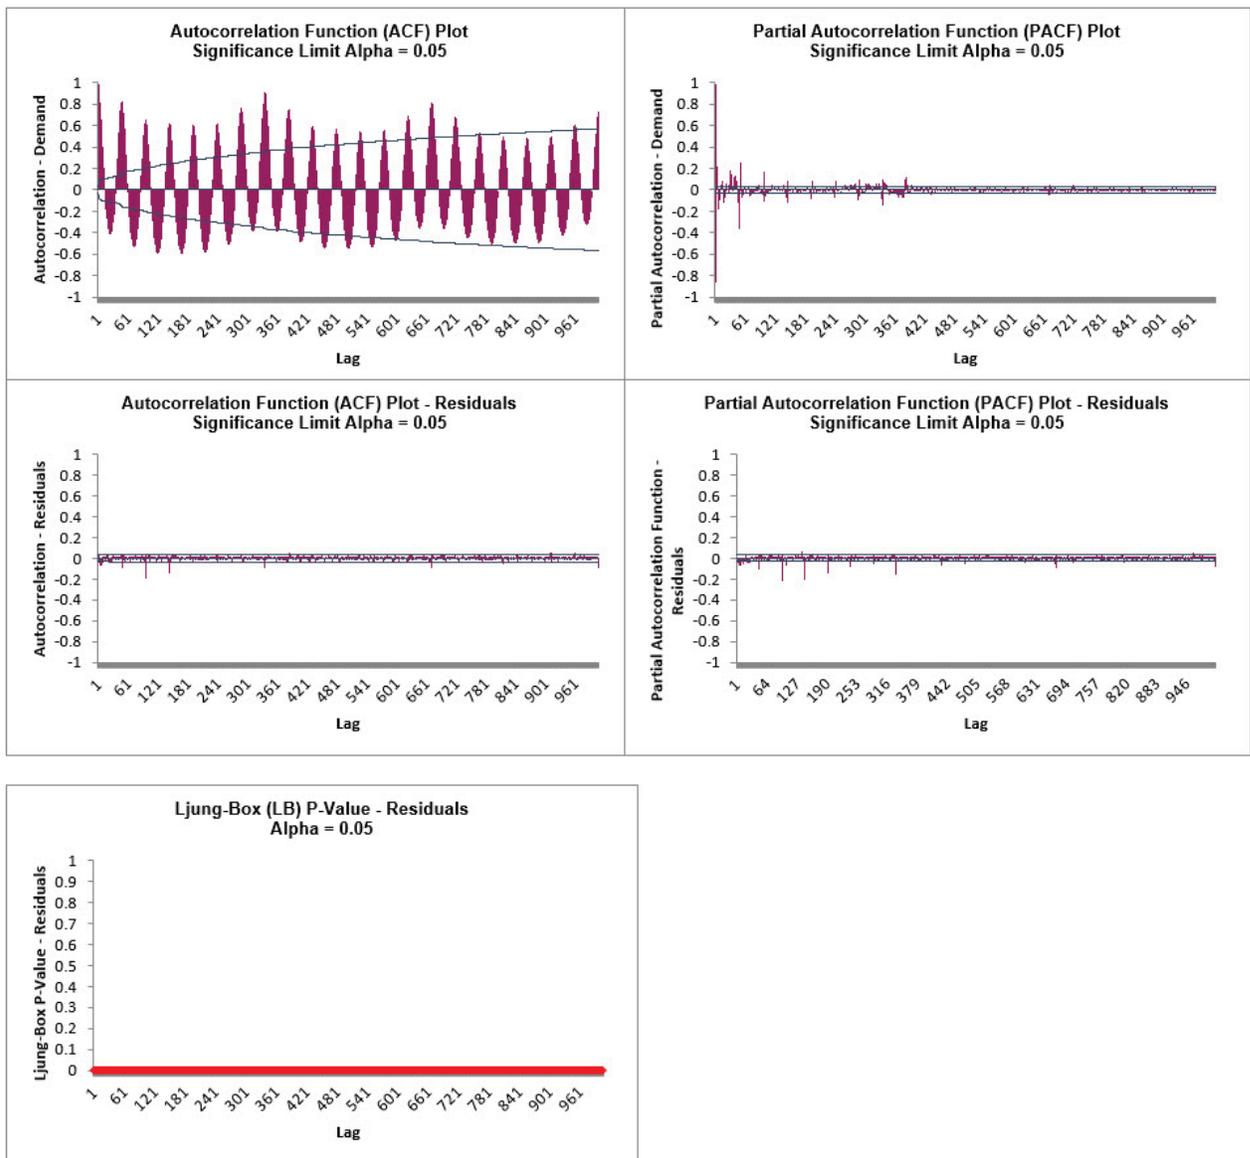
Forecast Accuracy metrics are calculated using the actual raw data versus forecast as displayed in the Forecast Chart and Table so, unlike the model statistics above, allows comparison across all forecast model types and transformations.

26. The Forecast Table is given as:

Forecast Table				
Period	Withhold Data	Multi-Step-Ahead Forecast	Lower 95.0% PI	Upper 95.0% PI
3937	24653	24595.68006	24399.19061	24792.1695
3938	23879	23839.01162	23561.1467	24116.87654
3939	23508	23546.05076	23205.74185	23886.35968
3940	23275	23300.43291	22907.48021	23693.38561
3941	22890	22735.89622	22296.56272	23175.22972
3942	22462	22349.79507	21868.5296	22831.06054

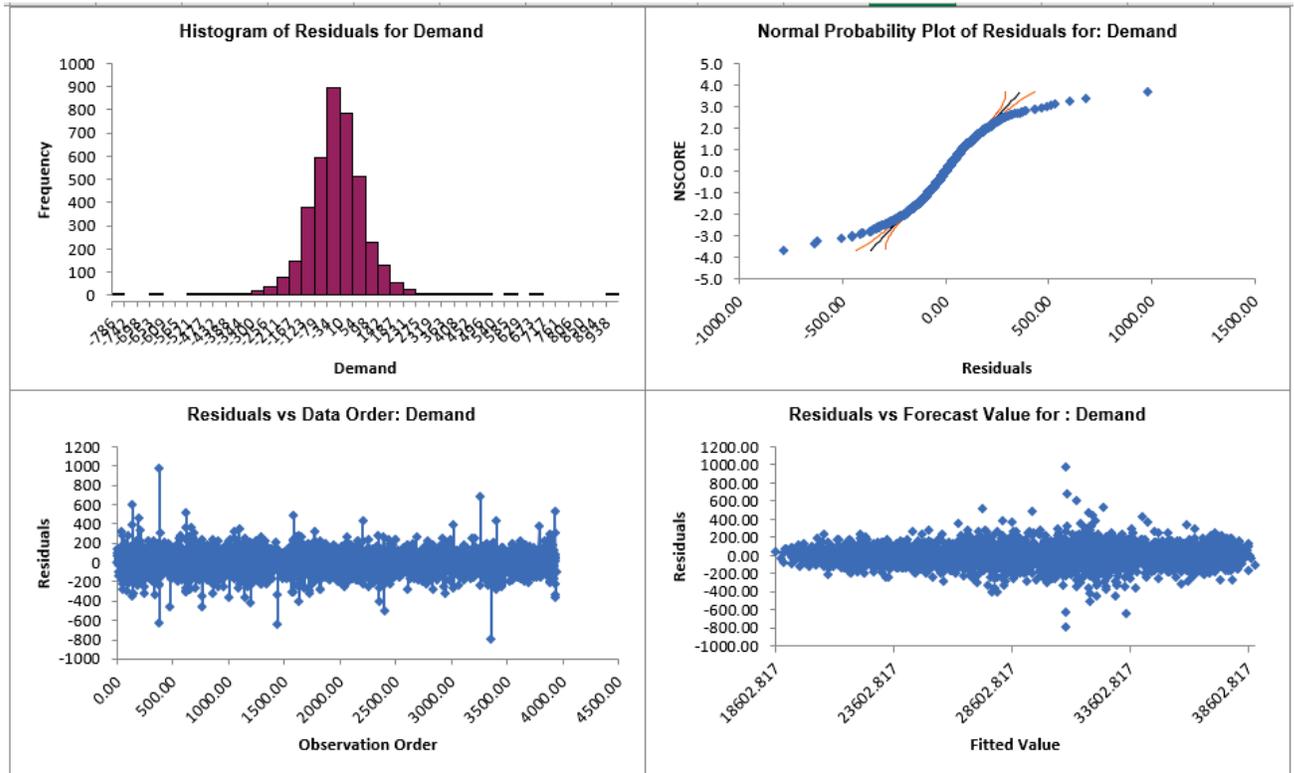
These are the same forecast and prediction interval values displayed in the Forecast Chart, but provided for further analysis or charting (e.g., a run chart of the forecast errors). The **Withhold Data** is also displayed.

27. Click on the **Exp Smooth MSD ACF PACF LB** sheet to view the ACF/PACF/LB Plots:



The ACF/PACF Residuals Plots indicate that much of the autocorrelation has been accounted for in the model, but the Ljung-Box plot shows that some significant autocorrelation still remains (the red P-Values are significant at $\alpha=.05$) - so the model can potentially be improved. This does not mean that the model is a bad model, it can still be very useful for prediction purposes, but the prediction intervals may not provide accurate coverage.

28. Click on the **Exp Smooth MSD Residuals** sheet to view the Residual Plots:



The residuals are not normally distributed and there are extreme outliers. These should be investigated with a control chart on the residuals. Outliers in Electricity Demand are often explained by Temperature and this is something we will look at later with a different data set (Daily Electricity Demand with Predictors – ElecDaily).

29. Note that Residuals for MSD are the final observed - predicted values, so there are no scaling differences if the model uses additive or multiplicative error. If a Box-Cox transformation is used, then the residuals are in transformed units.

Exponential Smoothing Control Chart

Statistical Process Control (SPC) For Autocorrelated Data

An Individuals control chart is created using the residuals of the Exponential Smoothing forecast model.

The Moving Limits chart uses the one step prediction as the center line, so the control limits will move with the center line. If a Box-Cox transformation is used then an inverse transformation is applied to calculate the control limits. If the residuals are exponential smoothing multiplicative, the control limits are approximate and out-of-control signals may not exactly match the Individuals Chart. If that occurs, the Individuals Chart should be used to determine what points are out-of-control.

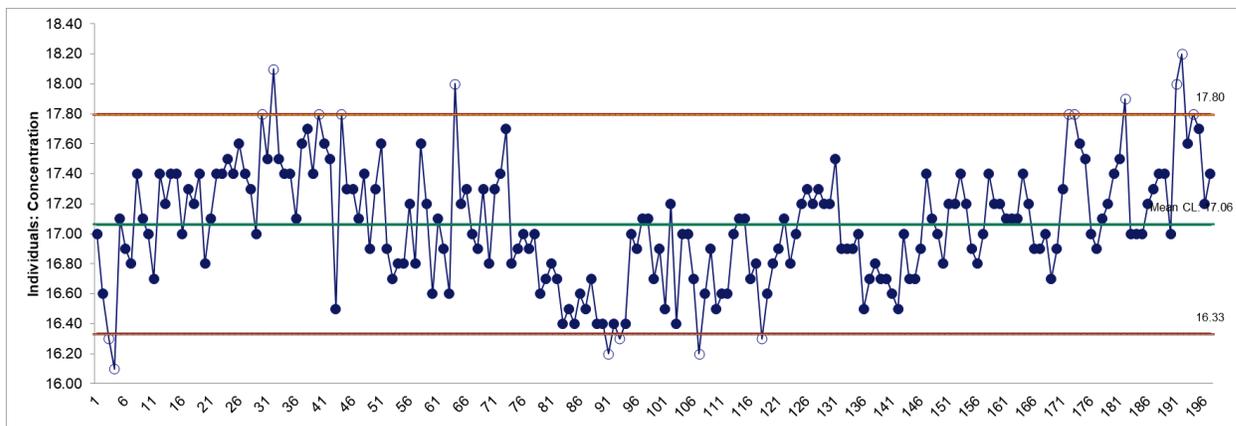
The popular “Add Data”, “Show Last 30” and “Scroll” features in SigmaXL Chart Tools are available for these control charts. For “Add Data”, the time series models are not refitted, but used to compute the residual values for the new data.

For further details and references, see the Appendix: [Control Charts for Autocorrelated Data](#).

Note that a Moving Range Chart and Tests for Special Causes are not available here, but the user can store and select Residuals, then create with **SigmaXL > Control Charts > Individuals & Moving Range**.

Chemical Process Concentration - Series A

1. Open **Chemical Process Concentration – Series A.xlsx (Sheet 1 tab)**. This is the Series A data from Box and Jenkins, a set of 197 concentration values from a chemical process taken at two-hour intervals. See the [Run Chart](#), [ACF/PACF Plots](#), [Spectral Density Plot](#) and [Seasonal Trend Decomposition Plots](#) for this data.
2. Earlier we saw that this process has significant autocorrelation. In order to see the impact on a control chart, we will construct an Individuals chart on the raw data. Click **SigmaXL > Control Charts > Individuals**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
3. Select *Concentration*, click **Numeric Data Variable (Y) >>**. Click **OK**. An Individuals Control Chart is produced:



There are 17 out-of-control data points, largely due to the autocorrelation. Searching for assignable causes using this chart as is would be futile.

4. Now click **Sheet 1 tab** and **SigmaXL > Time Series Forecasting > Exponential Smoothing Control Chart > Control Chart**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
5. Select *Concentration*, click **Numeric Time Series Data (Y) >>**. Uncheck **Display ACF/PACF/LB Plots**. Leave **Display Residual Plots**, **Specify Model Periods**, **Seasonal Frequency** and **Box-Cox Transformation** unchecked.

Exponential Smoothing Control Chart

Observation No.

Numeric Time Series Data (Y) >> Concentration

Optional Time Axis Labels >>

<< Remove

OK >>
Cancel
Help

Specify Model Periods

Start Model/Control Limit Calculations at Period 1

Withhold Periods 0

End Model/Control Limit Calculations at Period

Model Options

Seasonal Frequency

Specify 12

Select 4-Quarterly

Automatically Detect

Display ACF/PACF/LB Plots

Display Residual Plots

Box-Cox Transformation

Rounded Lambda

Optimal Lambda

Lambda & Threshold (Shift)

Since we will be running the same (A, N, N) model as used earlier, we will not need the ACF/PACF/LB and Residuals Plots.

- Click **Model Options**. Select **Specify Model**.

Exponential Model Selection

Automatic Model Selection

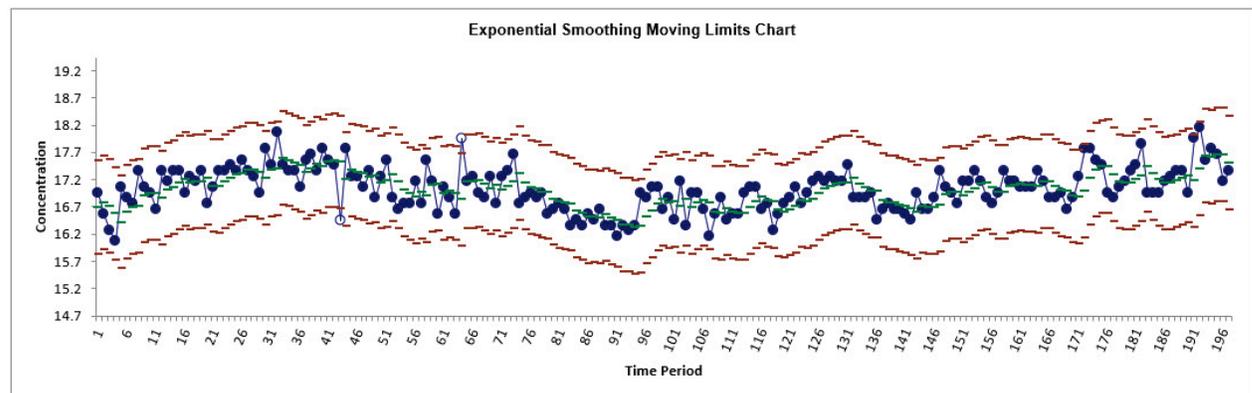
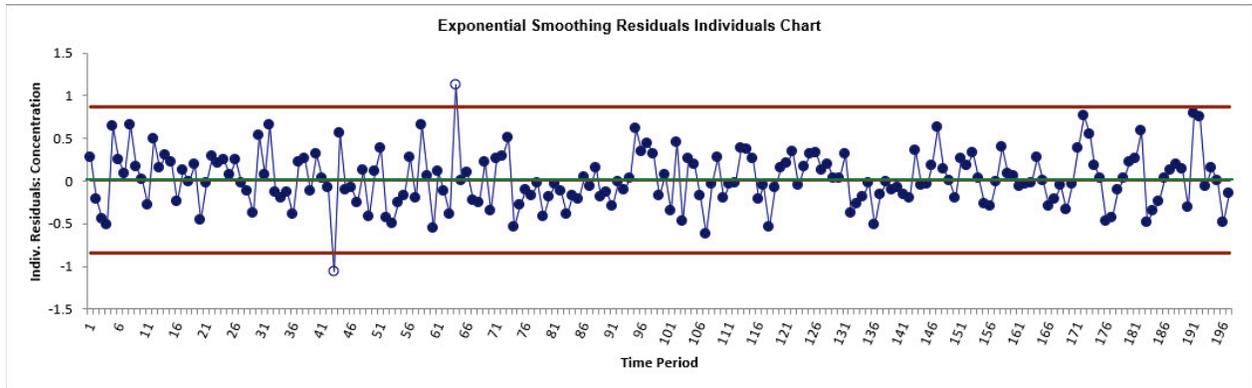
Specify Model

OK >>
Cancel
Help

Error	Trend	Seasonal
<input checked="" type="radio"/> Additive	<input checked="" type="radio"/> None	<input checked="" type="radio"/> None
<input type="radio"/> Multiplicative	<input type="radio"/> Additive	<input type="radio"/> Additive
	<input type="radio"/> Additive Damped	<input type="radio"/> Multiplicative

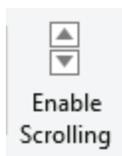
Simple Exponential Smoothing with Additive Errors (A, N, N) - Exponentially Weighted Moving Average (EWMA)

- We will use the default **Error: Additive** and **Trend: None**, which is a simple exponential smoothing model, or Exponentially Weighted Moving Average (EWMA). Click **OK** to return to the Exponential Smoothing Control Chart dialog. Click **OK**. The exponential smoothing control charts are produced:



Now we only have two out-of-control data points on the Individuals chart to investigate. The Moving Limits chart uses the one step prediction as the center line, so the control limits move with the center line. The reduction in out-of-control data points is due to the removal of the autocorrelation as noted earlier in the [ACF/PACF/LB Plots](#) for this model.

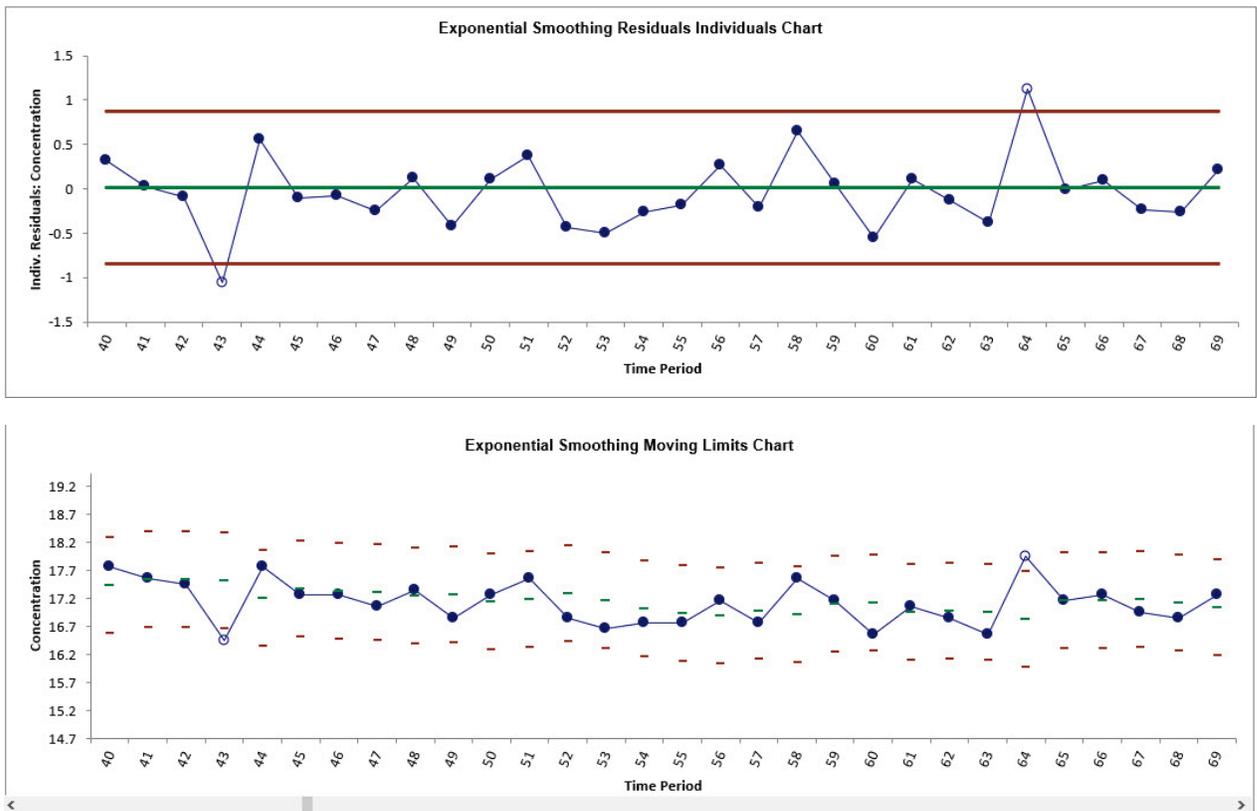
- You can scroll through the chart data points. Click **SigmaXL Chart Tools > Enable Scrolling**



You may be prompted with a warning message that custom formatting on the chart will be cleared. You can avoid seeing this warning by checking **Save this choice as default and do not show this form again**.

- Click **OK**. The scroll dialog appears allowing you to specify the **Start Subgroup** and **Window Width**. Enter **Start Subgroup** = 40 and **Window Width** = 30 to view the two out-of-control data points.

10. Click **OK**. This allows us to zoom in on the out-of-control points at 43 and 64.



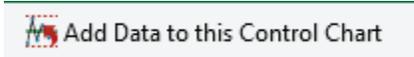
Observation **43** is lower than expected from the exponential smoothing forecast model.
 Observation **64** is higher than expected.

11. Click **Cancel** to exit the scroll dialog.

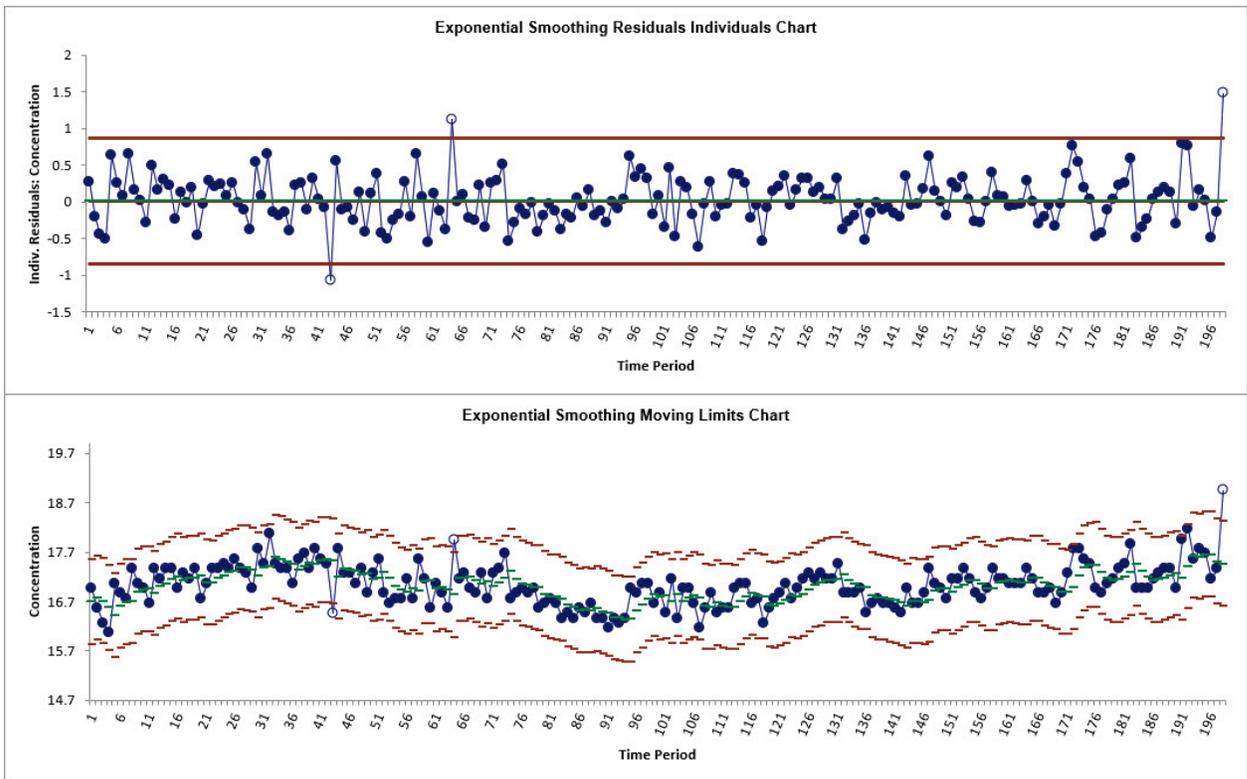
12. Now we will add a new data point to the Series A Concentration Data. The residuals will be computed using the same model as above without re-estimation of the model parameters or recalculation of the control limits. This is also known as the “Phase II” application of a Control Chart, where an out-of-control signal should lead to an investigation into the assignable cause and corrective action or process adjustment applied. Click Sheet1, enter the value 19 as shown in cell **B199** (and optionally Observation number 198 in cell **A199**).

196	195	17.7
197	196	17.2
198	197	17.4
199	198	19

- Click **Exp. Smoothing Control Charts** tab (if more than one control chart sheet exists in the workbook, please select the chart where the data will be added).
- Click **SigmaXL Chart Tools > Add Data to this Control Chart**



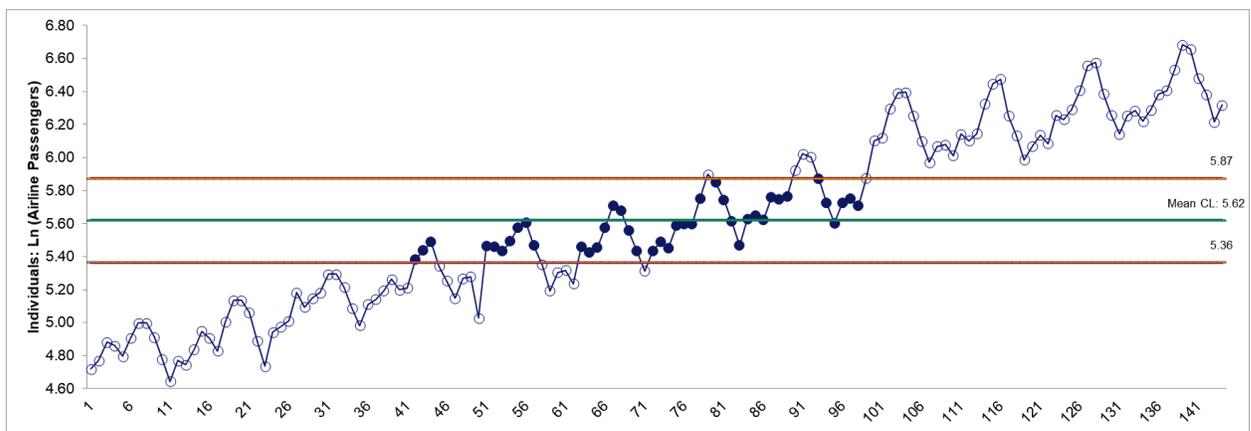
- The Residuals Individuals Control Chart and Moving Limits Charts are now updated with the new data, showing this as an out-of-control data point:



- We recommend renaming the workbook to **Chemical Process Concentration – Series A_AddData.xlsx**, so that later use of the Concentration data does not include the added data point.

Monthly Airline Passengers – Modified for Control Charts

17. Open **Monthly Airline Passengers – Modified for Control Charts.xlsx (Sheet 1 tab)**. This is based on the Series G data from Box and Jenkins, monthly total international airline passengers for January 1949 to December 1960. A Ln transformation is applied (avoiding the need for a Box-Cox transformation), a negative outlier is added at 50 (-.25) and a level shift applied (+.25), starting at 100. Coded variables were added to help distinguish an outlier versus a shift, but they will be analyzed later using ARIMA Forecast with Predictors. Exponential Smoothing does not support predictors.
18. Earlier we saw that this process has significant autocorrelation with a strong trend and seasonality. In order to see the impact on a control chart, we will construct an Individuals chart on the raw data. Click **SigmaXL > Control Charts > Individuals**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
19. Select *Ln (Airline Passengers-Modified)*, click **Numeric Data Variable (Y) >>**. Click **OK**. An Individuals Control Chart is produced:



With strong trend, seasonality and positive autocorrelation, this control chart is meaningless.

20. Now click **Sheet 1 tab** and **SigmaXL > Time Series Forecasting > Exponential Smoothing Control Chart > Control Chart**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
21. Select *Ln(Airline Passengers-Modified)*, click **Numeric Time Series Data (Y) >>**. Uncheck **Display ACF/PACF/LB Plots** and **Display Residual Plots**. Check **Seasonal Frequency with Specify = 12**. Leave **Specify Model Periods** and **Box-Cox Transformation** unchecked.

Exponential Smoothing Control Chart

Obs. No.
 Outlier 50
 Shift 50
 Outlier 100
 Shift 100

Numeric Time Series Data (Y) >> Ln (Airline Passengers-Mc)

Optional Time Axis Labels >>

<< Remove

OK >>
 Cancel
 Help

Specify Model Periods

Start Model/Control Limit Calculations at Period 1

Withhold Periods 0

End Model/Control Limit Calculations at Period

Model Options

Seasonal Frequency

Specify 12

Select 4 - Quarterly

Automatically Detect

Display ACF/PACF/LB Plots

Display Residual Plots

Box-Cox Transformation

Rounded Lambda

Optimal Lambda

Lambda & Threshold (Shift)

22. Click **Model Options**.

Exponential Model Selection

Automatic Model Selection

Specify Model

Model Selection Criterion

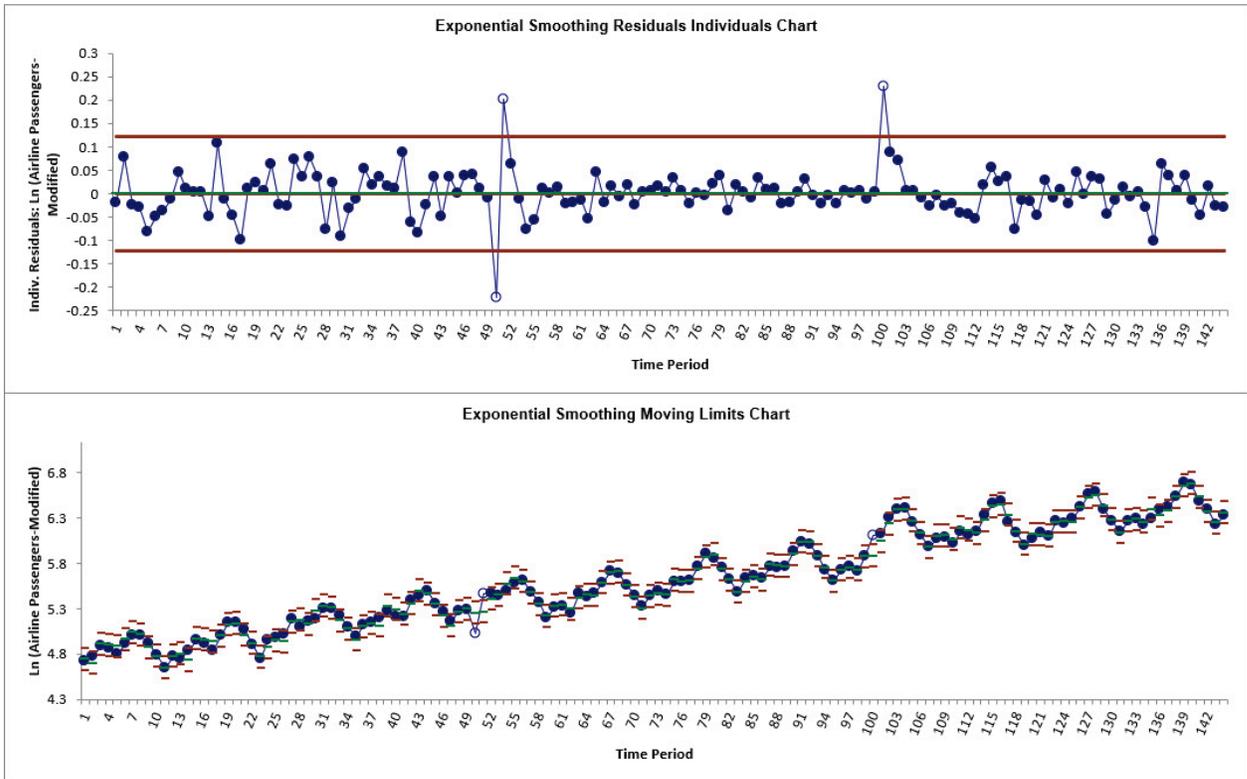
AICc - Akaike information criterion with small sample size correction

AIC - Akaike information criterion

BIC - Bayesian information criterion

OK >>
 Cancel
 Help

23. We will use the default **Automatic Model Selection** with **AICc** as the **Model Selection Criterion**. Click **OK** to return to the Exponential Smoothing Control Chart dialog. Click **OK**. The exponential smoothing control charts are produced:



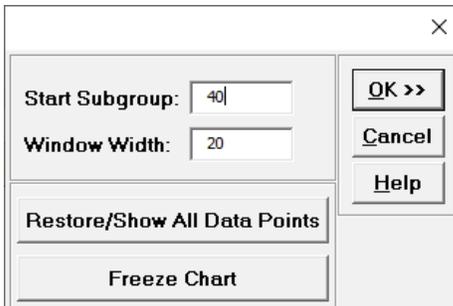
Now we can clearly see the out-of-control data points at 50, 51 and 100 on the Individuals chart. In order to view the points on the Moving Limits chart we will use scrolling.

24. Click **SigmaXL Chart Tools > Enable Scrolling**

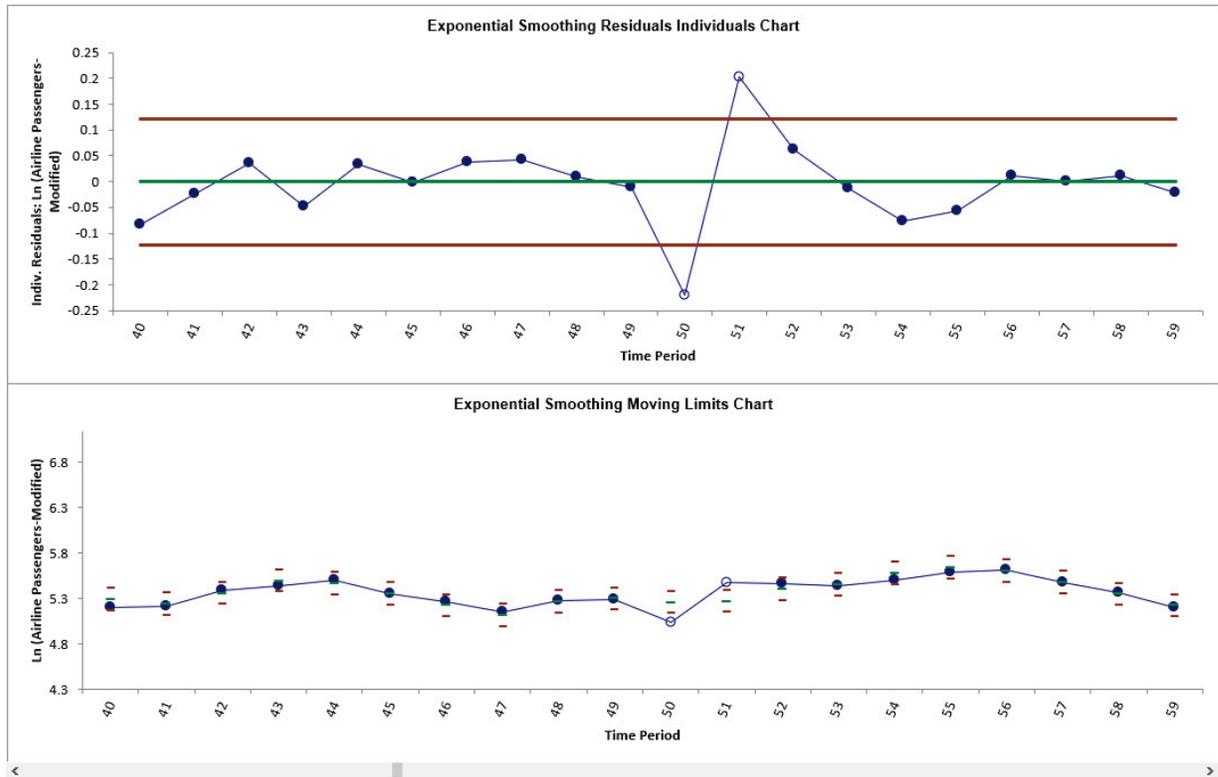


You may be prompted with a warning message that custom formatting on the chart will be cleared. You can avoid seeing this warning by checking **Save this choice as default and do not show this form again**.

25. Click **OK**. The scroll dialog appears allowing you to specify the **Start Subgroup** and **Window Width**. Enter **Start Subgroup = 40** and **Window Width = 20** to view the first two out-of-control data points.



26. Click **OK**. This allows us to zoom in on the out-of-control points at 50 and 51.



Observation **50** is lower than expected from the exponential smoothing forecast model. Observation **51** is higher than expected. Later investigation will reveal that this is a single negative outlier.

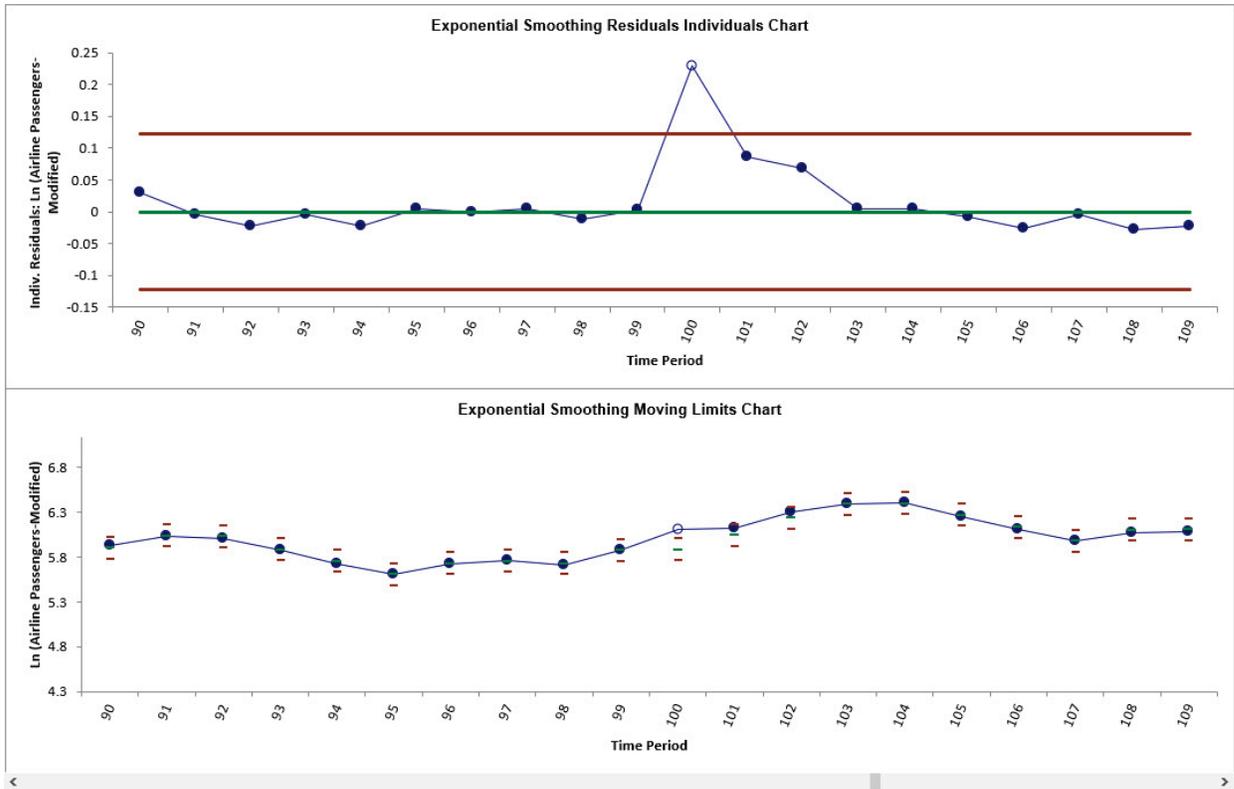
Tip: Scrolling keeps the original Y axis minimum and maximum setting. You may wish to change this to auto by clicking on the Y axis, right click **Format Axis**, click Bounds Minimum **Reset** and Bounds Maximum **Reset**. This changes the axis settings to Auto so when you scroll or Update the Y axis will automatically adjust as well.

27. Now enter **Start Subgroup** = 90 and **Window Width** = 20 to view the third out-of-control data point.

The screenshot shows a dialog box with the following controls:

- Start Subgroup:** Input field containing the value 90.
- Window Width:** Input field containing the value 20.
- Update** button
- Cancel** button
- Help** button
- Restore/Show All Data Points** button
- Freeze Chart** button

28. Click **Update**.



Observation **100** is higher than expected from the exponential smoothing forecast model. Later investigation will reveal that this is a shift in the mean.

29. Click **Cancel** to exit the scroll dialog.

30. Scroll down to view the Exponential Smoothing Model header:

Exponential Smoothing Model: Additive Trend, Additive Seasonal Method with Additive Errors (Holt-Winters) (A, A, A) - Model Automatically Selected
 Model Periods: All observations are used in the Exponential Smoothing model estimation. No withhold periods available for out-of-sample forecast accuracy evaluation.

The model **Additive Trend, Additive Seasonal Method with Additive Errors (Holt-Winters) (A, A, A)** was automatically selected as the best fit for the Modified Ln Airline Passenger data based on the AICc criterion.

31. The Parameter Estimates and Exponential Smoothing Model Statistics are slightly different than our [earlier analysis](#) because we have introduced an outlier and a shift, as well here we are using all of the data, i.e., there are no withhold periods. Note that earlier we used a Box-Cox Transformation with Lambda=0 and here we are using Ln of the data.

32. The Forecast Accuracy metrics are given as:

Forecast Accuracy		
Metric	In-Sample (Estimation) One-Step-Ahead Forecast	Out-of-Sample (Withhold) One-Step-Ahead Forecast
N	144	
RMSE	0.049508351	
MAE	0.033506668	
MAPE	0.609538124	
MASE	0.230072112	

Note that these forecast errors are very different than our [earlier analysis](#) where the forecast errors were calculated on the raw data versus final predicted values, but here we are using Ln of the Airline Passenger data.

Exponential Smoothing Multiple Seasonal Decomposition (MSD) Control Chart

Exponential Smoothing is limited to a maximum seasonal frequency of 24. For higher frequencies use Exponential Smoothing – Multiple Seasonal Decomposition (MSD). The seasonal component is first removed through decomposition, a nonseasonal exponential smooth model fitted to the remainder (+trend), and then the seasonal component is added back in.

As the name implies, Multiple Seasonal Decomposition (MSD) also accommodates multiple seasonality, for example the half-hourly data with a seasonal frequency of 48 observations per day and 336 observations per week.

An Individuals control chart of the residuals is created for this forecast method. The Moving Limits chart uses the one step prediction as the center line, so the control limits will move with the center line. If a Box-Cox transformation is used then an inverse transformation is applied to calculate the control limits. If the residuals are exponential smoothing multiplicative, the moving control limits are approximate and out-of-control signals may not exactly match the Individuals Chart. If that occurs, the Individuals Chart should be used to determine what points are out-of-control.

The popular “Add Data”, “Show Last 30” and “Scroll” features in SigmaXL Chart Tools are available for these control charts. For “Add Data”, the time series models are not refitted, but used to compute the residual values for the new data.

Monthly Airline Passengers – Modified for Control Charts

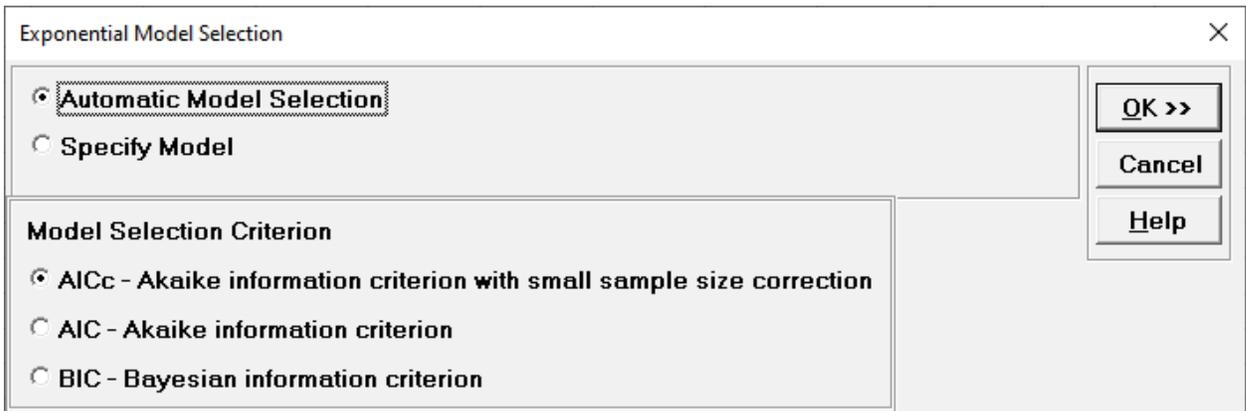
1. Open **Monthly Airline Passengers – Modified for Control Charts.xlsx (Sheet 1 tab)**. This is based on the Series G data from Box and Jenkins, monthly total international airline passengers for January 1949 to December 1960. A Ln transformation is applied (avoiding the need for a Box-Cox transformation), a negative outlier is added at 50 (-.25) and a level shift applied (+.25), starting at 100. The Multiple Seasonal Decomposition (MSD) option is not necessary for this data, but by way of introduction, we will use this to compare to the previous analysis.
2. Click **SigmaXL > Time Series Forecasting > Exponential Smoothing Control Chart > Multiple Seasonal Decomposition Control Chart**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
3. Select *Ln(Airline Passengers-Modified)*, click **Numeric Time Series Data (Y) >>**. Uncheck **Display ACF/PACF/LB Plots** and **Display Residual Plots**. Check **Seasonal Frequency with Specify = 12**. Leave **Specify Model Periods** and **Box-Cox Transformation** unchecked.

Exponential Smoothing MSD Control Chart

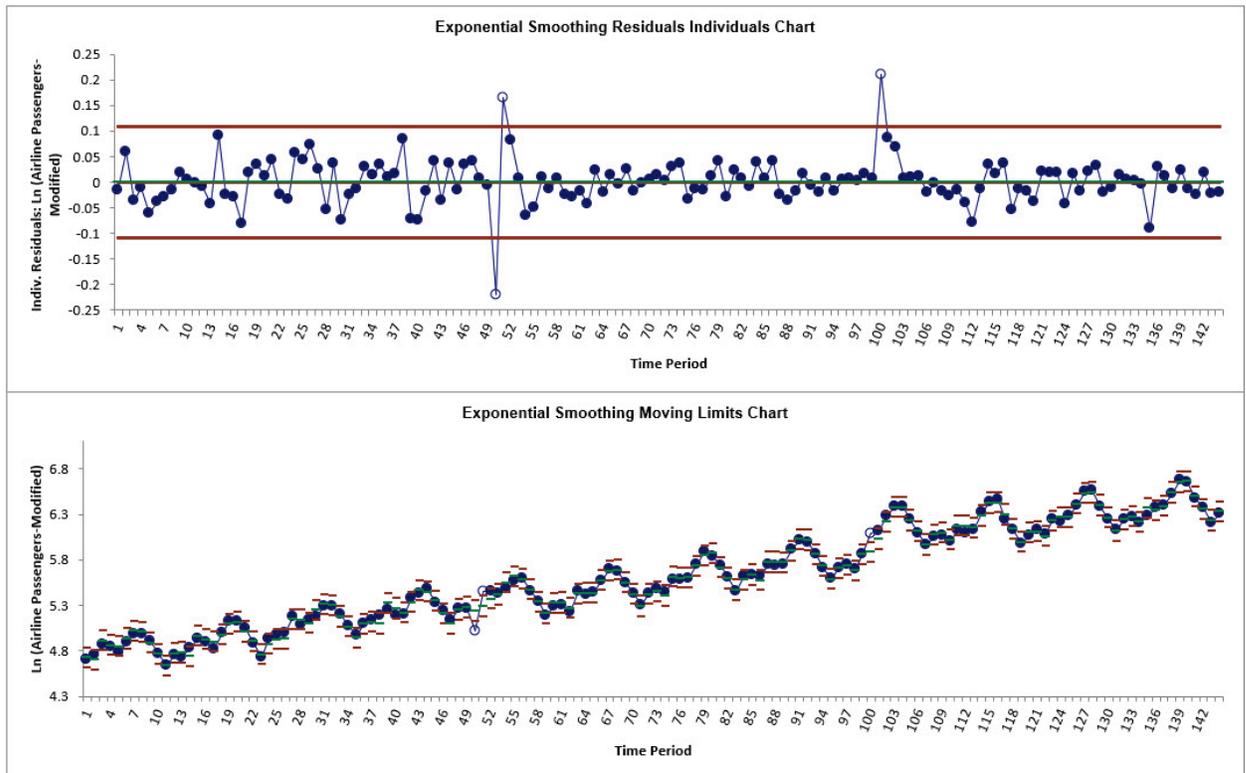
Obs. No. Outlier 50 Shift 50 Outlier 100 Shift 100	N umeric Time Series Data (Y) >>	Ln (Airline Passengers-Mc	O K >>
	O ptional Time Axis Labels >>		C ancel
	<< R emove		H elp

<input type="checkbox"/> S pecify Model Periods	M odel Options	<input type="checkbox"/> D isplay ACF/PACF/LB Plots
Start Model/Control Limit Calculations at Period <input type="text" value="1"/>	S easonal Frequency	<input type="checkbox"/> D isplay Residual Plots
<input checked="" type="radio"/> W ithhold Periods <input type="text" value="0"/>	<input checked="" type="radio"/> S pecify <input type="text" value="12"/>	<input type="checkbox"/> B ox-Cox Transformation
End Model/Control Limit Calculations at Period <input type="text"/>	<input type="radio"/> S elect <input type="text" value="4-Quarterly"/>	<input checked="" type="radio"/> R ounded L ambda
	<input type="radio"/> A utomatically Detect	<input type="radio"/> O ptimal Lambda
		<input type="radio"/> L ambda & T hreshold (Shift)

4. Click **Model Options**.



5. We will use the default **Automatic Model Selection** with **AICc** as the **Model Selection Criterion**. Click **OK** to return to the Exponential Smoothing Control Chart dialog. Click **OK**. The exponential smoothing (MSD) control charts are produced:



We can clearly see the out-of-control data points at 50, 51 and 100 on the Residuals Individuals chart. This matches what we observed [previously](#) with regular Exponential Smoothing Control Charts.

6. Scroll down to view the Exponential Smoothing Model header:

Exponential Smoothing Model (Multiple Seasonal Decomposition): Additive Trend Method with Additive Errors (Holt's Linear) (A, A, N) - Model Automatically Selected
 Model Periods: All observations are used in the Exponential Smoothing model estimation. No withhold periods available for out-of-sample forecast accuracy evaluation.

The model **Additive Trend Method with Additive Errors (Holt's Linear) (A, A, N)** was automatically selected as the best fit for the deseasonalized Modified Ln Airline Passenger data based on the AICc criterion. There are no withhold periods.

7. The Exponential Smoothing Model Information is given as:

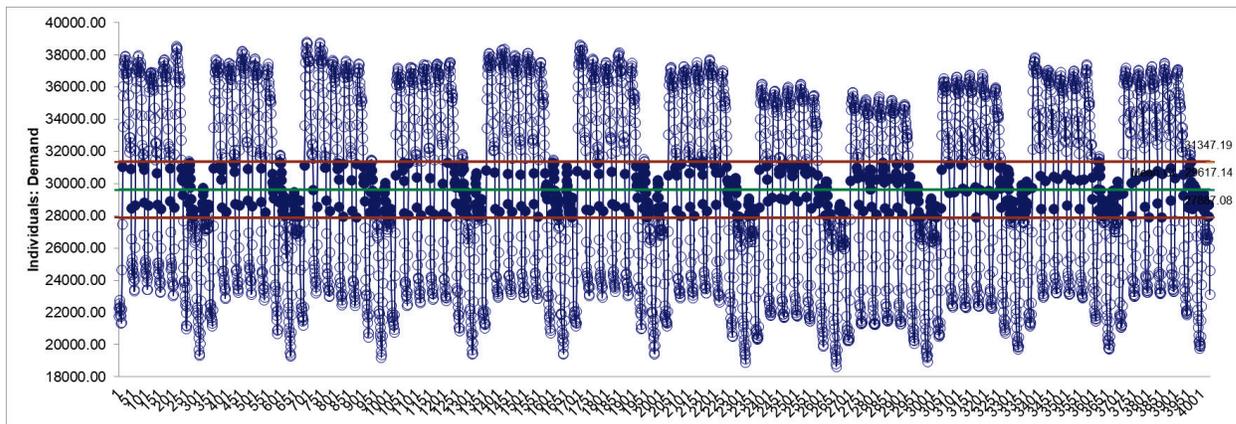
Exponential Smoothing Model Information	
Seasonal Frequency	12 Decomposition
Model Selection Criterion	AICc
Box-Cox Transformation	N/A
Lambda	
Threshold	

This is a summary of model information with Seasonal Frequency = 12 using Decomposition and Model Selection Criterion = "AICc". The Box-Cox Transformation is "N/A".

8. We will not review the Parameter Estimates, Model Statistics, and Forecast Accuracy as they are close to the MSD values given [earlier](#), although note that slight differences are due to the introduction of an outlier and a shift, as well now we are using all of the data, i.e., there are no withhold periods. Earlier we used a Box-Cox Transformation with Lambda=0 and here we are using Ln of the data.

Half-Hourly Multiple Seasonal Electricity Demand – Taylor

9. Open **Half-Hourly Multiple Seasonal Electricity Demand - Taylor.xlsx (Sheet 1 tab)**. This is half-hourly electricity demand (MW) in England and Wales from Monday, June 5, 2000 to Sunday, August 27, 2000 (taylor, R forecast). This data has multiple seasonality with frequency = 48 (observations per day) and 336 (observations per week), with a total of 4032 observations. See the [Run Chart](#), [ACF/PACF Plots](#), [Spectral Density Plot](#) and [Seasonal Trend Decomposition Plots](#) for this data.
10. We will first construct a classical Individuals Control Chart on the raw data. Click **SigmaXL > Control Charts > Individuals**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
11. Select *Demand*, click **Numeric Data Variable (Y) >>**. Click **OK**. An Individuals Control Chart is produced:



With the high frequency seasonality, this control chart is meaningless.

12. Click **SigmaXL > Time Series Forecasting > Exponential Smoothing Control Chart > Multiple Seasonal Decomposition Control Chart**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
13. Select *Demand*, click **Numeric Time Series Data (Y) >>**. Uncheck **Display ACF/PACF/LB Plots** and **Display Residual Plots**. Check **Seasonal Frequency** with **Specify = 48 336**. Leave **Specify Model Periods** and **Box-Cox Transformation** unchecked.

Exponential Smoothing MSD Control Chart

Obs. No.

Numeric Time Series Data (Y) >> Demand

Optional Time Axis Labels >>

<< Remove

OK >>
Cancel
Help

Specify Model Periods

Start Model/Control Limit Calculations at Period 1

Withhold Periods 0

End Model/Control Limit Calculations at Period

Model Options

Seasonal Frequency

Specify 48 336

Select 4 - Quarterly

Automatically Detect

Display ACF/PACF/LB Plots

Display Residual Plots

Box-Cox Transformation

Rounded Lambda

Optimal Lambda

Lambda & Threshold (Shift)

14. Click **Model Options**.

Exponential Model Selection

Automatic Model Selection

Specify Model

Model Selection Criterion

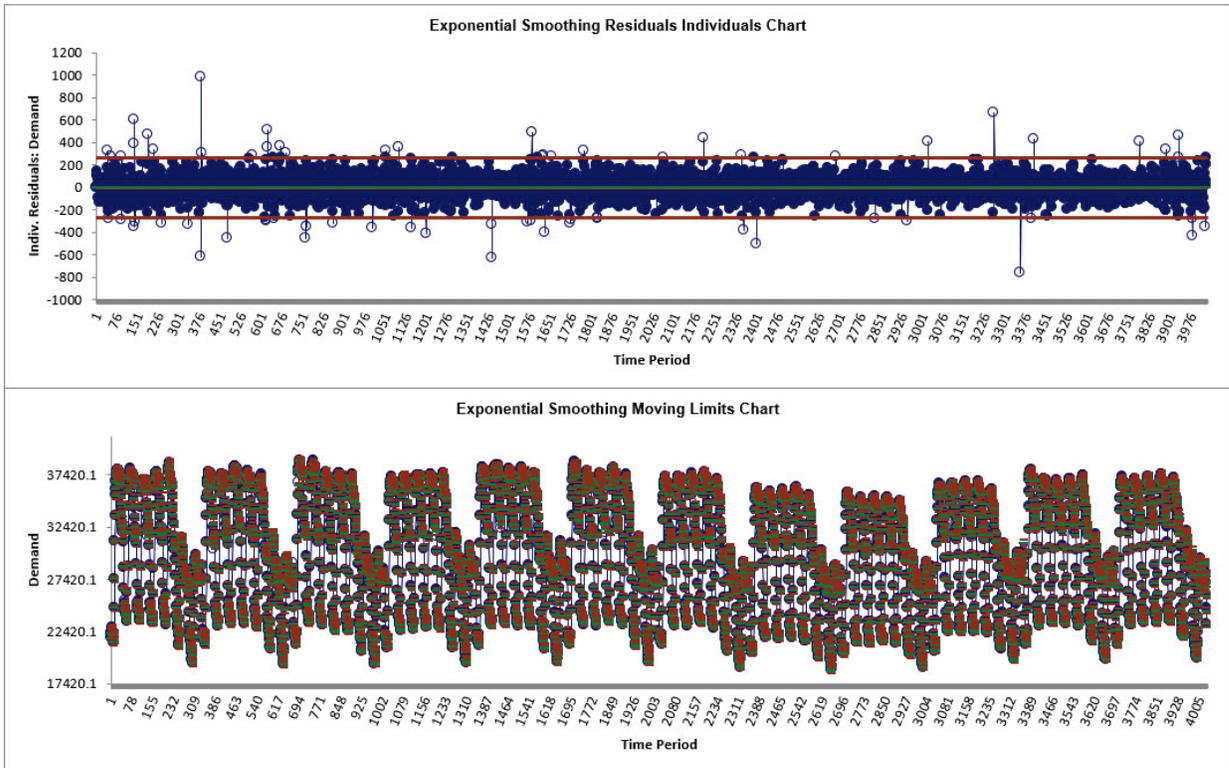
AICc - Akaike information criterion with small sample size correction

AIC - Akaike information criterion

BIC - Bayesian information criterion

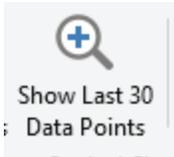
OK >>
Cancel
Help

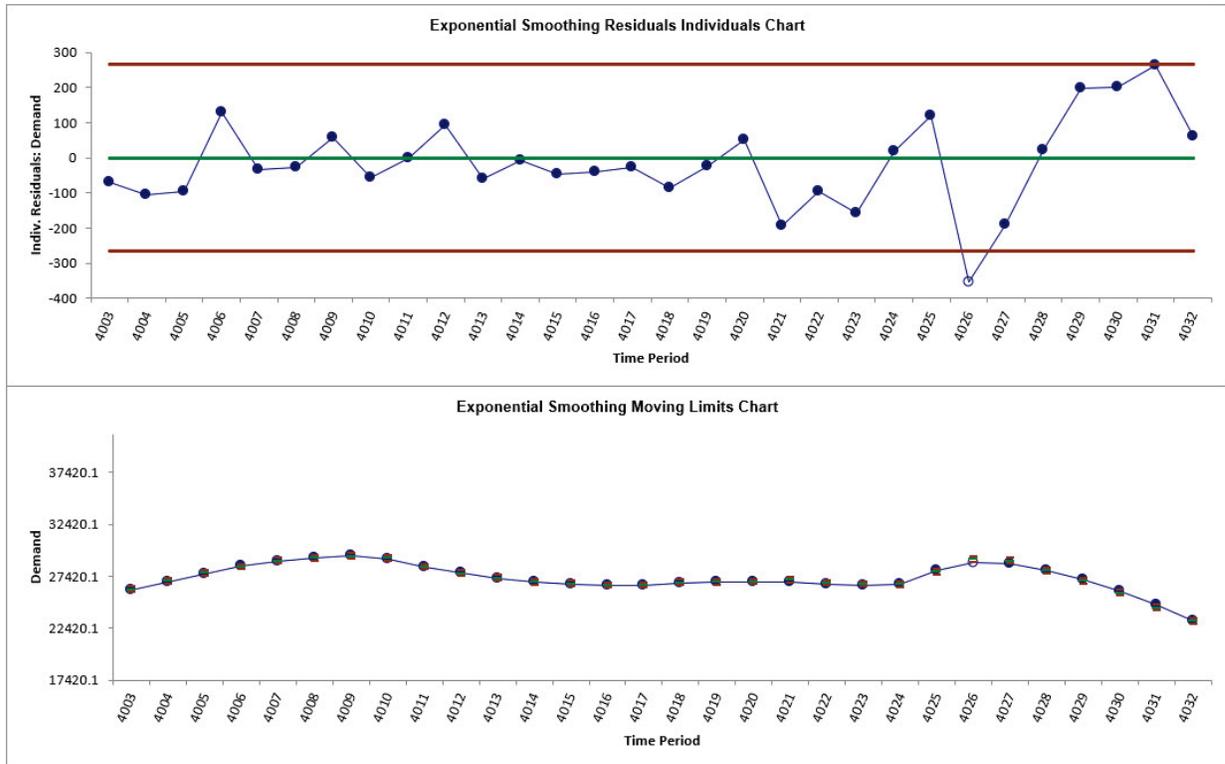
15. We will use the default **Automatic Model Selection** with **AICc** as the **Model Selection Criterion**. Click **OK** to return to the Exponential Smoothing MSD Control Chart dialog. Click **OK**. The exponential smoothing MSD control charts are produced:



Now we have a chart that can be used to identify assignable causes. The number of out-of-control signals have been dramatically reduced.

16. You can enable scrolling or zoom in to view the last 30 points on the Control Charts. Here we will do the latter. Click **SigmaXL Chart Tools > Show Last 30 Data Points**.

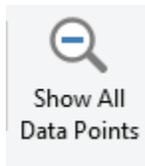




The out-of-control data point at 4026 is now clearly visible.

Tip: The Y axis scaling for the Moving Limits chart makes the same point difficult to see. You may wish to change the maximum and minimum values, right click **Format Axis**, and change the Bounds Minimum Maximum values.

- To restore the chart, click **SigmaXL Chart Tools > Show All Data Points**.



- Scroll down to view the Exponential Smoothing Model header:

Exponential Smoothing Model (Multiple Seasonal Decomposition): Simple Exponential Smoothing with Multiplicative Errors (M, N, N) - Model Automatically Selected
Model Periods: All observations are used in the Exponential Smoothing model estimation. No withhold periods available for out-of-sample forecast accuracy evaluation.

The model **Simple Exponential Smoothing with Multiplicative Errors (M, N, N)** was automatically selected as the best fit for the deseasonalized Demand data based on the AICc criterion. There are no withhold periods.

This is the same model used [previously](#) in Exponential Smoothing MSD Forecast, but that used a withhold sample of 96.

ARIMA Forecast

The Autoregressive Integrated Moving Average (ARIMA) model was developed by Box and Jenkins. The default automatic determination of the best model order in SigmaXL uses the stepwise method of Hyndman and Khandakar (see fpp2).

Stationarity, Differencing and Constant

ARIMA assumes that the time series is stationary, i.e., it has the property that the mean, variance and autocorrelation structure do not change over time. If a time series mean is not stationary (e.g. trending), this can be corrected by differencing, computing the differences between consecutive observations for nonseasonal and between consecutive periods for seasonal data (e.g., Jan 2019 – Jan 2018, etc.). For nonseasonal, this will typically involve 1 or 2 orders of differencing. This order is the Integrated term d . For seasonal, this will typically involve 1 order of differencing. This order is the Seasonal Integrated term D . A constant term c is optional:

- If $d + D = 0$, a constant term in the model is the mean.
- If $d + D = 1$, a constant term in the model is a trend (drift).
- If $d + D > 1$, a constant term would be a quadratic or higher trend, so constant should not be included.
- It is recommended that $d + D$ should not be > 3 .
- If the variance is not stationary, use a Box-Cox transformation.

Autoregressive (AR) Model

In an autoregressive model, we forecast the variable of interest using a linear combination of past values of the variable. The term autoregressive indicates that it is a regression of the variable against itself:

$$y_t = c + \phi_1 y_{t-1} + \phi_2 y_{t-2} + \dots + \phi_p y_{t-p} + \varepsilon_t$$

where ε_t is white noise. This is like a multiple regression but with lagged values of y_t as predictors. We refer to this as an AR(p) model, an autoregressive model of order p (fpp2).

Moving Average (MA) Model

Rather than using past values of the forecast variable in a regression, a moving average model uses past forecast errors in a regression-like model:

$$y_t = c + \varepsilon_t + \theta_1 \varepsilon_{t-1} + \theta_2 \varepsilon_{t-2} + \dots + \theta_q \varepsilon_{t-q}$$

We refer to this as an MA(q) model, a moving average model of order q (fpp2).

Autoregressive Integrated Moving Average (ARIMA) Model

If we combine differencing with autoregression and a moving average model, we obtain a nonseasonal ARIMA model:

$$y'_t = c + \phi_1 y'_{t-1} + \dots + \phi_p y'_{t-p} + \theta_1 \varepsilon_{t-1} + \dots + \theta_q \varepsilon_{t-q} + \varepsilon_t$$

where y'_t is the differenced series. This is the ARIMA (p, d, q) model, where p is the number of autoregressive terms, d is the degree of differencing and q is the number of moving average terms.

Seasonal ARIMA

For seasonal, the model consists of terms that are similar to the nonseasonal components of the model, but involve backshifts of the seasonal period. The seasonal model is ARIMA (P, D, Q) and combined we have ARIMA $(p, d, q) (P, D, Q)$.

ARIMA Model Order

Model order may be automatically determined or user specified. The **Stepwise Procedure** utilizes the stepwise method of Hyndman and Khandakar (see Appendix: [Automatic Model Selection](#)):

- Seasonal D (0 or 1) is determined using a Seasonal Strength test, or user specified.
- Nonseasonal d (0, 1, or 2) is determined using a modified KPSS unit root test, or user specified.
- AR (p), MA (q), with orders from 0 to a maximum of 5.
- Seasonal, SAR (P), SMA (Q), with orders from 0 to a maximum of 2.
- Constant (included or not included) if $d + D \leq 2$; not included if $d + D > 2$.

An **Extended Model Search** will search over the same range of order values as Stepwise but do so using all combinations, subject to the following constraints for consistency and computational efficiency:

- d and D are determined using the same methods as Stepwise, or specified by the user.
- Constant (included or not included) if $d + D \leq 2$; not included if $d + D > 2$.
- $p + q + P + Q \leq 7$.
- Computation time limit is specified by user.

In general, the default **Stepwise Procedure** is recommended over the **Extended Model Search**, as it is much faster and usually finds the best ARIMA model, or a simpler one that is close to the best ARIMA model.

Model Parameter Estimation and Missing Values

Model parameters are solved using nonlinear maximization of the Log-Likelihood function. Two general models are available - the conditional sum of squares (CSS) and the state space Kalman maximum likelihood. The CSS is always used for initial estimates and is used if $n > 500$ or seasonal frequency > 12 for computational efficiency. Kalman Filters permit exact calculations and can handle missing values. For CSS, if missing values are encountered, the largest contiguous range is used. For further details and formulas, see Appendix: [Autoregressive Integrated Moving Average - ARIMA](#).

ARIMA Model Statistics and Information Criteria for Model Comparison

The ARIMA model statistics are similar to those used in Exponential Smoothing. Log-Likelihood is related to $-\ln(\text{Sum-of-Squares Error})$, so is maximized. Information Criteria AICc, AIC and BIC are calculated using $-2 * \text{Log-Likelihood}$ and incorporate a penalty for the number of terms in the model, so smaller is better. These are used in automatic model selection. AICc is the default Information Criterion, based on forecast error performance with competition data. For further details, see [Information Criteria for Model Comparison](#).

Forecast Accuracy Metrics

The Forecast Accuracy Metrics for ARIMA are the same as those used in [Exponential Smoothing](#):

$$\text{Root mean squared error: RMSE} = \sqrt{\text{mean}(e_t^2)}$$

$$\text{Mean absolute error: MAE} = \text{mean}(|e_t|)$$

$$\text{Mean absolute percentage error: MAPE} = \text{mean}\left(\left|\frac{100e_t}{y_t}\right|\right)$$

$$\text{Mean absolute scaled error: MASE} = \text{mean}(|e_t|)/\text{scale}$$

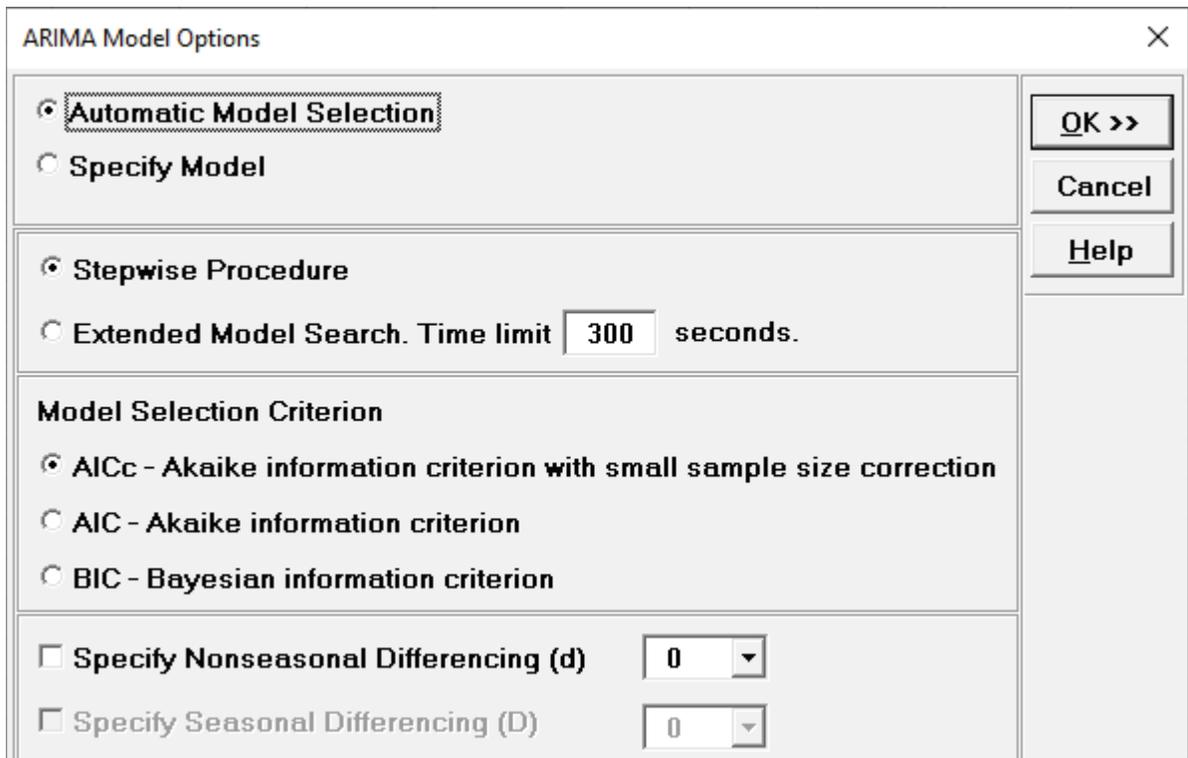
Chemical Process Concentration - Series A

1. Open **Chemical Process Concentration – Series A.xlsx (Sheet 1 tab)**. This is the Series A data from Box and Jenkins, a set of 197 concentration values from a chemical process taken at two-hour intervals. See the [Run Chart](#), [ACF/PACF Plots](#), [Spectral Density Plot](#) and [Seasonal Trend Decomposition Plots](#) for this data.
2. Click **SigmaXL > Time Series Forecasting > ARIMA Forecast > Forecast**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
3. Select *Concentration*, click **Numeric Time Series Data (Y) >>**. Check **Display ACF/PACF/LB Plots** and **Display Residual Plots**. Leave **Specify Model Periods**, **Seasonal Frequency** and **Box-Cox Transformation** unchecked. We will use the default **No. of Forecast Periods = 24** and **Prediction Interval = 95.0 %**.

- **Optional Time Axis Labels** will be displayed on the forecast chart time axis. If used, dates for the forecast periods should also be included, otherwise the time axis will be blank for the forecast periods.
- **No. of Forecast Periods** are the number of time series values to be predicted (forecast horizon). The most accurate forecast will be for the first predicted value (one-step-ahead).
- **Prediction Interval %** is the confidence level for the individual predictions. For example, a 95% prediction interval contains a range of values which should include the actual future value with 95% probability. The interval will get larger the further out you predict.

- **Model Options** opens another dialog which allows you to set automatic options or to specify a model.
- **Display ACF/PACF/LB** option will produce ACF and PACF plots for the raw data as well as for the model residuals. The LB plot is a plot of Ljung-Box test P-Values for various lags and is used to determine if a group of autocorrelations are significant, (i.e., the autocorrelations do not come from a white noise series). For further details, see [Ljung-Box Test](#).
- **Display Residual Plots** will produce a table of model residuals and the usual model residual plots: histogram, normal probability plot, residuals versus data order, and residuals vs forecast value. Note that if a Box-Cox transformation is applied, the residuals are transformed so will not be equal to forecast - actual.
- **Specify Model Periods** are used to specify a start period, end period or withhold sample. The withheld data is not used in model estimation, so this is very useful for model validation and comparison. This will be used in a later example.
- **Seasonal Frequency** and **Box-Cox Transformation** will be used in a later example.

4. Click **Model Options**.



- **Automatic Model Selection** will be used later. It is the default selection.
- **Stepwise Procedure** selects the use of the [stepwise method](#).
- **Extended Model Search** is described [above](#). The **Time limit** may need to be increased for seasonal models with high seasonal frequency and/or large number of observations.

- **Model Selection Criterion** is the information criterion metric to be used in automatic model selection. AICc is the default selection.
 - **Specify Nonseasonal Differencing (d)** = 0, 1, or 2, overrides the automatic nonseasonal differencing. This is useful to compare models for borderline cases that are “nearly nonstationary” (see Box and Jenkins).
 - **Specify Seasonal Differencing (D)** = 0 or 1, overrides the automatic seasonal differencing. It is greyed out here because **Seasonal Frequency** was unchecked in the previous dialog.
 - Clicking **OK** accepts the settings and returns you to the previous dialog. Clicking **Cancel** will cancel any changes and return you to the previous dialog.
5. Select **Specify Model**. Specify **Nonseasonal Order I – Integrated/Differencing (d)** = 1 and **MA – Moving Average (q)** = 1. Leave **Include Constant** unchecked.

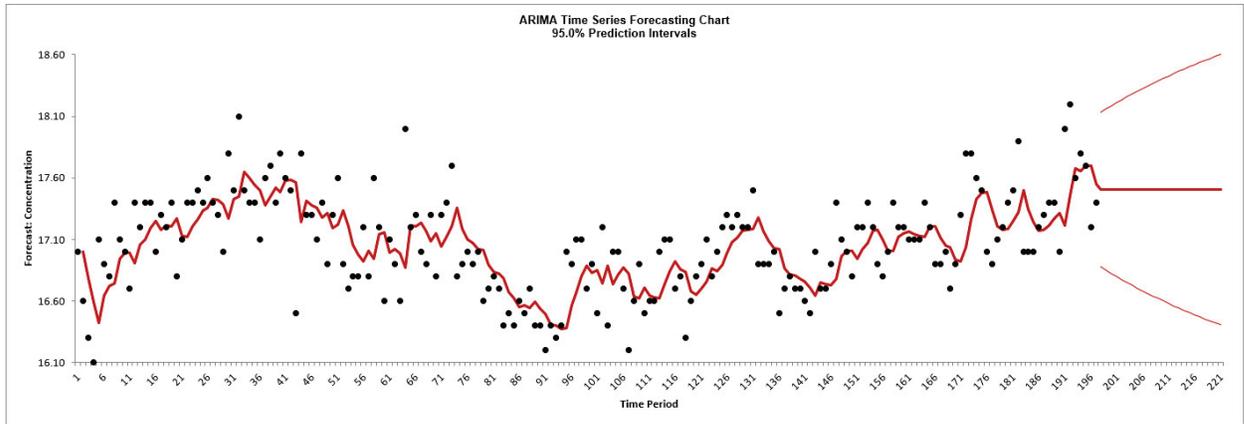
Nonseasonal Order		Seasonal Order	
AR - Autoregressive (p)	0	SAR - Seasonal Autoregressive (P)	0
I - Integrated/Differencing (d)	1	SI - Seasonal Integrated/Differencing (D)	0
MA - Moving Average (q)	1	SMA - Seasonal Moving Average (Q)	0

Include Constant (Mean if d + D = 0; Trend/Drift if d or D = 1)

- **Specify Model** allows you to manually specify the **Nonseasonal Order** and **Seasonal Order** values and option for **Include Constant (Mean if d + D = 0; Trend/Drift if d or D = 1)**. If d + D = 0, then the constant term is the mean; if d or D = 1, then the constant term is a Trend; if d + D > 1, then the constant term is quadratic or higher – this is not recommended.
- **Seasonal Order** is greyed out because **Seasonal Frequency** was unchecked in the previous dialog.

The specified ARIMA (0,1,1) is equivalent to a simple exponential smoothing model (with slight differences due to estimation of the initial value).

6. Click **OK** to return to the ARIMA Forecast dialog. Click **OK**. The ARIMA forecast report is given:



As expected, this is very similar to the [exponential smoothing forecast chart](#) that was produced using the **Simple Exponential Smoothing with Additive Errors (A, N, N) – Exponentially Weighted Moving Average (EWMA)** model. The initial in-sample predicted value for ARIMA is slightly different and starts at the second time period due to differencing.

7. Scroll down to view the ARIMA Model header:

ARIMA Model: Concentration - User Specified Model

Model Periods: All observations are used in the ARIMA model estimation. No withhold periods available for out-of-sample forecast accuracy evaluation.

If we had checked Specify Model Periods in the main dialog, the start, end or withhold selection would be summarized here as well.

8. The ARIMA Model Summary is given as:

ARIMA Model Summary	
AR Order (p)	0
I Order (d)	1
MA Order (q)	1
SAR Order (P)	0
SI Order (D)	0
SMA Order (Q)	0
Seasonal Frequency	1
Include Constant	0
No. of Predictors	0
Model Selection Criterion	Specified
Box-Cox Transformation	N/A
Lambda	
Threshold	

This is a summary of model information: ARIMA (0,1,1) with no constant and no predictors. Seasonal Frequency = 1 (nonseasonal); Model Selection Criterion = "Specified" because the model was user specified; and Box-Cox Transformation = "N/A" because Box-Cox Transformation was unchecked.

9. The Parameter Estimates are:

Parameter Estimates				
Term	Coefficient	SE Coefficient	T	P
MA_1	0.699386992	0.064505199	10.8423	0.0000

The MA_1 parameter coefficient value is approximately equal to $1 - \alpha = 1 - 0.2948 = 0.7052$ in [Exponential Smoothing Parameter Estimates](#). The slight difference is due to estimation of the initial value.

ARIMA Parameter Estimates include significance tests; P-Values < .05 are significant and highlighted in red. This may be useful for model refinement with multiple predictors (and will be demonstrated later). Note that for AR/MA model order selection, minimum AICc should be used, rather than significance tests (see [Kostenko, A.V. and Hyndman, R.J.](#)).

10. The ARIMA Model Statistics are:

ARIMA Model Statistics	
No. of Observations	197
DF	195
StDev	0.317609332
Variance	0.100875688
Log-Likelihood	-53.5087905
AICc	111.0797573
AIC	111.0175811
BIC	117.5738104

Degrees of freedom (DF) = $n - 2$ (1 term in the model, 1 order of differencing). See [ARIMA Model Statistics and Information Criteria for Model Comparison](#).

Comparing to the [Exponential Smoothing Model Statistics](#), we see that the StDev and Variance are approximately equal, but the Log-Likelihood, AICc, AIC and BIC are very different. This is due to different formulas being used in the Likelihood function. You cannot use Information Criteria to compare ARIMA and Exponential Smooth models to determine which model has the best fit.

11. The In-Sample Forecast Accuracy metrics are:

Forecast Accuracy		
Metric	In-Sample (Estimation) One-Step-Ahead Forecast	Out-of-Sample (Withhold) One-Step-Ahead Forecast
N	196	
RMSE	0.318324153	
MAE	0.248181659	
MAPE	1.45163544	
MASE	0.900807504	

MASE is less than one, so it is a better forecast than would be obtained from a naïve forecast (set all forecasts to be the value of the last observation). See [Forecast Accuracy Metrics](#).

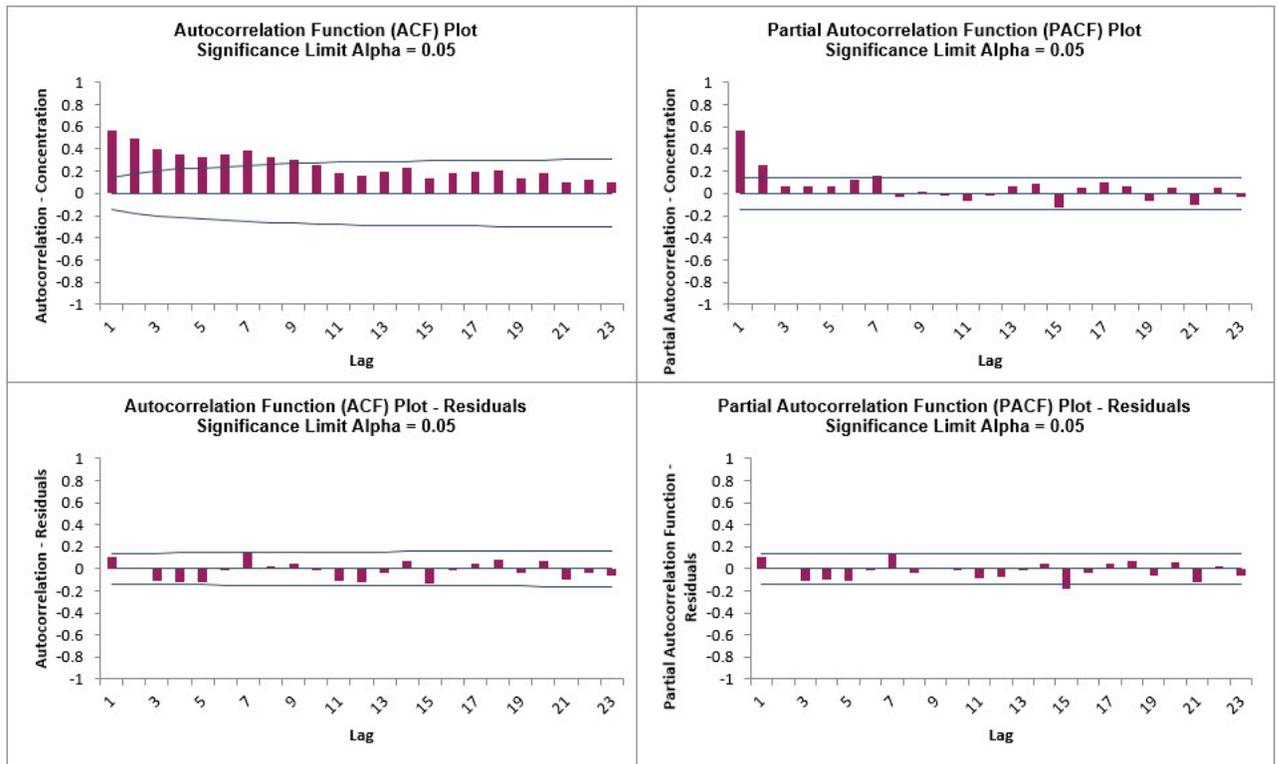
Comparing to the Exponential Smoothing [Exponential Smoothing In-Sample Forecast Accuracy metrics](#), as expected, we see the forecast errors are approximately the same. Note that we lose one observation (N=196) since we do not have a predicted value at time period = 1 due to differencing.

12. The Forecast Table is given as:

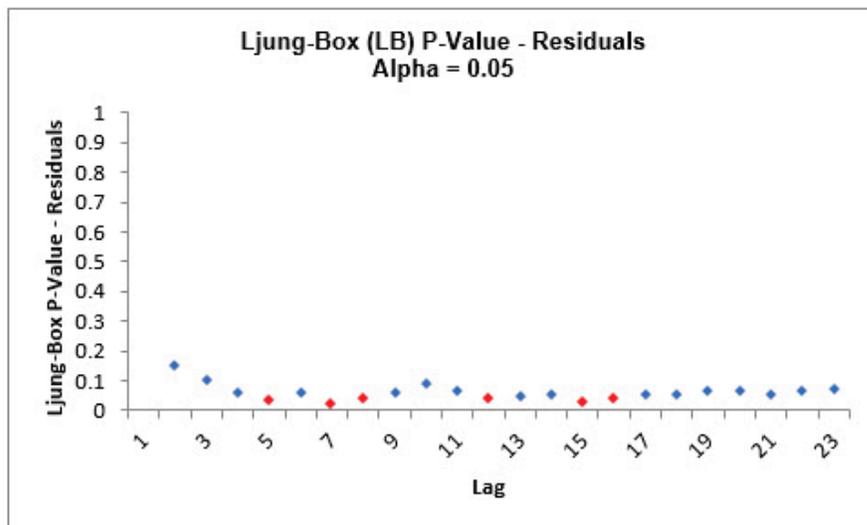
Forecast Table				
Period	Withhold Data	Forecast	Lower 95.0% PI	Upper 95.0% PI
198		17.50392441	16.87755395	18.13029487
199		17.50392441	16.84986402	18.1579848
200		17.50392441	16.82329967	18.18454915
201		17.50392441	16.79773386	18.21011496
202		17.50392441	16.77306181	18.23478701
203		17.50392441	16.74919586	18.25865296
204		17.50392441	16.72606181	18.28178701
205		17.50392441	16.70359618	18.30425264
206		17.50392441	16.68174418	18.32610464
207		17.50392441	16.66045813	18.34739069
208		17.50392441	16.63969619	18.36815263
209		17.50392441	16.61942146	18.38842736

These are the same forecast and prediction interval values displayed in the Forecast Chart but provided for further analysis or charting. If Withhold Periods are specified, the Withhold Data will be displayed as well.

13. Click on the **ARIMA ACF PACF LB** sheet to view the ACF/PACF/LB Plots:



These are approximately the same as what we obtained previously with [Exponential Smoothing ACF/PACF/LB Plots](#). We can see that all of the autocorrelation has been removed by the exponential smoothing model (with the exception of lag 15 in the PACF), so this is a good fit to the time series data.

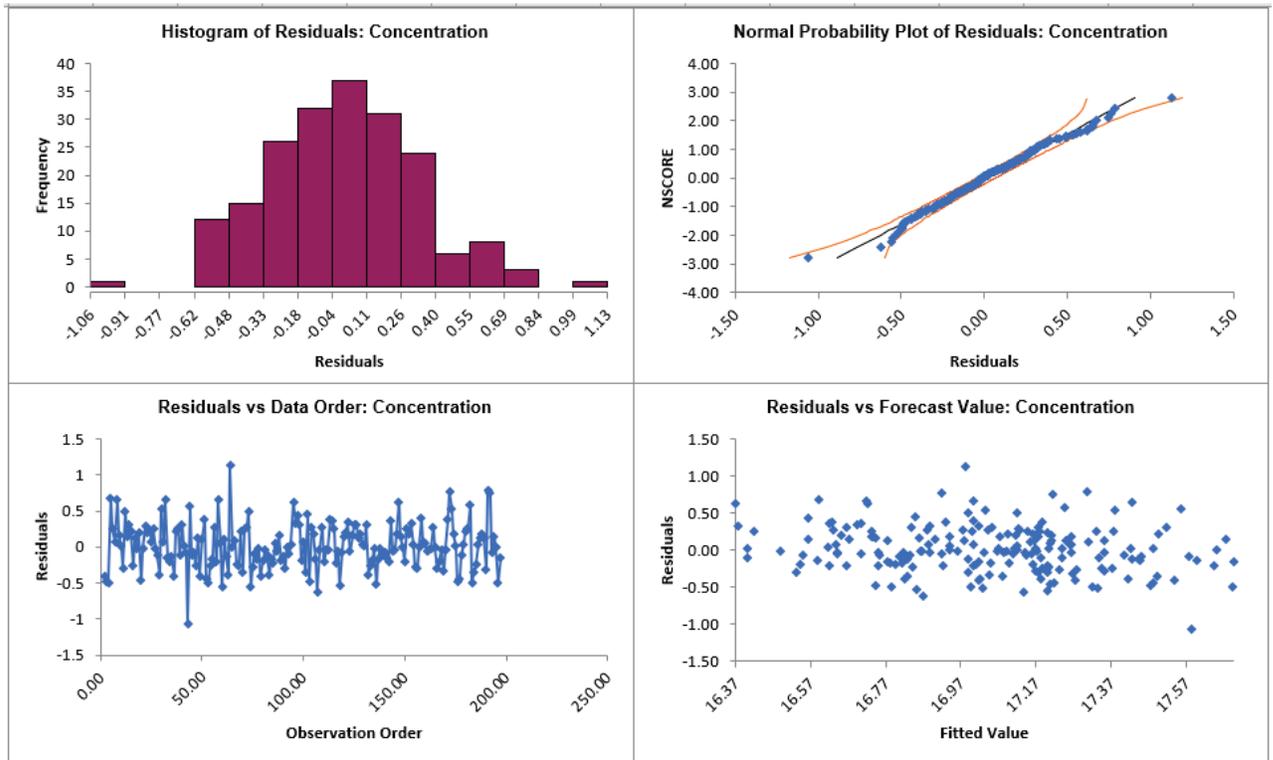


The LB plot is a plot of Ljung-Box test P-Values for various lags and is used to determine if a group of autocorrelations are significant, (i.e., the autocorrelations do not come from a white noise series). For further details, see [Ljung-Box Test](#).

The red P-Values are significant ($\alpha=.05$) and the blue P-Values are not significant. It is desirable that all P-Values be blue. The ACF/PACF plots indicated that almost all of the correlation has been accounted for in the model, but the Ljung-Box plot shows that some significant autocorrelation still remains - so the model can potentially be improved. This does not mean that the model is a bad model, it can still be very useful for prediction purposes, but the prediction intervals may not provide accurate coverage.

There does appear to be fewer significant P-Values than we obtained previously with the [Exponential Smoothing LB plot](#), but this may not be a practical difference, given the similarity of all the other statistics.

14. Click on the **ARIMA Residuals** sheet to view the Residual Plots:



The residuals are approximately normally distributed, with a roughly straight line on the normal probability plot. There are no obvious extreme outliers or patterns in the charts. Later, we will apply a control chart to the residuals to formally test for significant outliers or assignable causes.

These are approximately the same as what we obtained previously with [Exponential Smoothing Residual Plots](#).

15. Now we will rerun ARIMA Forecast on the Concentration data with Automatic Model Selection and Specify Withhold Periods. Click **Recall SigmaXL Dialog** menu or press **F3** to recall last dialog. Check **Specify Model Periods**. Set **Withhold Periods** = 24. We will use the default **Withhold Forecast Type: One-Step-Ahead with Prediction Interval at: Start of Withhold**.

ARIMA Forecast

Observation No.

Numeric Time Series Data (Y) >> Concentration

Optional Time Axis Labels >>

<< Remove

OK >>

Cancel

Help

No. of Forecast Periods 24

Prediction Interval 95.0 %

Specify Model Periods

Start Model at Period 1

Withhold Periods 24

End Model at Period

Model Options

Seasonal Frequency

Specify 12

Select 4-Quarterly

Automatically Detect

Display ACF/PACF/LB Plots

Display Residual Plots

Box-Cox Transformation

Rounded Lambda

Optimal Lambda

Lambda & Threshold (Shift)

Withhold Forecast Type:

One-Step-Ahead with Prediction Interval at: Start of Withhold

Include in Residuals

Multi-Step-Ahead with Prediction Interval at Start of Withhold.

- **Specify Model Periods** option allows you to specify the start and end periods used in automatic model identification and parameter estimation. Typically, **Start Model at Period** is kept = 1 and **Withhold Periods** specifies the number of periods to be withheld for out-of-sample testing. **End Model at Period** specifies the end period, so the withhold sample size would be: total number of observations – end period.
- **Withhold Forecast Type: One-Step-Ahead** will exclude the withhold sample from automatic model identification and parameter estimation, but uses the withhold data to update the predicted one-step ahead forecast. This is useful to assess forecast error when you only care about the short-term one-step ahead prediction.
- **Withhold Forecast Type: One-Step-Ahead with Prediction Interval at: Start of Withhold** will display the prediction interval for the duration of the withhold sample. Note that the length of the prediction interval is determined by the number of withhold periods, so overrides the specified **No. of Forecast Periods**.
- **Withhold Forecast Type: One-Step-Ahead with Prediction Interval at: End of Withhold** will display the prediction interval at the end of the withhold sample. The length of the prediction interval is determined by the specified **No. of Forecast Periods**.

- **Include in Residuals** will treat the one-step-ahead forecast errors as residuals (even though they were not part of the model estimation process) and will be included in the ACF/PACF/LB Residual Plots along with the Residuals report and graphs. Typically, this is kept unchecked.
- **Withhold Forecast Type: Multi-Step-Ahead with Prediction Interval at Start of Withhold** will exclude the withhold sample from automatic model identification and parameter estimation and does not use the withhold data to update the predicted one-step ahead forecast. This is useful to assess forecast error when you are interested in a long-term forecast window (horizon). The prediction interval will be displayed for the duration of the withhold sample. Note that the length of the prediction interval is determined by the number of withhold periods, so overrides the specified **No. of Forecast Periods**. These forecast errors are not included in ACF/PACF/LB Residual Plots or the Residuals report and graphs.

16. Click **Model Options**. Select **Automatic Model Selection**. We will use the defaults: **Stepwise Procedure** and **Model Selection Criterion: AICc – Akaike information criterion with small sample size correction**, leave **Specify Nonseasonal Differencing (d)** unchecked.

ARIMA Model Options

Automatic Model Selection

Specify Model

Stepwise Procedure

Extended Model Search. Time limit seconds.

Model Selection Criterion

AICc - Akaike information criterion with small sample size correction

AIC - Akaike information criterion

BIC - Bayesian information criterion

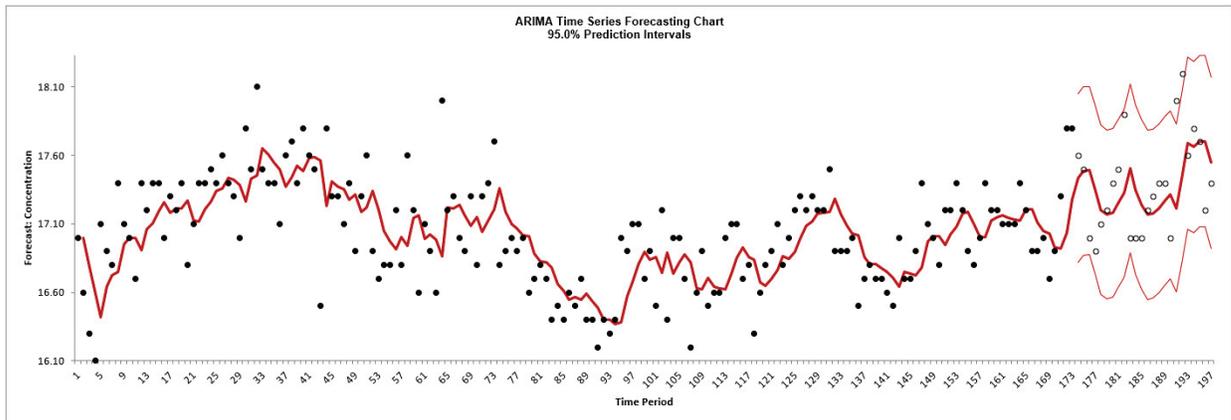
Specify Nonseasonal Differencing (d)

Specify Seasonal Differencing (D)

OK >>
Cancel
Help

Tip: When using **Recall SigmaXL Dialog** and if there are no changes to the **Model Option** settings, the previous settings will be used. It is not necessary to repeat this step.

17. Click **OK** to return to the ARIMA Forecast dialog. Click **OK**. The ARIMA forecast report is given:



The blank dots are the data values in the withhold sample with a one-step-ahead forecast and prediction intervals displayed at the start of the withhold sample.

This is very similar to the [exponential smoothing forecast chart](#) that was produced using the **Simple Exponential Smoothing with Multiplicative Errors (M, N, N)** model. The initial in-sample predicted value for ARIMA is slightly different and starts at the second time period due to differencing.

18. Scroll down to view the ARIMA Model header:

ARIMA Model: Concentration - Model Automatically Selected
Model Periods: Model parameter estimates calculated excluding 24 withhold periods.

19. The ARIMA Model Summary is given as:

ARIMA Model Summary	
AR Order (p)	0
I Order (d)	1
MA Order (q)	1
SAR Order (P)	0
SI Order (D)	0
SMA Order (Q)	0
Seasonal Frequency	1
Include Constant	0
No. of Predictors	0
Model Selection Criterion	AICc
Box-Cox Transformation	N/A
Lambda	
Threshold	

The ARIMA (0,1,1) model that we manually specified above, was also automatically selected based on the AICc criterion.

20. The Parameter Estimates are:

Parameter Estimates				
Term	Coefficient	SE Coefficient	T	P
MA_1	0.690028724	0.065305747	10.56613	0.0000

This is close to the parameter estimate obtained [above](#) (which used all of the data).

21. The ARIMA Model Statistics are:

ARIMA Model Statistics	
No. of Observations	173
DF	171
StDev	0.312472686
Variance	0.09763918
Log-Likelihood	-44.18071519
AICc	92.4324363
AIC	92.36143039
BIC	98.65641934

These are fairly close to the model statistics obtained [above](#) using all of the data. Here we are using only 173 of the 197 observations.

22. The Forecast Accuracy metrics are:

Forecast Accuracy		
Metric	In-Sample (Estimation) One-Step-Ahead Forecast	Out-of-Sample (Withhold) One-Step-Ahead Forecast
N	172	24
RMSE	0.313295746	0.352257753
MAE	0.243952039	0.273585123
MAPE	1.431266547	1.568577996
MASE	0.881507368	0.988584899

As expected, the **Out-of-Sample (Withhold) One-Step-Ahead Forecast** errors are larger than the **In-Sample (Estimation) One-Step-Ahead Forecast** errors.

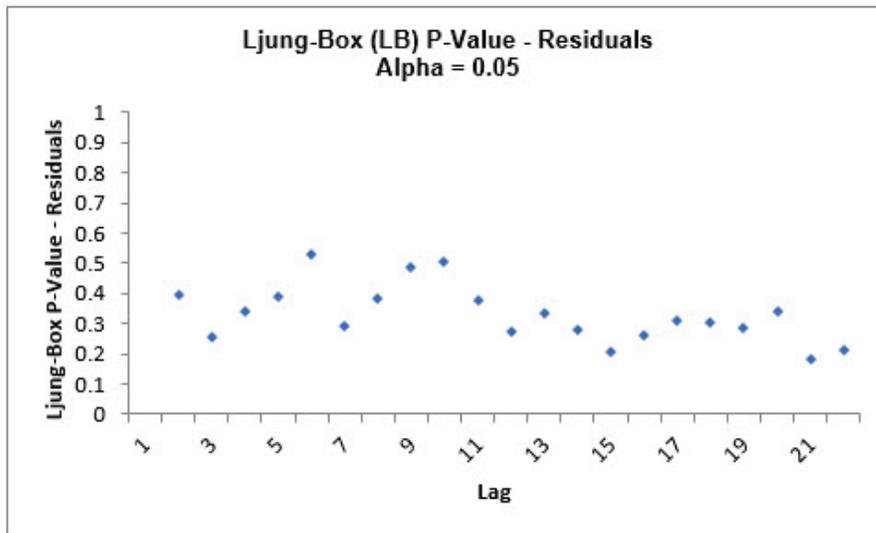
These are very similar to the [Exponential Smooth Forecast Accuracy Metrics](#) that were produced using the **Simple Exponential Smoothing with Multiplicative Errors (M, N, N)** model. Note that we lose one observation on the In-Sample (N=172) since we do not have a predicted value at time period = 1 due to differencing.

23. The Forecast Table is given as:

Forecast Table				
Period	Withhold Data	One-Step-Ahead Forecast	Lower 95.0% PI	Upper 95.0% PI
174	17.6	17.43699636	16.82022147	18.05377125
175	17.5	17.48752281	16.87207955	18.10296607
176	17	17.49139038	16.87775042	18.10503034
177	16.9	17.33907348	16.72282244	17.95532451
178	17.1	17.20297331	16.58502829	17.82091833
179	17.2	17.17105454	16.55470128	17.78740781
180	17.4	17.1800268	16.56542588	17.79462772
181	17.5	17.24821218	16.63450051	17.86192384
182	17.9	17.32625917	16.71316165	17.93935668
183	17	17.50410235	16.88693756	18.12126714
184	17	17.3478451	16.72798565	17.96770455
185	17	17.24002311	16.61980167	17.86024455
186	17.2	17.16562284	16.54613085	17.78511483
187	17.3	17.17627877	16.55847435	17.79408319
188	17.4	17.2146288	16.59825692	17.83100068
189	17.4	17.27208855	16.6568128	17.8873643
190	17	17.31173742	16.69785289	17.92562196
191	18	17.21510778	16.60123859	17.82897696
192	18.2	17.45840182	16.83590796	18.08089568
193	17.6	17.68827595	17.05843876	18.31811315
194	17.8	17.66091294	17.03262123	18.28920466
195	17.7	17.70402594	17.07708046	18.33097141
196	17.2	17.70277801	17.07747894	18.32807709
197	17.4	17.54693127	16.91921585	18.17464669

These are the same forecast and prediction interval values displayed in the Forecast Chart, but provided for further analysis or charting. The **Withhold Data** is also displayed.

24. The ACF/PACF/LB Residual Plots and Residual Plots are based on the in-sample data. The plots look similar to the complete data [above](#), except for the Ljung-Box P-Values:



As was the case with [exponential smoothing](#), the ARIMA (0,1,1) model is a better fit to the subset than the complete data, with all P-Values being blue ($> .05$).

25. If **Include in Residuals** was checked then the residuals would also include the **Out-of-Sample (Withhold) One-Step-Ahead Forecast** errors.
26. The above analysis can be rerun using **Withhold Forecast Type: Multi-Step-Ahead with Prediction Interval at Start of Withhold** (click **Recall SigmaXL Dialog** menu or press **F3** to recall last dialog), but we will not do so here. The results would be very similar to those obtained with [exponential smoothing](#).

Monthly Airline Passengers - Series G

27. Open **Monthly Airline Passengers - Series G.xlsx (Sheet 1 tab)**. This is the Series G data from Box and Jenkins, monthly total international airline passengers for January 1949 to December 1960. See the [Run Chart](#), [ACF/PACF Plots](#), [Spectral Density Plot](#) and [Seasonal Trend Decomposition Plots](#) for this data.
28. Click **SigmaXL > Time Series Forecasting > ARIMA Forecast > Forecast**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
29. Select *Monthly Airline Passengers*, click **Numeric Time Series Data (Y) >>**. Select *Date*, click **Optional Time Axis Labels >>**. Check **Display ACF/PACF/LB Plots** and **Display Residual Plots**. Check **Specify Model Periods**. Set **Withhold Periods = 24**. Select **Withhold Forecast Type: Multi-Step-Ahead with Prediction Interval at Start of Withhold**. Check **Seasonal Frequency** with **Specify = 12**. Check **Box-Cox Transformation** and select **Rounded Lambda** (selected because the [Run Chart](#) showed an increase in the seasonal variance over time). We will use the default **Prediction Interval = 95.0 %**.

ARIMA Forecast

Obs. No.
Ln (Airline Passengers)

Numeric Time Series Data (Y) >> Monthly Airline Passenger

Optional Time Axis Labels >> Date

<< Remove

OK >>
Cancel
Help

No. of Forecast Periods 24

Prediction Interval 95.0 %

Model Options

Display ACF/PACF/LB Plots

Display Residual Plots

Specify Model Periods

Seasonal Frequency

Box-Cox Transformation

Start Model at Period 1

Withhold Periods 24

End Model at Period

Specify 12

Select 4 - Quarterly

Automatically Detect

Rounded Lambda

Optimal Lambda

Lambda & Threshold (Shift)

Withhold Forecast Type:

One-Step-Ahead with Prediction Interval at: Start of Withhold

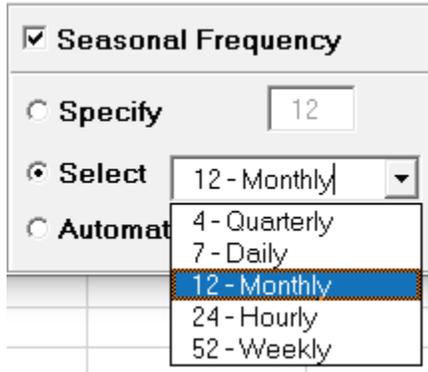
Include in Residuals

Multi-Step-Ahead with Prediction Interval at Start of Withhold.

- **Withhold Forecast Type: Multi-Step-Ahead with Prediction Interval at Start of Withhold** will exclude the withhold sample from automatic model identification and parameter estimation and does not use the withhold data to update the predicted one-step ahead forecast. This is useful to assess forecast error when you are interested in a long-term forecast window

(horizon). The prediction interval will be displayed for the duration of the withhold sample. Note that the length of the prediction interval is determined by the number of withhold periods, so overrides the specified **No. of Forecast Periods**.

- **Seasonal Frequency Specify** is used to specify the seasonal frequency.
- **Seasonal Frequency Select** gives a drop-down list of commonly used seasonal frequencies:



- **Seasonal Frequency Automatically Detect** should be used if uncertain what the seasonal frequency value is (or do a Spectral Density Plot prior to the Seasonal Trend Decomposition Plots).
- **Box-Cox Transformation with Rounded Lambda** will select $\Lambda = 0$ (Ln), 0.5 (SQRT) or 1 (Untransformed). Threshold (Shift) is computed automatically if the time series data includes 0 or negative values, otherwise it is 0.
- **Box-Cox Transformation with Optimal Lambda** uses the range of 0 to 1 for Lambda. Threshold is computed automatically if the time series data includes 0 or negative values.
- **Box-Cox Transformation with Lambda & Threshold (Shift)** if left blank, will compute optimal lambda and threshold. The user may also specify Lambda and Threshold. Lambda may be specified outside of the 0 to 1 range, but practically for time series analysis, should be limited to -1 to 2. Threshold is typically 0, but if the time series data includes 0 or negative values, a negative threshold value should be entered that is smaller than the minimum data value. This value will be subtracted from the data resulting in positive time series values.

30. Click **Model Options**. Select **Automatic Model Selection**. We will use the defaults: **Stepwise Procedure** and **Model Selection Criterion: AICc – Akaike information criterion with small sample size correction**, leave **Specify Nonseasonal Differencing (d)** and **Specify Seasonal Differencing (D)** unchecked.

ARIMA Model Options

Automatic Model Selection

Specify Model

Stepwise Procedure

Extended Model Search. Time limit seconds.

Model Selection Criterion

AICc - Akaike information criterion with small sample size correction

AIC - Akaike information criterion

BIC - Bayesian information criterion

Specify Nonseasonal Differencing (d)

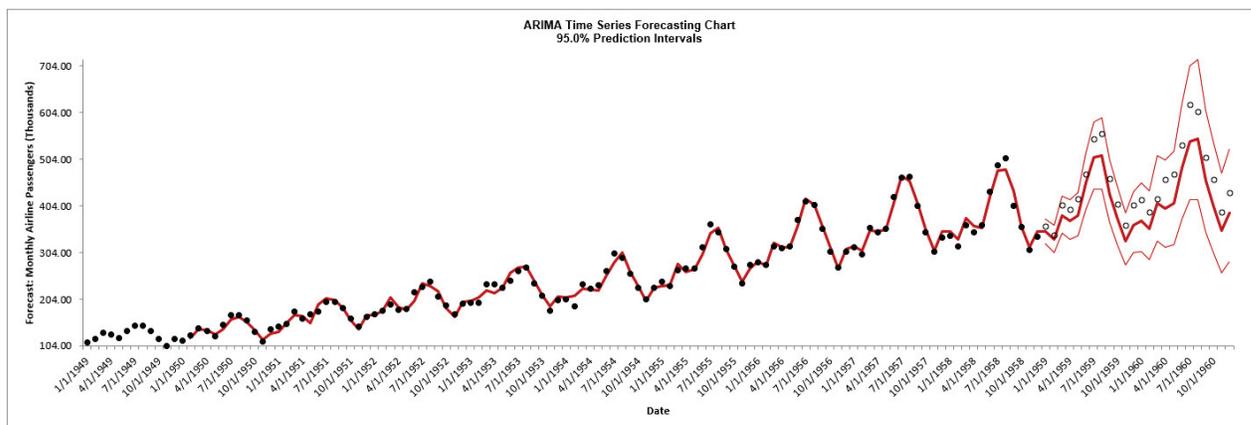
Specify Seasonal Differencing (D)

OK >>

Cancel

Help

31. Click **OK** to return to the ARIMA Forecast dialog. Click **OK**. The ARIMA forecast report is given:



The blank dots are the data values in the withhold sample with a multi-step forecast and prediction intervals displayed at the start of the withhold sample. Note that the data has been transformed using Box-Cox, but the inverse transformation is applied to produce this chart with the original units.

This is similar to the [exponential smoothing forecast chart](#) that was produced using the **Additive Trend, Additive Seasonal Method with Additive Errors (Holt-Winters) (A, A, A)**

model. The initial in-sample predicted value for ARIMA starts at the 14th time period due to nonseasonal and seasonal differencing.

32. Scroll down to view the ARIMA Model header:

ARIMA Model: Monthly Airline Passengers (Thousands) - Model Automatically Selected
Model Periods: Model parameter estimates calculated excluding 24 withhold periods.

33. The ARIMA Model Summary is given as:

ARIMA Model Summary	
AR Order (p)	0
I Order (d)	1
MA Order (q)	1
SAR Order (P)	0
SI Order (D)	1
SMA Order (Q)	1
Seasonal Frequency	12
Include Constant	0
No. of Predictors	0
Model Selection Criterion	AICc
Box-Cox Transformation	Rounded Lambda
Lambda	0
Threshold	0

The ARIMA (0,1,1) (0,1,1) model (with no constant) was automatically selected based on the AICc criterion. This agrees with the results shown in the [ACF/PACF Plots](#) of the manually differenced data and is the model used in Box and Jenkins, Chapter 9., “Analysis of Seasonal Time Series”.

There are no predictors. Seasonal Frequency = 12; Model Selection Criterion = “AICc” and Box-Cox Transformation = “Rounded Lambda” with Lambda = 0 (i.e., a Ln transformation).

34. The Parameter Estimates are:

Parameter Estimates				
Term	Coefficient	SE Coefficient	T	P
MA_1	0.342313249	0.100902427	3.392517	0.0010
SMA_1	0.540469465	0.087677292	6.164304	0.0000

Unlike the previous example with nonseasonal and nontrended Concentration data, we cannot compare these parameters to the exponential smooth model.

35. The ARIMA Model Statistics are:

ARIMA Model Statistics	
No. of Observations	120
DF	105
StDev	0.037414431
Variance	0.00139984
Log-Likelihood	197.5047754
AICc	-388.776541
AIC	-389.009551
BIC	-380.991064

- The number of observations, $n = 144 - 24$ (*withhold*) = 120
- Degrees of freedom (DF) = $120 (n) - 15$ (2 terms in the model, 1 nonseasonal difference, 12 seasonal differences) = 105
- Note that the model statistics are based on the Ln transformed data, not the original data.

Comparing to the [Exponential Smoothing Model Statistics](#), we see that the StDev and Variance are approximately equal, but the Log-Likelihood, AICc, AIC and BIC are very different. This is due to different formulas being used in the Likelihood function. You cannot use Information Criteria to compare ARIMA and Exponential Smooth models to determine which model has the best fit.

36. The Forecast Accuracy metrics are:

Forecast Accuracy		
Metric	In-Sample (Estimation) One-Step-Ahead Forecast	Out-of-Sample (Withhold) Multi-Step-Ahead Forecast
N	107	24
RMSE	9.44337388	43.18833723
MAE	7.384036891	39.45185993
MAPE	3.001604734	8.51734062
MASE	0.258417364	1.380687256

Comparing to [Exponential Smoothing Forecast Accuracy](#), we have higher Out-of-Sample forecast errors. MASE is greater than one, so it is a worse forecast than would be obtained from a seasonal naïve forecast. However, it is important to keep in mind that we are forecasting out for 24 months, so this is still quite reasonable. Selecting a 12 month Multi-Step Ahead Withhold Sample (or a One-Step-Ahead Withhold Sample) would dramatically improve the forecast.

As expected, the **Out-of Sample (Withhold) Multi-Step-Ahead Forecast** errors are larger than the **In-Sample (Estimation) One-Step-Ahead Forecast** errors. Note, if we were primarily interested in a short term one-step ahead forecast, then we would have selected **Withhold Forecast Type: One-Step-Ahead** and this table would show **Out-of Sample (Withhold) One-Step-Ahead Forecast** errors.

Note that we lose 13 observations on the In-Sample (N=107) since we do not have predicted values at time period = 1 to 13, due to differencing.

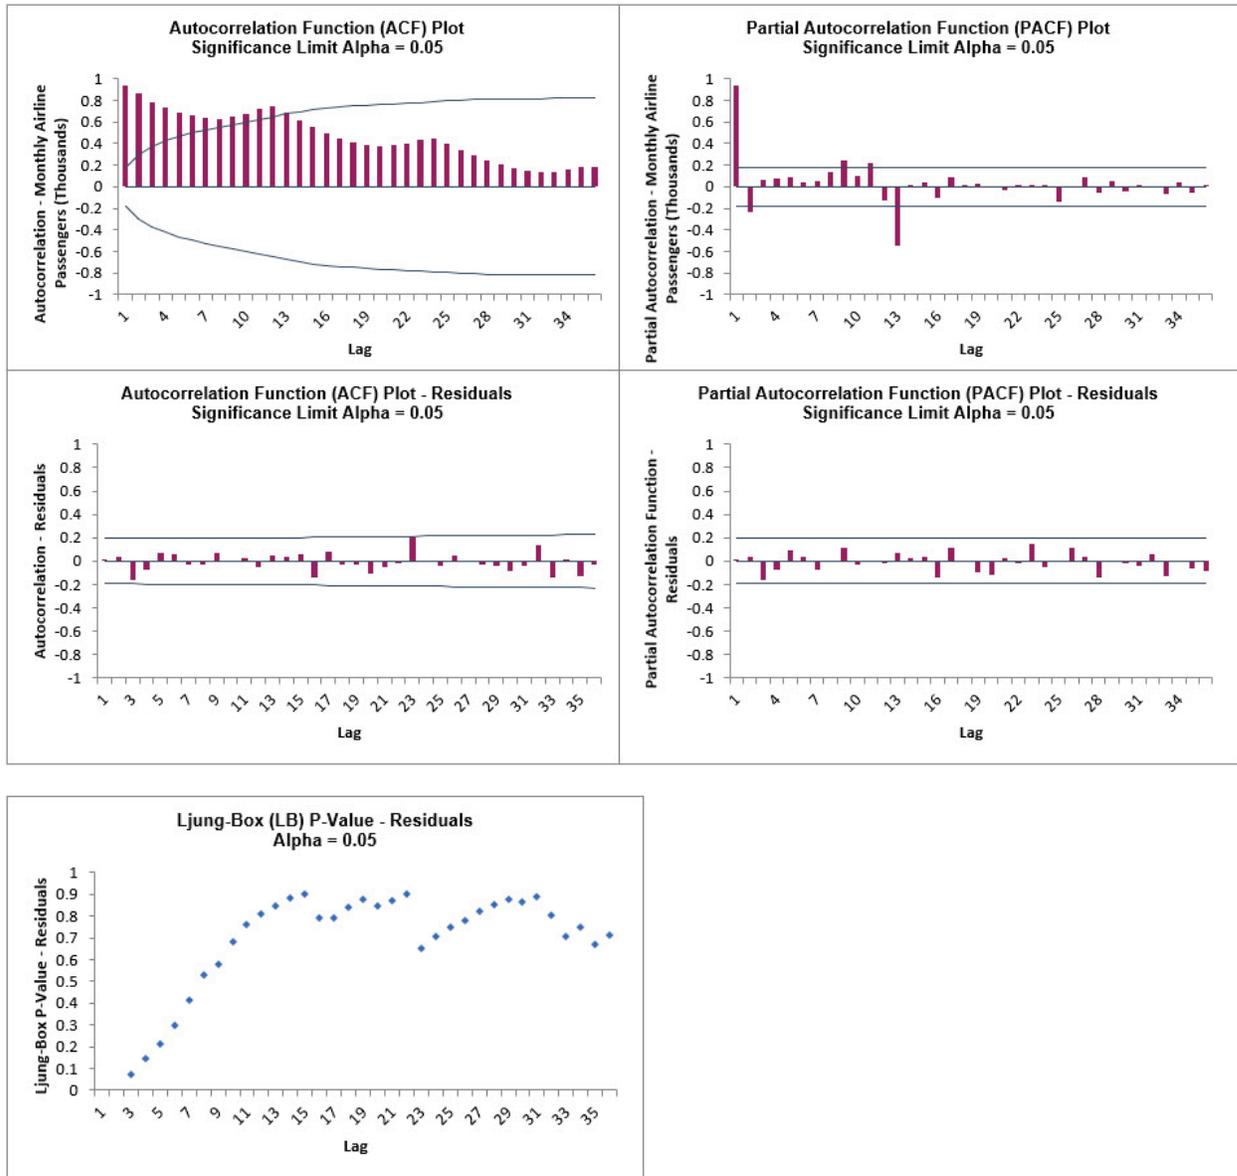
Forecast Accuracy metrics are calculated using the actual raw data versus inverse transformed forecast as displayed in the Forecast Chart and Table, so allow comparison across all model types and transformations.

37. The Forecast Table is given as:

Forecast Table				
Period	Withhold Data	Multi-Step Ahead Forecast	Lower 95.0% PI	Upper 95.0% PI
121	360	348.5828062	323.6916935	375.3879856
122	342	331.5740082	303.4388404	362.3178982
123	406	383.1608225	346.2907533	423.9565003
124	396	372.606915	333.0281012	416.8894835
125	420	382.9486611	338.8267505	432.8161127
126	472	453.1350644	397.199074	516.9483013
127	548	507.198003	440.7280155	583.6929018
128	559	511.2341986	440.6029023	593.1881166
129	463	429.3989001	367.2058264	502.1255171
130	407	376.442152	319.5422535	443.4740391
131	362	328.1517697	276.5830729	389.3354095
132	405	362.9392672	303.8283644	433.5504089
133	417	372.7865814	306.4056918	453.5484783
134	391	354.596781	287.8353252	436.8431047
135	419	409.7655151	328.7189301	510.7943655
136	461	398.4787992	316.1075385	502.3143522
137	472	409.5386223	321.4319641	521.7959068
138	535	484.5984041	376.4753418	623.7742214
139	622	542.4151917	417.2724015	705.0891434
140	606	546.7316396	416.6289379	717.4621313
141	508	459.2141241	346.749965	608.1546736
142	461	402.5803351	301.3041689	537.8980544
143	390	350.9369201	260.4037852	472.9452062
144	432	388.1398803	285.6122259	527.4724013

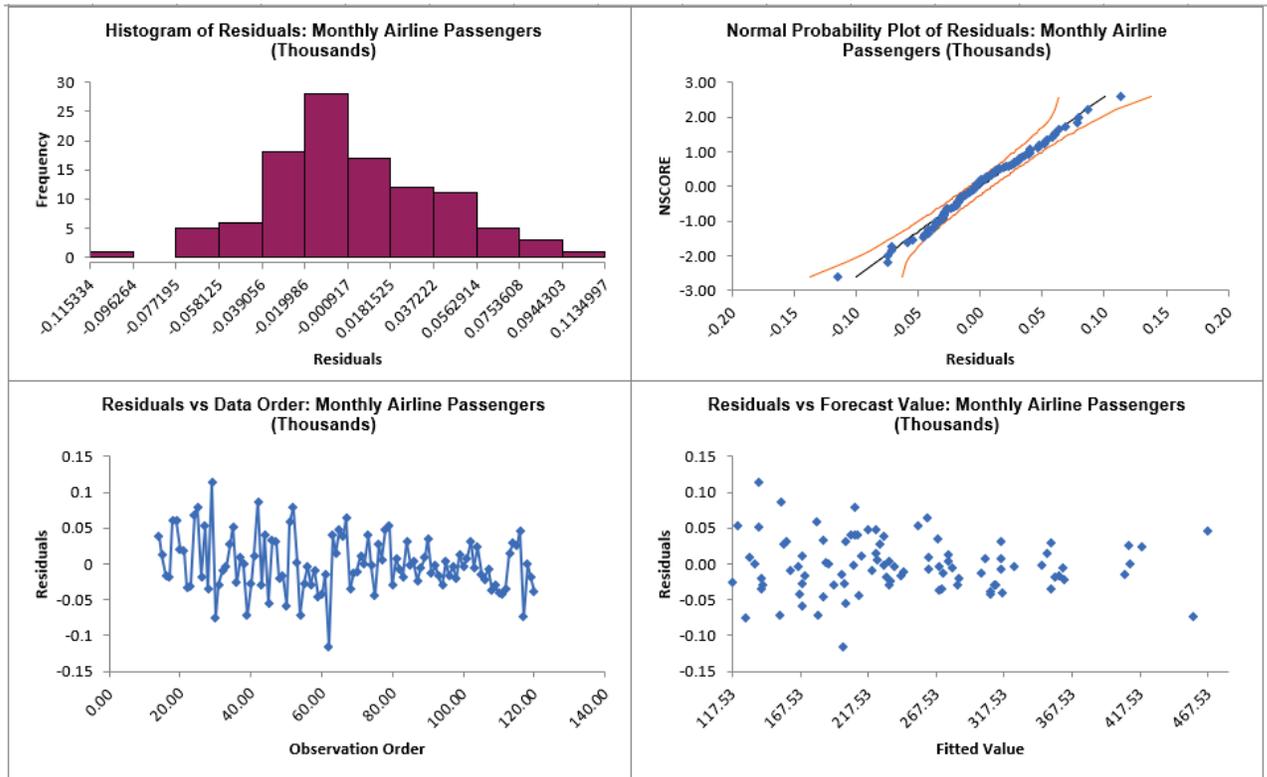
These are the same forecast and prediction interval values displayed in the Forecast Chart but provided for further analysis or charting. Note that this table gives the Forecast Period number whereas the Chart displays the Date.

38. Click on the **ARIMA ACF PACF LB** sheet to view the ACF/PACF/LB Plots:



The ACF/PACF Residuals Plots indicate that the autocorrelation has been accounted for in the model. The Ljung-Box plot confirms that no significant autocorrelation remains (all blue P-Values $> .05$), so this is a good fit to the time series data (as compared to the [Exponential Smoothing LB Plot](#)).

39. Click on the **ARIMA Residuals** sheet to view the Residual Plots:



The residuals are approximately normally distributed, with a roughly straight line on the normal probability plot. There are no obvious extreme outliers or patterns in the charts.

Note that the residuals are based on the Ln transformed data, not the original data. The **ARIMA Model Information** to the right of the plots shows the Box-Cox Transformation information.

ARIMA Model Summary	
AR Order (p)	0
I Order (d)	1
MA Order (q)	1
SAR Order (P)	0
SI Order (D)	1
SMA Order (Q)	1
Seasonal Frequency	12
Include Constant	0
No. of Predictors	0
Model Selection Criterion	AICc
Box-Cox Transformation	Rounded Lambda
Lambda	0
Threshold	0

ARIMA Forecast with Predictors

The ARIMA model supports continuous or categorical predictors, similar to multiple regression in SigmaXL. For further details see Appendix: [ARIMA with Covariates \(Predictors\)](#).

In order to provide a forecast, additional predictor (X) values must be added to the dataset prior to running the analysis. The number of forecast periods will be equal to the number of additional predictor rows. Alternatively, the predictor values from a withhold sample may be used. If neither are provided, SigmaXL will use the last row of predictor values and compute one forecast period.

Continuous predictors are numeric, categorical predictors can be text or numeric. All predictors must have the same number of rows. A predictor cannot have all values the same.

As with multiple linear regression, predictors should not be strongly correlated. SigmaXL will automatically remove terms with very high variance inflation factors ($VIF > 100$) and give a warning message “Multicollinearity detected in predictors. The following predictors were removed...”.

Also as done in multiple linear regression, categorical predictors are automatically “dummy coded”. Predictors will have a level append to their term names in the Parameter Estimates table. The first level (sorted alphanumerically) is a hidden reference level. If there are three levels, only two will appear in the Parameter Estimates table.

Missing values, while permitted in the Time Series Data (Y), are not permitted in the predictors. A warning message will be given, “Missing values detected in predictors. The following predictors were removed...”. If all predictors have missing values an error message is returned. An exception is made for missing values in the first rows to accommodate indicators with lags.

Sometimes the impact of a predictor will not be simple and immediate. For example, an advertising campaign may impact sales for some time beyond the end of the campaign, and sales in one month will depend on the advertising expenditure in each of the past few months. In these situations, we need to allow for lagged effects of the predictor (Hyndman, fpp2, 9.6 Lagged Predictors). A Pre-Whitened CCF Plot will show which lags need to be included in the model. SigmaXL can accommodate lagged predictors, but they must be manually created using **SigmaXL > Utilities > Lag Data** (or simply: copy, shift down one row, paste new column, delete extra row at bottom, add column header for the predictor data, and include as a predictor in the ARIMA model). Note, Box and Jenkins, Ch. 11, describe a more complex model method called a Transfer Function, but this is not provided in SigmaXL.

Daily Electricity Demand with Predictors – ElecDaily

1. Open **Daily Electricity Demand with Predictors – ElecDaily.xlsx**. This is daily electricity demand (GW) for the state of Victoria, Australia, every day during 2014 (see Hyndman, fpp2, Section 9.3, **Example: Forecasting electricity demand**, <https://otexts.com/fpp2/forecasting.html>). Temp (C) is the maximum daily temperature in degrees Celsius for the city of Melbourne. TempSq is Temperature squared. WorkDay takes on the value 1 on work days and 0 otherwise. This data has a seasonal frequency = 7 (observations per week). See the [Run Chart](#), [ACF/PACF Plots](#), [Spectral Density Plot](#) and [Seasonal Trend Decomposition Plots](#) for this data.
2. Click the **Forecast 2 Weeks** Sheet tab. Following the referenced example, we will use the ARIMA model to forecast 14 days ahead starting from January 1, 2015 (a non-work-day being a public holiday for New Year’s Day). We could obtain weather forecasts for those 14 days, but for the sake of illustration, we will set the temperature for the 14 days to a constant 26 degrees and TempSq to 676. Scroll down to view the added data. Note that Date, Temp (C), TempSq and WorkDay are added for these 14 days, but the Demand values are blank:

1	Date	Demand	Temp (C)	TempSq	WorkDay
366	12/31/2014	186.370181	25.5	650.25	1
367	1/1/2015		26	676	0
368	1/2/2015		26	676	1
369	1/3/2015		26	676	0
370	1/4/2015		26	676	0
371	1/5/2015		26	676	1
372	1/6/2015		26	676	1
373	1/7/2015		26	676	1
374	1/8/2015		26	676	1
375	1/9/2015		26	676	1
376	1/10/2015		26	676	0
377	1/11/2015		26	676	0
378	1/12/2015		26	676	1
379	1/13/2015		26	676	1
380	1/14/2015		26	676	1

3. Click **SigmaXL > Time Series Forecasting > ARIMA Forecast > Forecast with Predictors**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.

4. Select *Demand*, click **Numeric Time Series Data (Y) >>**; select *Date*, click **Optional X-Axis Labels >>**; select *Temp (C)* and *TempSq*, click **Optional Continuous Pred. >>**; select *WorkDay*, click **Optional Categorical Pred. >>**. Check **Display ACF/PACF/LB Plots** and **Display Residual Plots**. Check **Seasonal Frequency** with **Select = 7 - Daily** (or **Specify = 7**). Leave **Specify Model Periods** and **Box-Cox Transformation** unchecked. We will use the default **Prediction Interval = 95.0 %**.

The screenshot shows the 'ARIMA with Predictors Forecast' dialog box. The main area contains four rows of predictor selection:

- Numeric Time Series Data (Y) >>**: Demand
- Optional Time Axis Labels >>**: Date
- Optional Continuous Pred. (X) >>**: Temp (C), TempSq
- Optional Categorical Pred. (X) >>**: WorkDay

Buttons for 'OK >>', 'Cancel', and 'Help' are on the right, and '<< Remove' is at the bottom center. The 'Model Options' section is expanded and contains:

- No. of Forecast Periods**: 24
- Prediction Interval**: 95.0 %
- Specify Model Periods**: (unchecked)
- Start Model at Period**: 1
- Withhold Periods**: 0
- End Model at Period**: (empty)
- Model Options**:
 - Seasonal Frequency**: (checked)
 - Specify**: (unchecked)
 - Select**: 7 - Daily (dropdown menu)
 - Automatically Detect**: (unchecked)
- Display ACF/PACF/LB Plots**: (checked)
- Display Residual Plots**: (checked)
- Box-Cox Transformation**: (unchecked)
- Rounded Lambda**: (checked)
- Optimal Lambda**: (unchecked)
- Lambda & Threshold (Shift)**: (unchecked)

- **No. of Forecast Periods** is greyed out because they are determined by the number of additional predictor rows that are provided, in this example, 14.
 - **Optional Continuous Pred. (X)** are continuous predictors.
 - **Optional Categorical Pred. (X)** are categorical predictors. In this example *WorkDay* could be either continuous or categorical since it is coded as 0,1.
5. Click **Model Options**. Select **Automatic Model Selection**. We will use the defaults: **Stepwise Procedure** and **Model Selection Criterion: AICc – Akaike information criterion with small sample size correction**, leave **Specify Nonseasonal Differencing (d)** and **Specify Seasonal Differencing (D)** unchecked.

ARIMA Model Options

Automatic Model Selection

Specify Model

Stepwise Procedure

Extended Model Search. Time limit seconds.

Model Selection Criterion

AICc - Akaike information criterion with small sample size correction

AIC - Akaike information criterion

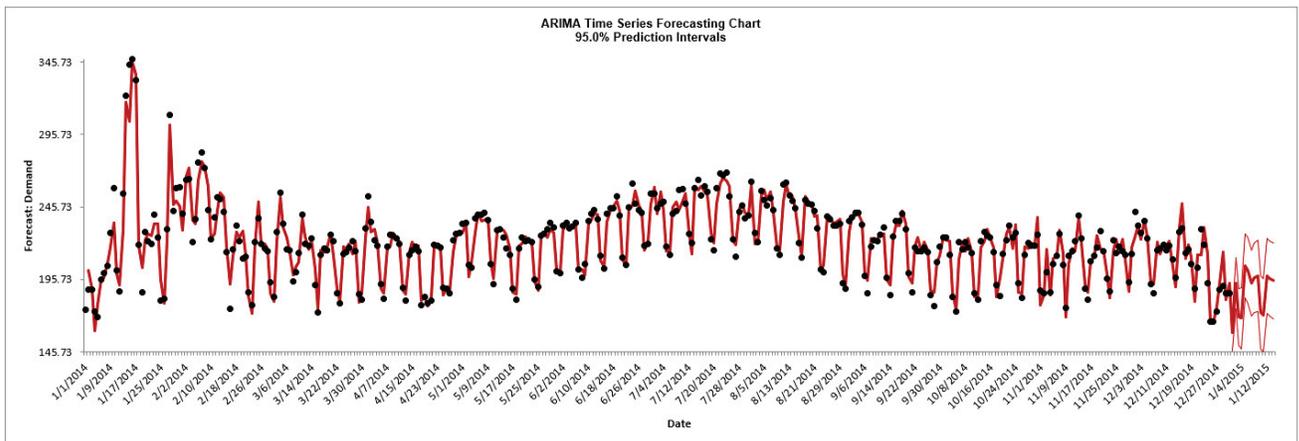
BIC - Bayesian information criterion

Specify Nonseasonal Differencing (d)

Specify Seasonal Differencing (D)

OK >>
Cancel
Help

- Click **OK** to return to the ARIMA Forecast dialog. Click **OK**. This is a complex model, so computation time will be approximately one to two minutes. The ARIMA forecast report is given:



This agrees with Figure 9.8 given in Hyndman, fpp2, Section 9.3, **Example: Forecasting electricity demand**, <https://otexts.com/fpp2/forecasting.html>. There are slight differences in the initial predicted values.

- Scroll down to view the ARIMA Model header:

ARIMA Model: Demand - Model Automatically Selected
Model Periods: All observations are used in the ARIMA model estimation. No withhold periods available for out-of-sample forecast accuracy evaluation.

If we had checked Specify Model Periods in the main dialog, the start, end or withhold selection would be summarized here as well.

8. The ARIMA Model Summary is given as:

ARIMA Model Summary	
AR Order (p)	2
I Order (d)	1
MA Order (q)	2
SAR Order (P)	2
SI Order (D)	0
SMA Order (Q)	0
Seasonal Frequency	7
Include Constant	0
No. of Predictors	3
Model Selection Criterion	AICc
Box-Cox Transformation	N/A
Lambda	
Threshold	

This is a summary of the model information: ARIMA (2,1,2) (2,0,0) with no constant and 3 predictors. Seasonal Frequency = 7; Model Selection Criterion = "AICc" and Box-Cox Transformation = "N/A" because Box-Cox Transformation was unchecked.

This agrees with the model given in Hyndman, fpp2, Section 9.3, **Example: Forecasting electricity demand**, <https://otexts.com/fpp2/forecasting.html>.

9. The Parameter Estimates are:

Parameter Estimates				
Term	Coefficient	SE Coefficient	T	P
AR_1	-0.063223451	0.075658448	0.83564	0.4039
AR_2	0.673128346	0.067270503	10.0063	0.0000
MA_1	0.022660844	0.043288704	0.52348	0.6010
MA_2	0.929862871	0.039474102	23.5563	0.0000
SAR_1	0.200902989	0.053912363	3.72647	0.0002
SAR_2	0.402632085	0.05676416	7.09307	0.0000
Temp (C)	-7.501559029	0.446098708	16.8159	0.0000
TempSq	0.17890261	0.008530253	20.9727	0.0000
WorkDay_1	30.56943168	1.295720007	23.5926	0.0000

ARIMA Parameter Estimates include significance tests; P-Values < .05 are significant and highlighted in red. All of the predictors are significant. The AR_1 and MA_1 terms are not significant but they must remain in the model due to hierarchy. Note that for AR/MA model

order selection, minimum AICc should be used, rather than significance tests (see [Kostenko, A.V. and Hyndman, R.J.](#)).

The categorical predictor *WorkDay_1* has 0 as the hidden reference level. If there were three category levels, only two would appear in the table.

Note: SE Coefficients, T and P-Values may be slightly different due to machine precision.

10. The ARIMA Model Statistics are:

ARIMA Model Statistics	
No. of Observations	365
DF	355
StDev	6.565744215
Variance	43.1089971
Log-Likelihood	-1200.73256
AICc	2422.088353
AIC	2421.465123
BIC	2460.436662

Degrees of freedom (DF) = $n - 10$ (9 terms in the model, 1 order of differencing). See [ARIMA Model Statistics and Information Criteria for Model Comparison](#).

11. The In-Sample Forecast Accuracy metrics are:

Forecast Accuracy		
Metric	In-Sample (Estimation) One-Step-Ahead Forecast	Out-of-Sample (Withhold) One-Step-Ahead Forecast
N	364	
RMSE	6.615605184	
MAE	4.768089714	
MAPE	2.162224989	
MASE	0.327382095	

MASE is less than one, so it is a better forecast than would be obtained from a naïve forecast (set all forecasts to be the value of the last observation). See [Forecast Accuracy Metrics](#).

The analysis can be rerun (using Recall SigmaXL Dialog) with a withhold sample to obtain Out-of-Sample One-Step-Ahead or Multi-Step-Ahead Forecast errors, but we will not do so here.

12. The Forecast Table is given as:

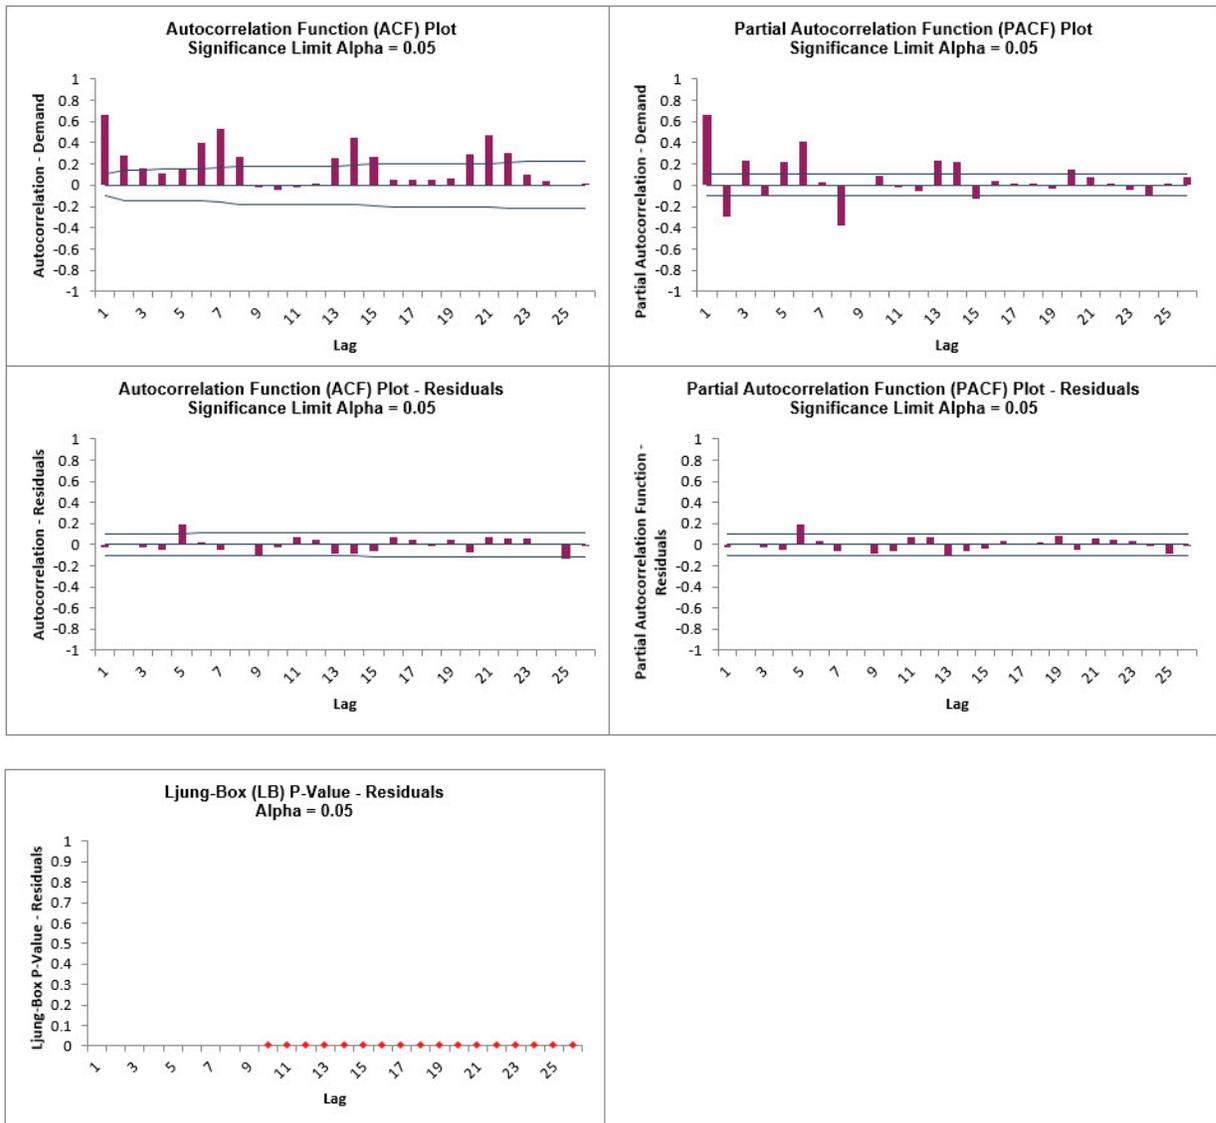
Forecast Table				
Period	Withhold Data	Forecast	Lower 95.0% PI	Upper 95.0% PI
366		159.2517864	146.3402335	172.1633392
367		193.2139395	175.7207745	210.7071045
368		169.8531097	150.3787979	189.3274216
369		168.9969534	147.9371527	190.0567541
370		205.1573872	183.2956912	227.0190833
371		201.9616978	179.3846259	224.5387697
372		193.1293236	170.1638907	216.0947565
373		197.1190445	173.1860174	221.0520716
374		198.0338842	173.4650501	222.6027184
375		173.0122062	147.9556892	198.0687231
376		171.1330507	145.7330385	196.5330629
377		198.1491778	172.4687919	223.8295637
378		196.3934701	170.5020251	222.2849151
379		194.6921179	168.6187184	220.7655175

These are the same forecast and prediction interval values displayed in the Forecast Chart but provided for further analysis or charting. If Withhold Periods are specified, the Withhold Data will be displayed as well.

13. The Predictor Values for Forecast are the additional predictor (X) values used to obtain the 14 period forecast in the Forecast Table:

Predictor Values for Forecast		
Temp (C)	TempSq	WorkDay_1
26	676	0
26	676	1
26	676	0
26	676	0
26	676	1
26	676	1
26	676	1
26	676	1
26	676	1
26	676	1
26	676	0
26	676	0
26	676	1
26	676	1
26	676	1

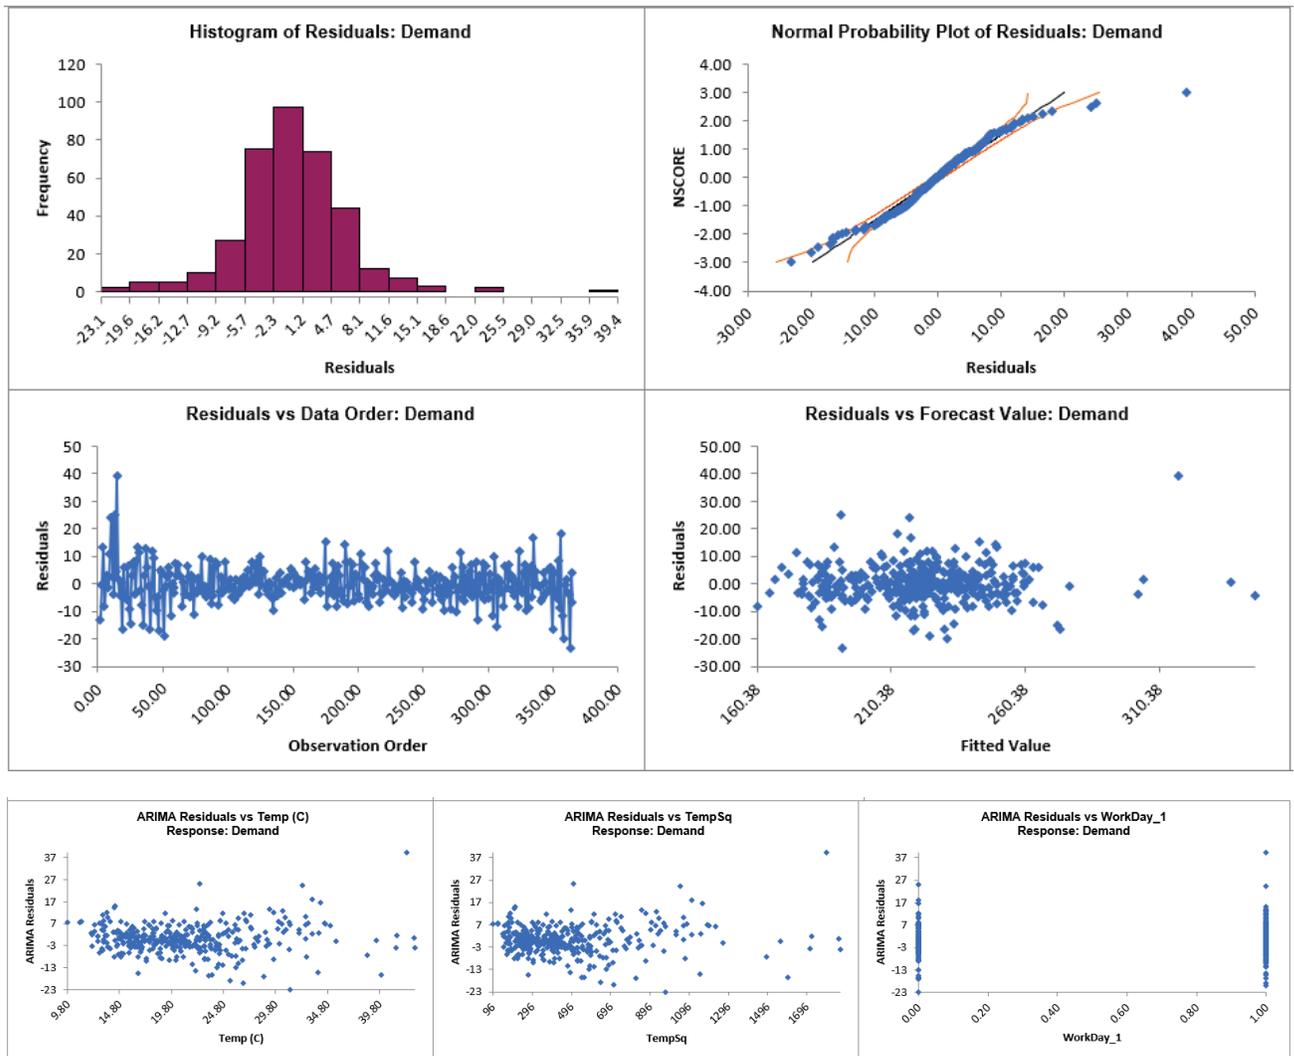
14. Click on the **ARIMA ACF PACF LB** sheet to view the ACF/PACF/LB Plots:



We can see that much of the autocorrelation has been removed by the ARIMA with Predictors model (with the exception of lag 5 and 25 in the ACF, lag 5 in the PACF).

The Ljung-Box plot shows that some significant autocorrelation still remains (the red P-Values are significant at $\alpha=0.05$) - so the model can potentially be improved. This does not mean that the model is a bad model, it can still be very useful for prediction purposes, but the prediction intervals may not provide accurate coverage.

15. Click on the **ARIMA Residuals** sheet to view the Residual Plots:

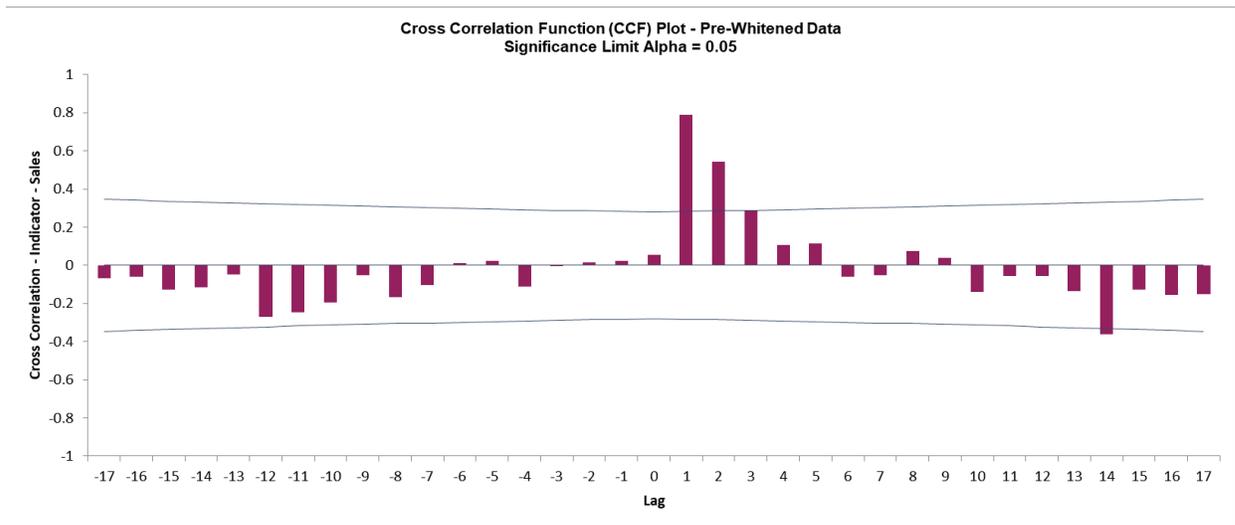


Looking at the histogram and normal probability plot, there are some outliers (but smaller than would have been the case if we did not include the predictors). Later, these will be investigated with a control chart on the residuals.

Note that additional residual plots are provided for each predictor.

Sales with Indicator - Modified Series M

16. Open **Sales with Indicator - Modified Series M.xlsx**. This is modified Series M data from Box and Jenkins. Originally, a set of 150 monthly corporate sales values along with a leading indicator, the data was modified by converting it to quarterly values by averaging every three months, so 50 quarters, labelled as Q1-Y1, Q2-Y1, etc. This was done in order to simplify the analysis of the leading indicator. Although the data was monthly and summarized to quarterly, it will be treated as nonseasonal, as done in Box and Jenkins. See the [Run Chart](#), [ACF/PACF Plots](#), [CCF Plot](#), [Spectral Density Plot](#) and [Trend Decomposition Plots](#) for this data.
17. Recall that the Pre-Whitened CCF Plot for this data is:



Pre-Whitening the data has dramatically altered the CCF plot, allowing us to see the underlying cross correlation pattern. Lags 1 and 2 are significantly positive, and Lag 3 is just on the significance line. Use this as a guide to assist with what lags to include in the model but it is possible that some additional lags may be significant. In this example, we will initially include up to Lag 5.

Note that while X is called a leading indicator, i.e., X comes before Y in time, the positive lag means that the X variable is lagging the Y variable in terms of correlation structure. SigmaXL uses this convention as given in Box and Jenkins (2016, pp. 437-440).

18. Click the **Indicator with Lags** Sheet tab. **SigmaXL > Time Series Forecasting > ARIMA Forecast > Forecast with Predictors**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
19. Select *Sales*, click **Numeric Time Series Data (Y) >>**; select *Qtr-Year*, click **Optional Time Axis Labels >>**; select *Indicator* to *Indicator Lag 5*, click **Optional Continuous Pred. >>**. Check **Display ACF/PACF/LB Plots** and **Display Residual Plots**. Check **Specify Model Periods**. Set **Withhold Periods = 12** (i.e., 3 years). Select **Withhold Forecast Type: Multi-Step-Ahead with Prediction Interval at Start of Withhold**. Leave **Specify Model Periods**, **Seasonal Frequency** and **Box-Cox Transformation** unchecked. We will use the default **Prediction Interval = 95.0 %**.

ARIMA with Predictors Forecast

Numeric Time Series Data (Y) >> Sales

Optional Time Axis Labels >> Qtr-Year

Optional Continuous Pred. (X) >> Indicator Lag 3
Indicator Lag 4
Indicator Lag 5

Optional Categorical Pred. (X) >>

<< **Remove**

OK >>
Cancel
Help

No. of Forecast Periods	24	Model Options	<input checked="" type="checkbox"/> D<u>i</u>splay ACF/PACF/LB Plots
Prediction Interval	95.0 %		<input checked="" type="checkbox"/> D<u>i</u>splay Residual Plots
<input checked="" type="checkbox"/> S<u>pec</u>ify Model Periods		<input type="checkbox"/> S<u>ea</u>sonal Frequency	<input type="checkbox"/> B<u>ox</u>-C<u>ox</u> Transformation
Start Model at Period	1	<input checked="" type="radio"/> S<u>pec</u>ify 12	<input checked="" type="radio"/> R<u>ounded</u> <u>L</u>ambda
<input checked="" type="radio"/> W<u>ith</u>hold Periods	12	<input type="radio"/> S<u>elect</u> 4-Quarterly	<input type="radio"/> O<u>pti</u>mal <u>L</u>ambda
<input type="radio"/> E<u>nd</u> Model at Period		<input type="radio"/> A<u>ut</u>omatically Detect	<input type="radio"/> L<u>ambda</u> & <u>T</u>hreshold (Shift)
Withhold Forecast Type:			
<input type="radio"/> O<u>ne</u>-Step-Ahead with Prediction Interval at: Start of Withhold			
<input type="checkbox"/> I<u>n</u>clude in Residuals			
<input checked="" type="radio"/> M<u>ulti</u>-Step-Ahead with Prediction Interval at Start of Withhold.			

20. Click **Model Options**. Select **Automatic Model Selection**. We will use the defaults: **Stepwise Procedure** and **Model Selection Criterion: AICc – Akaike information criterion with small sample size correction**, leave **Specify Nonseasonal Differencing (d)** and **Specify Seasonal Differencing (D)** unchecked.

ARIMA Model Options ✕

Automatic Model Selection

Specify Model

Stepwise Procedure

Extended Model Search. Time limit seconds.

Model Selection Criterion

AICc - Akaike information criterion with small sample size correction

AIC - Akaike information criterion

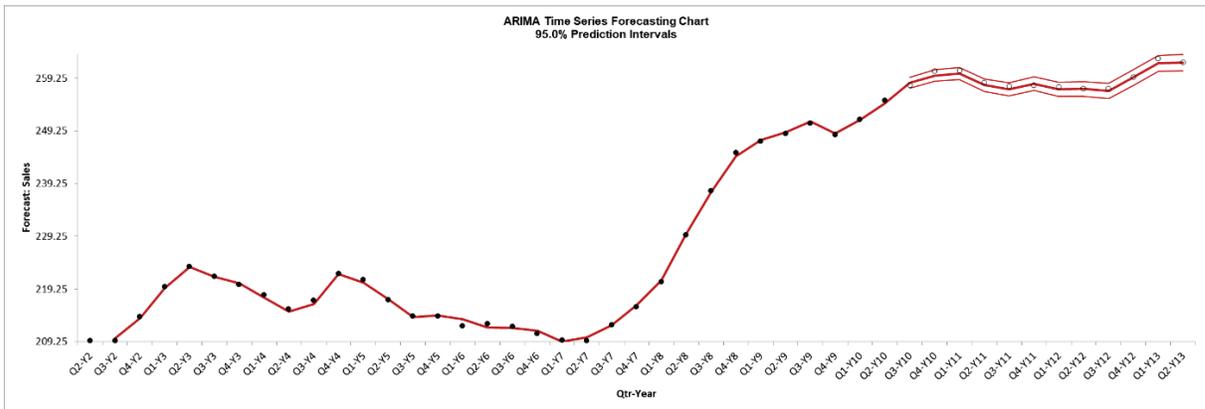
BIC - Bayesian information criterion

Specify Nonseasonal Differencing (d)

Specify Seasonal Differencing (D)

Tip: When using **Recall SigmaXL Dialog** and if there are no changes to the **Model Option** settings, the previous settings will be used. It is not necessary to repeat this step.

- Click **OK** to return to the ARIMA Forecast dialog. Click **OK**. The ARIMA with Predictors forecast report is given:



The blank dots are the data values in the withhold sample with a multi-step forecast and prediction intervals displayed at the start of the withhold sample. The model uses the Indicator Lag values in the withhold sample and does quite well at predicting the 12 quarters of Sales values.

Note that the chart starts at Q2-Y2. This is due to the first 5 rows of X and Y data being deleted since there are missing values in Indicator Lags 1 to 5.

22. Scroll down to view the ARIMA Model header:

ARIMA Model: Sales - Model Automatically Selected

Model Periods: Model parameter estimates calculated excluding 12 withhold periods. Missing values in initial rows of predictors. Deleted first 5 rows of X and Y data.

23. The ARIMA Model Summary is given as:

ARIMA Model Summary	
AR Order (p)	0
I Order (d)	1
MA Order (q)	1
SAR Order (P)	0
SI Order (D)	0
SMA Order (Q)	0
Seasonal Frequency	1
Include Constant	1
No. of Predictors	6
Model Selection Criterion	AICc
Box-Cox Transformation	N/A
Lambda	
Threshold	

This is a summary of the model information: ARIMA (0,1,1) with no constant and 6 predictors. Seasonal Frequency = 1 (nonseasonal); Model Selection Criterion = "AICc" and Box-Cox Transformation = "N/A" because Box-Cox Transformation was unchecked.

24. The Parameter Estimates are:

Parameter Estimates				
Term	Coefficient	SE Coefficient	T	P
MA_1	0.661826535	0.171947916	3.84899	0.0008
Const:Trend	0.107385184	0.033507981	3.20476	0.0038
Indicator	-0.193804762	0.279261631	0.69399	0.4944
Indicator Lag 1	8.199265436	0.304329926	26.942	0.0000
Indicator Lag 2	5.392904003	0.2952757	18.264	0.0000
Indicator Lag 3	2.481788513	0.301027513	8.24439	0.0000
Indicator Lag 4	0.885376557	0.324714746	2.72663	0.0118
Indicator Lag 5	0.154053202	0.264999213	0.58133	0.5664

ARIMA Parameter Estimates include significance tests; P-Values < .05 are significant and highlighted in red. The CCF Plot suggested that only Lags 1 and 2, possibly 3, would be significant, but here we see that Lag 4 is also significant (hence why CCF is just an approximate guide). Later we will rerun the model and remove Indicator & Indicator Lag 5 and manually

compare AICc values. Note: SE Coefficients, T and P-Values may be slightly different due to machine precision.

25. The ARIMA with Predictors Model Statistics are:

ARIMA Model Statistics	
No. of Observations	33
DF	24
StDev	0.505833251
Variance	0.255867277
Log-Likelihood	-19.8991901
AICc	65.98019828
AIC	57.7983801
BIC	70.99000323

- The number of observations, $n = 50 - 12$ (*withhold*) $- 5$ (*deleted rows*) = 33
- Degrees of freedom (DF) = 33 (n) $- 9$ (8 terms in the model, 1 nonseasonal difference) = 24

26. The Forecast Accuracy metrics are:

Forecast Accuracy		
Metric	In-Sample (Estimation) One-Step-Ahead Forecast	Multi-Step-Ahead Forecast
N	32	12
RMSE	0.441579854	0.519299269
MAE	0.355486014	0.433503476
MAPE	0.160638285	0.167034359
MASE	0.123200929	0.150239472

The Out-of-Sample forecast errors are only slightly larger than the In-Sample, so this is a good prediction.

MASE is less than one, so it is a better forecast than would be obtained from a naïve forecast (set all forecasts to be the value of the last observation). See [Forecast Accuracy Metrics](#).

27. The Forecast Table is given as:

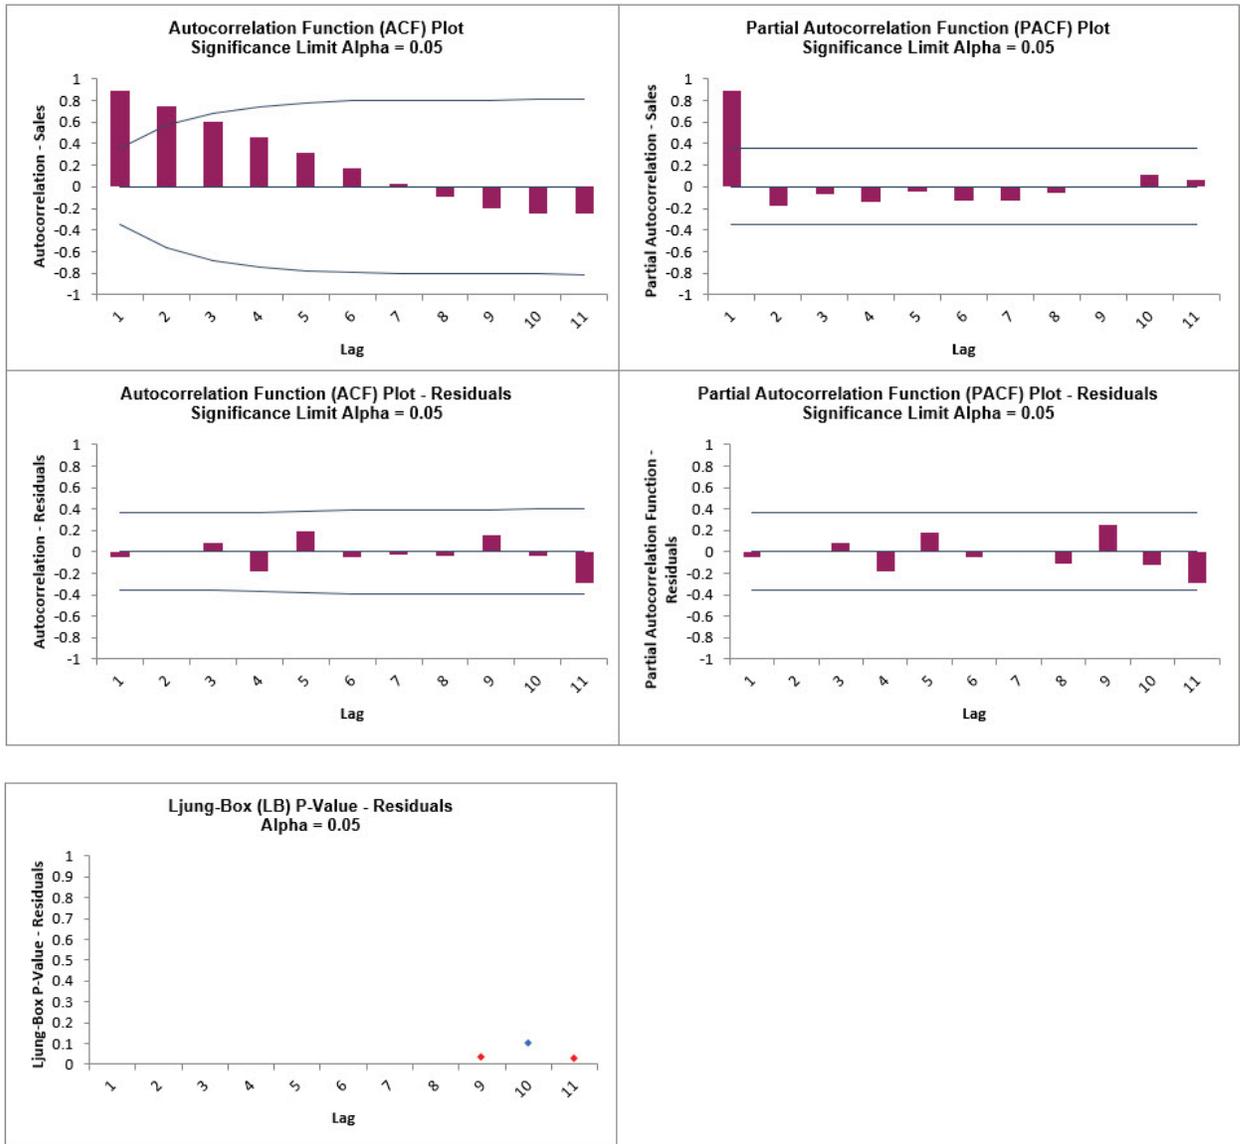
Forecast Table				
Period	Withhold Data	Multi-Step Ahead Forecast	Lower 95.0% PI	Upper 95.0% PI
34	257.9333333	258.3217379	257.2913893	259.3520865
35	260.5666667	259.7168718	258.6292016	260.8045419
36	260.7333333	260.1163414	258.974223	261.2584597
37	258.4333333	257.9037045	256.7096181	259.097791
38	257.6666667	257.0834025	255.8395172	258.3272877
39	257.9333333	258.1603883	256.8686227	259.4521539
40	257.5666667	257.1341925	255.7962588	258.4721261
41	257.2666667	257.1889296	255.8063687	258.5714904
42	257.3	256.7947588	255.3689669	258.2205507
43	259.4	259.3538829	257.8861328	260.8216331
44	263	262.0730127	260.5644708	263.5815546
45	262.2333333	262.2149879	260.6667286	263.7632471

These are the same forecast and prediction interval values displayed in the Forecast Chart but provided for further analysis or charting. Note that this table gives the Forecast Period number whereas the Chart displays the Quarter-Year.

28. The Predictor Values for Forecast are the additional predictor (X) values used to obtain the 12 period forecast in the Forecast Table:

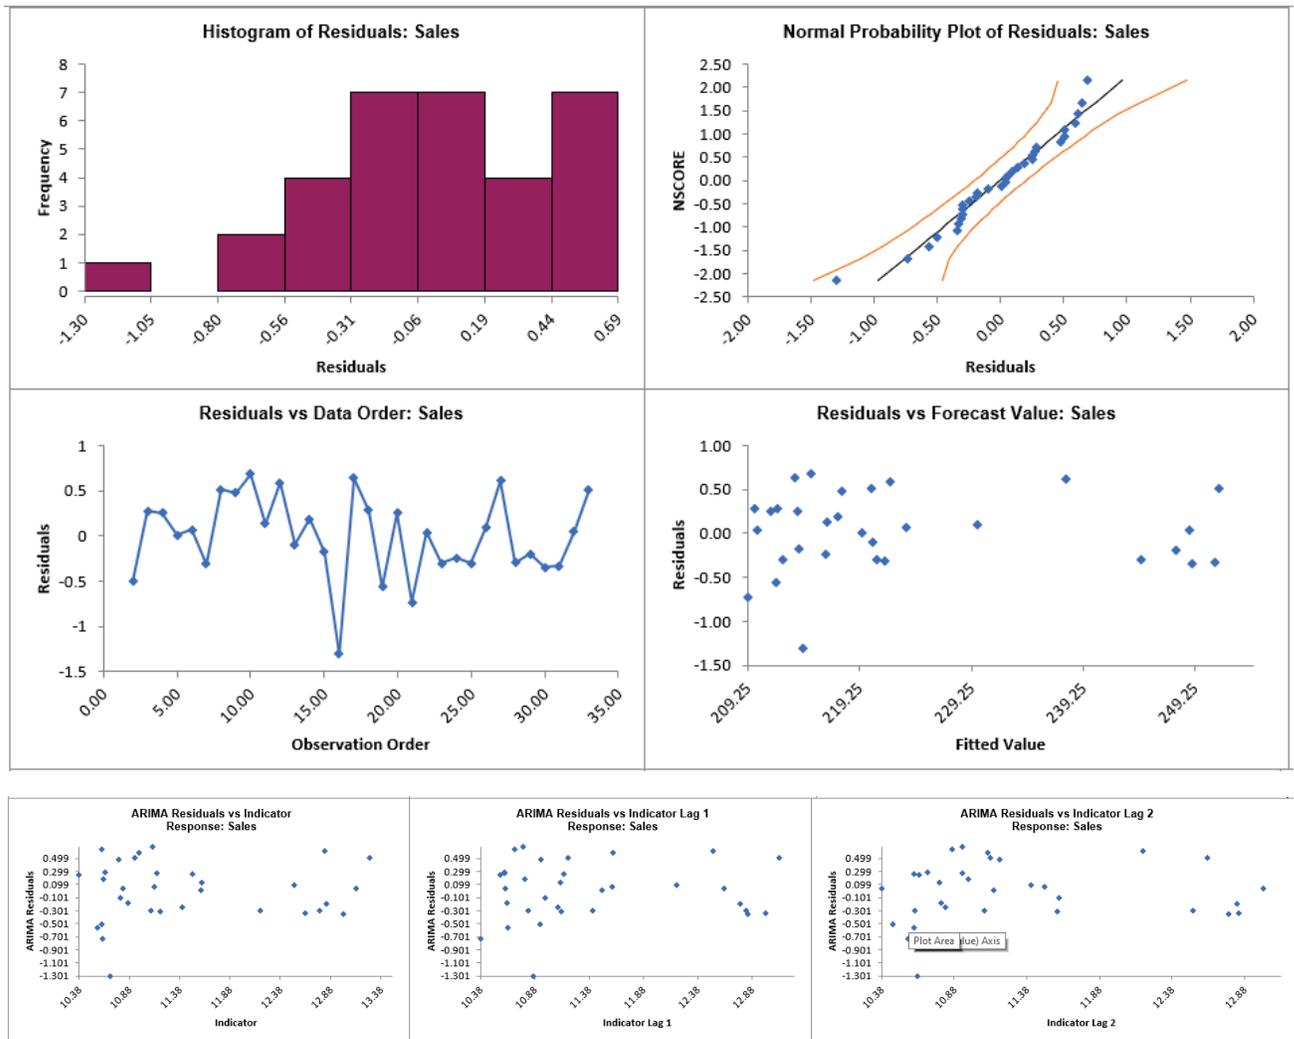
Predictor Values for Forecast					
Indicator	Indicator Lag 1	Indicator Lag 2	Indicator Lag 3	Indicator Lag 4	Indicator Lag 5
13.4166667	13.5	13.27	13.13333333	12.62666667	13.01
13.4066667	13.41666667	13.5	13.27	13.13333333	12.62666667
13.13	13.40666667	13.41666667	13.5	13.27	13.13333333
13.21	13.13	13.40666667	13.41666667	13.5	13.27
13.3566667	13.21	13.13	13.40666667	13.41666667	13.5
13.13	13.35666667	13.21	13.13	13.40666667	13.41666667
13.2233333	13.13	13.35666667	13.21	13.13	13.40666667
13.16	13.22333333	13.13	13.35666667	13.21	13.13
13.4966667	13.16	13.22333333	13.13	13.35666667	13.21
13.6033333	13.49666667	13.16	13.22333333	13.13	13.35666667
13.4433333	13.60333333	13.49666667	13.16	13.22333333	13.13
13.56	13.44333333	13.60333333	13.49666667	13.16	13.22333333

29. Click on the **ARIMA ACF PACF LB** sheet to view the ACF/PACF/LB Plots:



We can see that much of the autocorrelation has been removed by the ARIMA with Predictors model. The Ljung-Box plot shows that some significant autocorrelation still remains (the red P-Values are significant at alpha = .05) - so the model can potentially be improved. This does not mean that the model is a bad model, it can still be very useful for prediction purposes, but the prediction intervals may not provide accurate coverage.

30. Click on the **ARIMA Residuals** sheet to view the Residual Plots:

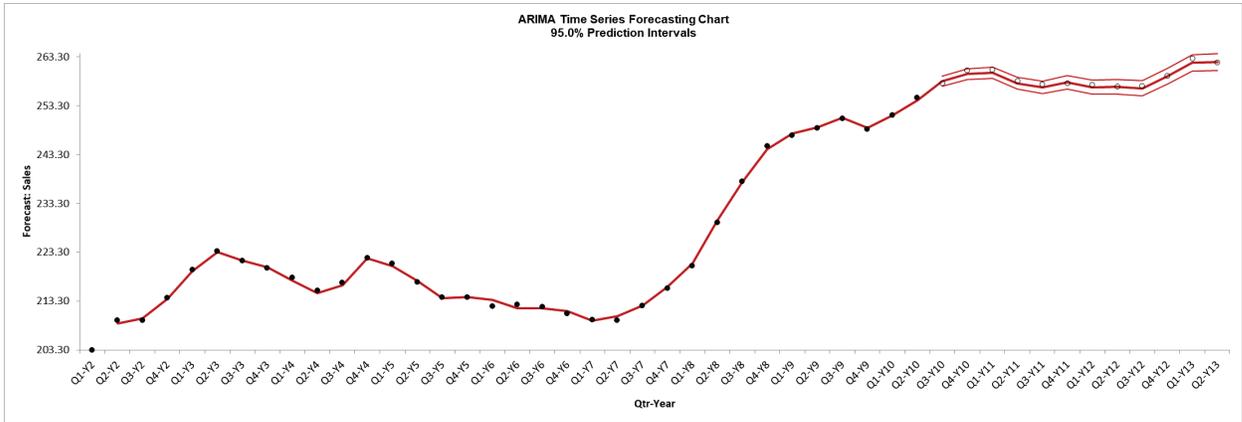


The residuals are approximately normally distributed, with a roughly straight line on the normal probability plot. There are no obvious extreme outliers or patterns in the charts. The histogram might suggest a left skew but the sample size is quite small. (Normality tests may also be applied on the residuals using SigmaXL > Statistical Tools > Descriptive Statistics: Options, Additional Normality Tests, and these would all show the residuals to be normal).

Note that additional residual plots are provided for each predictor.

31. Now we will rerun the analysis and remove the insignificant terms *Indicator* and *Indicator Lag 5*. Click **Recall SigmaXL Dialog** menu or press **F3** to recall last dialog.
32. Select *Indicator*, click Remove; select *Indicator Lag 5*, click Remove:

33. We will not make any changes to the **Model Options**. Click **OK**. The revised ARIMA with Predictors forecast report is given:



Note that the chart now starts at Q1-Y2. This is due to the first 4 rows of X and Y data being deleted since there are missing values in Indicator Lags 1 to 4.

34. Scroll down to view the ARIMA Model header:

ARIMA Model: Sales - Model Automatically Selected

Model Periods: Model parameter estimates calculated excluding 12 withhold periods. Missing values in initial rows of predictors. Deleted first 4 rows of X and Y data.

35. The ARIMA Model Summary is given as:

ARIMA Model Summary	
AR Order (p)	0
I Order (d)	1
MA Order (q)	1
SAR Order (P)	0
SI Order (D)	0
SMA Order (Q)	0
Seasonal Frequency	1
Include Constant	1
No. of Predictors	4
Model Selection Criterion	AICc
Box-Cox Transformation	N/A
Lambda	
Threshold	

This is a summary of the model information: ARIMA (0,1,1) without a constant has not changed. Now there are 4 predictors. Seasonal Frequency = 1 (nonseasonal); Model Selection Criterion = "AICc" and Box-Cox Transformation = "N/A" because Box-Cox Transformation was unchecked.

36. The Parameter Estimates are:

Parameter Estimates				
Term	Coefficient	SE Coefficient	T	P
MA_1	0.579458707	0.137596707	4.21128	0.0003
Const:Trend	0.110416608	0.039494253	2.79576	0.0094
Indicator Lag 1	8.14204032	0.260590425	31.2446	0.0000
Indicator Lag 2	5.381655285	0.288605873	18.6471	0.0000
Indicator Lag 3	2.521912726	0.293140527	8.60308	0.0000
Indicator Lag 4	0.932902504	0.262059255	3.55989	0.0014

ARIMA Parameter Estimates include significance tests; P-Values < .05 are significant and highlighted in red. The CCF Plot suggested that only Lags 1 and 2, possibly 3, would be significant, but here we see that Lag 4 is also significant (hence why CCF is just an approximate guide).

37. The ARIMA with Predictors Model Statistics are:

ARIMA Model Statistics	
No. of Observations	34
DF	27
StDev	0.491053186
Variance	0.241133231
Log-Likelihood	-20.6849129
AICc	59.84982585
AIC	55.36982585
BIC	65.84537878

- The number of observations, $n = 50 - 12$ (*withhold*) $- 4$ (*deleted rows*) = 34
- Degrees of freedom (DF) = 34 (n) $- 7$ (6 terms in the model, 1 nonseasonal difference) = 27
- AICc is lower than the [previous](#) model which included insignificant *Indicator* terms.

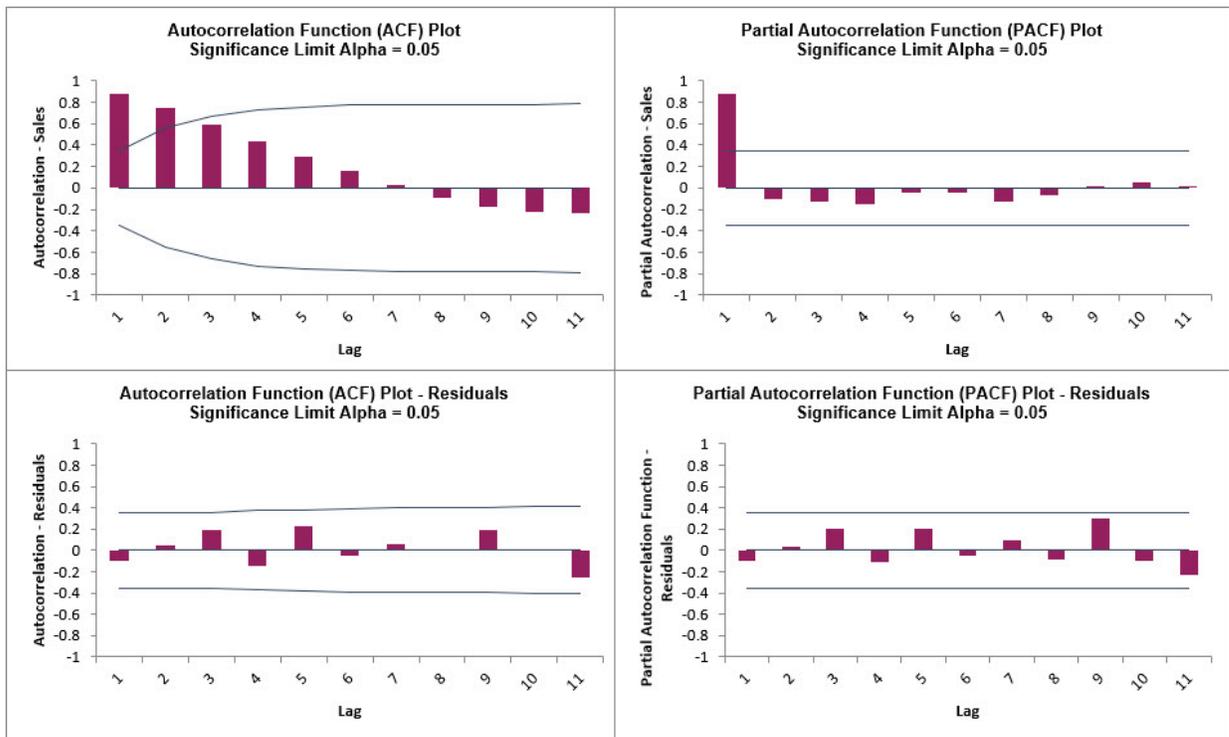
38. The Forecast Accuracy metrics are:

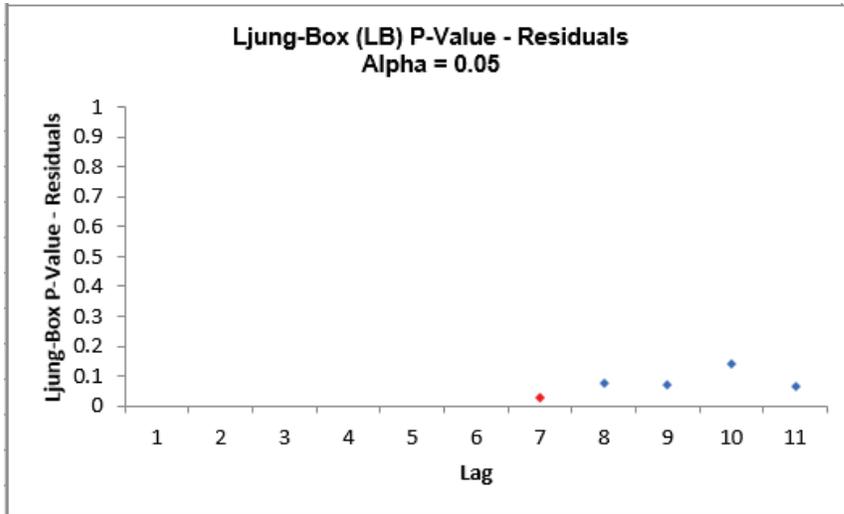
Forecast Accuracy		
Metric	In-Sample (Estimation) One-Step-Ahead Forecast	Out-of-Sample (Withhold) Multi-Step Ahead Forecast
N	33	12
RMSE	0.44993435	0.505510545
MAE	0.363143972	0.420790326
MAPE	0.164705945	0.162132467
MASE	0.121786088	0.141118707

Compared to the [previous](#) Forecast Accuracy metrics the Out-of-Sample forecast errors are slightly lower, so this is a good prediction with a simpler model.

MASE is less than one, so it is a better forecast than would be obtained from a naïve forecast (set all forecasts to be the value of the last observation). See [Forecast Accuracy Metrics](#).

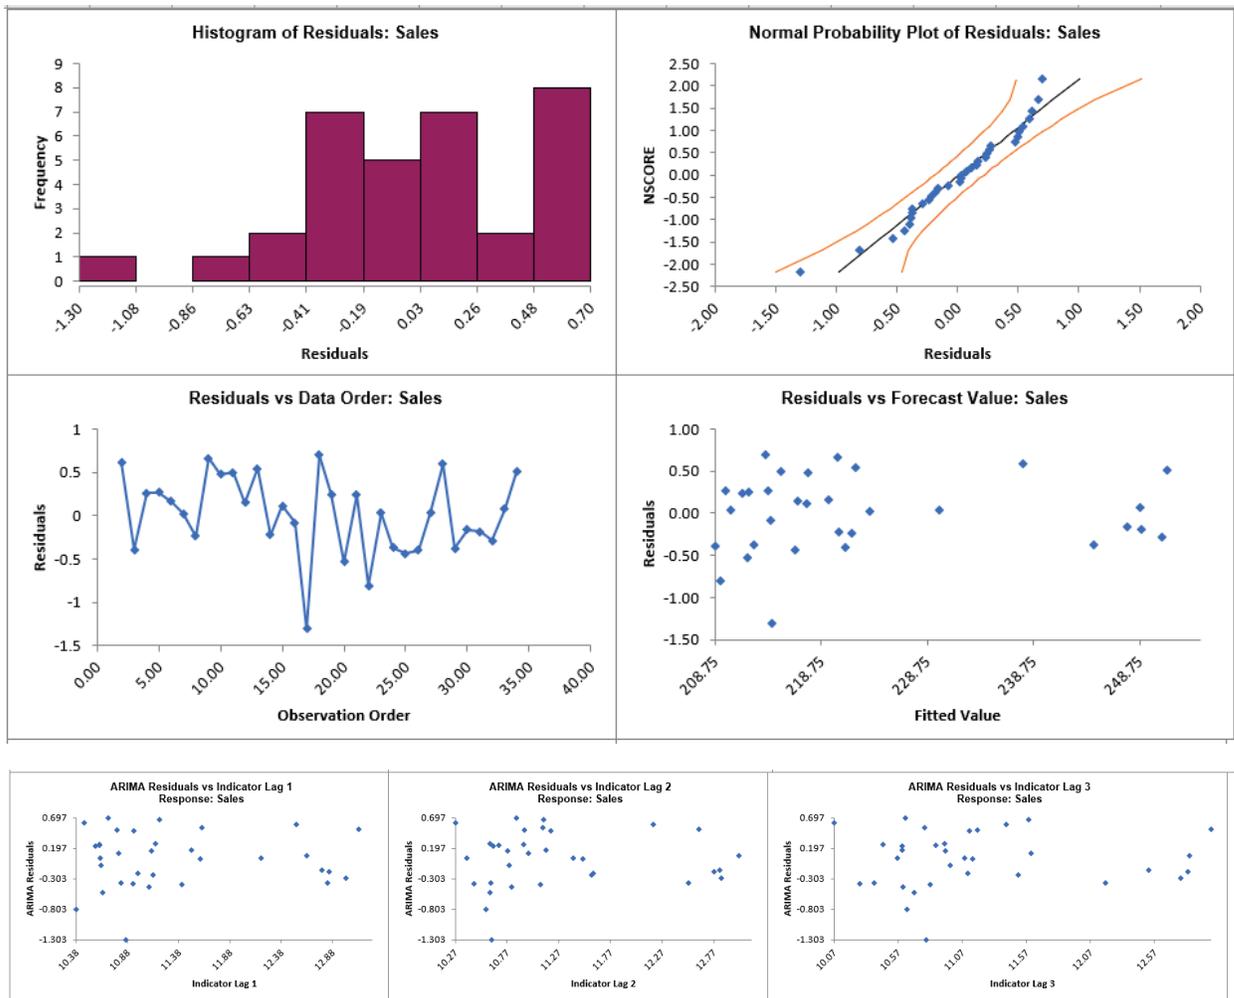
39. Click on the **ARIMA ACF PACF LB** sheet to view the ACF/PACF/LB Plots:





We can see that much of the autocorrelation has been removed by the ARIMA with Predictors model. The Ljung-Box plot shows fewer significant P-Values than the [previous](#) model.

40. Click on the **ARIMA Residuals** sheet to view the Residual Plots:



The residuals look similar to those in the [previous](#) model, approximately normally distributed, with a roughly straight line on the normal probability plot. There are no obvious extreme outliers or patterns in the charts. The histogram might suggest a left skew but the sample size is quite small. (Normality tests may also be applied on the residuals using SigmaXL > Statistical Tools > Descriptive Statistics: Options, Additional Normality Tests, and these would all show the residuals to be normal).

Note that additional residual plots are provided for each predictor, *Indicator Lag 1 to 4*.

ARIMA – Multiple Seasonal Decomposition (MSD) Forecast

ARIMA does not have a theoretical frequency limit, but for computational efficiency and to minimize the potential loss of observations through differencing, we recommend using ARIMA – Multiple Seasonal Decomposition (MSD) for seasonal frequency greater than 52 (or with multiple frequencies).

The seasonal component is first removed through decomposition, a nonseasonal ARIMA model fitted to the remainder (+trend), and then the seasonal component is added back in. For forecasting, a naïve seasonal forecast is used on the seasonal component. Note that the prediction intervals are derived from the ARIMA model and do not include uncertainty in the seasonal component.

As the name implies, Multiple Seasonal Decomposition (MSD) also accommodates multiple seasonality, for example the half-hourly data with a seasonal frequency of 48 observations per day and 336 observations per week. When using MSD, it is recommended to limit the forecast period to 2*dominant seasonal frequency.

For further details and formulas, see Appendix: [Seasonal Trend Decomposition](#).

Monthly Airline Passengers - Series G

1. Open **Monthly Airline Passengers - Series G.xlsx (Sheet 1 tab)**. This is the Series G data from Box and Jenkins, monthly total international airline passengers for January 1949 to December 1960. See the [Run Chart](#), [ACF/PACF Plots](#), [Spectral Density Plot](#) and [Seasonal Trend Decomposition Plots](#) for this data. The Multiple Seasonal Decomposition (MSD) option is not necessary for this data, but by way of introduction, we will use this to compare to the previous ARIMA analysis.
2. Click **SigmaXL > Time Series Forecasting > ARIMA Forecast > Multiple Seasonal Decomposition Forecast**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
3. Select *Monthly Airline Passengers*, click **Numeric Time Series Data (Y) >>**. Select *Date*, click **Optional Time Axis Labels >>**. Check **Display ACF/PACF/LB Plots** and **Display Residual Plots**. Check **Specify Model Periods**. Set **Withhold Periods = 24**. Select **Withhold Forecast Type: Multi-Step-Ahead with Prediction Interval at Start of Withhold**. Select **Seasonal Frequency Specify** and enter 12. Check **Box-Cox Transformation** and select **Rounded Lambda**. We will use the default **Prediction Interval = 95.0 %**.

ARIMA MSD Forecast

Obs. No. Ln (Airline Passengers)	Numeric Time Series Data (Y) >>	Monthly Airline Passenger	OK >>
	Optional Time Axis Labels >>	Date	Cancel
	<< Remove		Help

No. of Forecast Periods <input type="text" value="24"/>	Model Options	<input checked="" type="checkbox"/> Display ACF/PACF/LB Plots
Prediction Interval <input type="text" value="95.0 %"/>		<input checked="" type="checkbox"/> Display Residual Plots
<input checked="" type="checkbox"/> Specify Model Periods	Seasonal Frequency	<input checked="" type="checkbox"/> Box-Cox Transformation
Start Model at Period <input type="text" value="1"/>	<input checked="" type="radio"/> Specify <input type="text" value="12"/>	<input checked="" type="radio"/> Rounded Lambda
<input checked="" type="radio"/> Withhold Periods <input type="text" value="24"/>	<input type="radio"/> Select <input type="text" value="4-Quarterly"/>	<input type="radio"/> Optimal Lambda
<input type="radio"/> End Model at Period <input type="text" value=""/>	<input type="radio"/> Automatically Detect	<input type="radio"/> Lambda & Threshold (Shift)
Withhold Forecast Type:		
<input type="radio"/> One-Step-Ahead with Prediction Interval at: <input type="text" value="Start of Withhold"/>		
<input type="checkbox"/> Include in Residuals		
<input checked="" type="radio"/> Multi-Step-Ahead with Prediction Interval at Start of Withhold.		

- **Seasonal Frequency** can have multiple entries however, we recommend no more than 3 values.

4. Click **Model Options**.

ARIMA Model Options

Automatic Model Selection

Specify Model

Stepwise Procedure

Extended Model Search. Time limit seconds.

Model Selection Criterion

AICc - Akaike information criterion with small sample size correction

AIC - Akaike information criterion

BIC - Bayesian information criterion

Specify Nonseasonal Differencing (d)

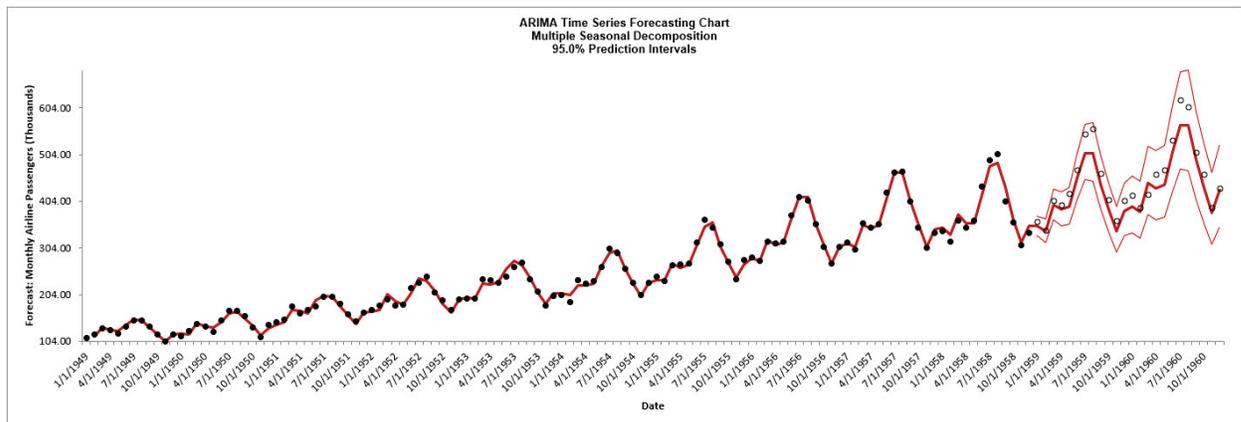
Specify Seasonal Differencing (D)

OK >>

Cancel

Help

5. We will use the default **Automatic Model Selection** with **AICc** as the **Model Selection Criterion**. Click **OK** to return to the ARIMA MSD Forecast dialog. Click **OK**. The ARIMA forecast report is given:



6. Scroll down to view the ARIMA Model header:

ARIMA Model (Multiple Seasonal Decomposition): Monthly Airline Passengers (Thousands) - Model Automatically Selected
Model Periods: Model parameter estimates calculated excluding 24 withhold periods.

7. The ARIMA Model Summary is given as:

ARIMA Model Summary	
AR Order (p)	0
I Order (d)	1
MA Order (q)	1
SAR Order (P)	0
SI Order (D)	0
SMA Order (Q)	0
Seasonal Frequency	12 Decomposition
Include Constant	1
No. of Predictors	0
Model Selection Criterion	AICc
Box-Cox Transformation	Rounded Lambda
Lambda	0
Threshold	0

This is a summary of the model information for the deseasonalized data: ARIMA (0,1,1) with a constant. Seasonal Frequency = 12 using Decomposition and Model Selection Criterion = "AICc". There are no seasonal terms in the model. The Box-Cox Transformation is "Rounded Lambda" with Lambda = 0 (Ln transformation).

8. The Parameter Estimates for the deseasonalized Airline Passenger data are:

Parameter Estimates				
Term	Coefficient	SE Coefficient	T	P
MA_1	0.322113898	0.095098865	3.38715	0.0010
Const:Trend	0.009427003	0.001881534	5.01027	0.0000

9. The ARIMA Model Statistics are:

ARIMA Model Statistics	
No. of Observations	120
DF	117
StDev	0.030409792
Variance	0.000924755
Log-Likelihood	247.7313206
AICc	-489.253946
AIC	-489.462641
BIC	-481.125271

- The number of observations, $n = 144 - 24$ (*withhold*) = 120
- Degrees of freedom (DF) = $120 (n) - 3$ (2 terms in the model, 1 nonseasonal difference) = 117
- Note that the model statistics are based on the Ln transformed data, not the original data.
- Comparing to the [Exponential Smoothing MSD Model Statistics](#), we see that the StDev and Variance are approximately equal, but the Log-Likelihood, AICc, AIC and BIC are very different. This is due to different formulas being used in the Likelihood function. You cannot use Information Criteria to compare ARIMA and Exponential Smooth models to determine which model has the best fit.

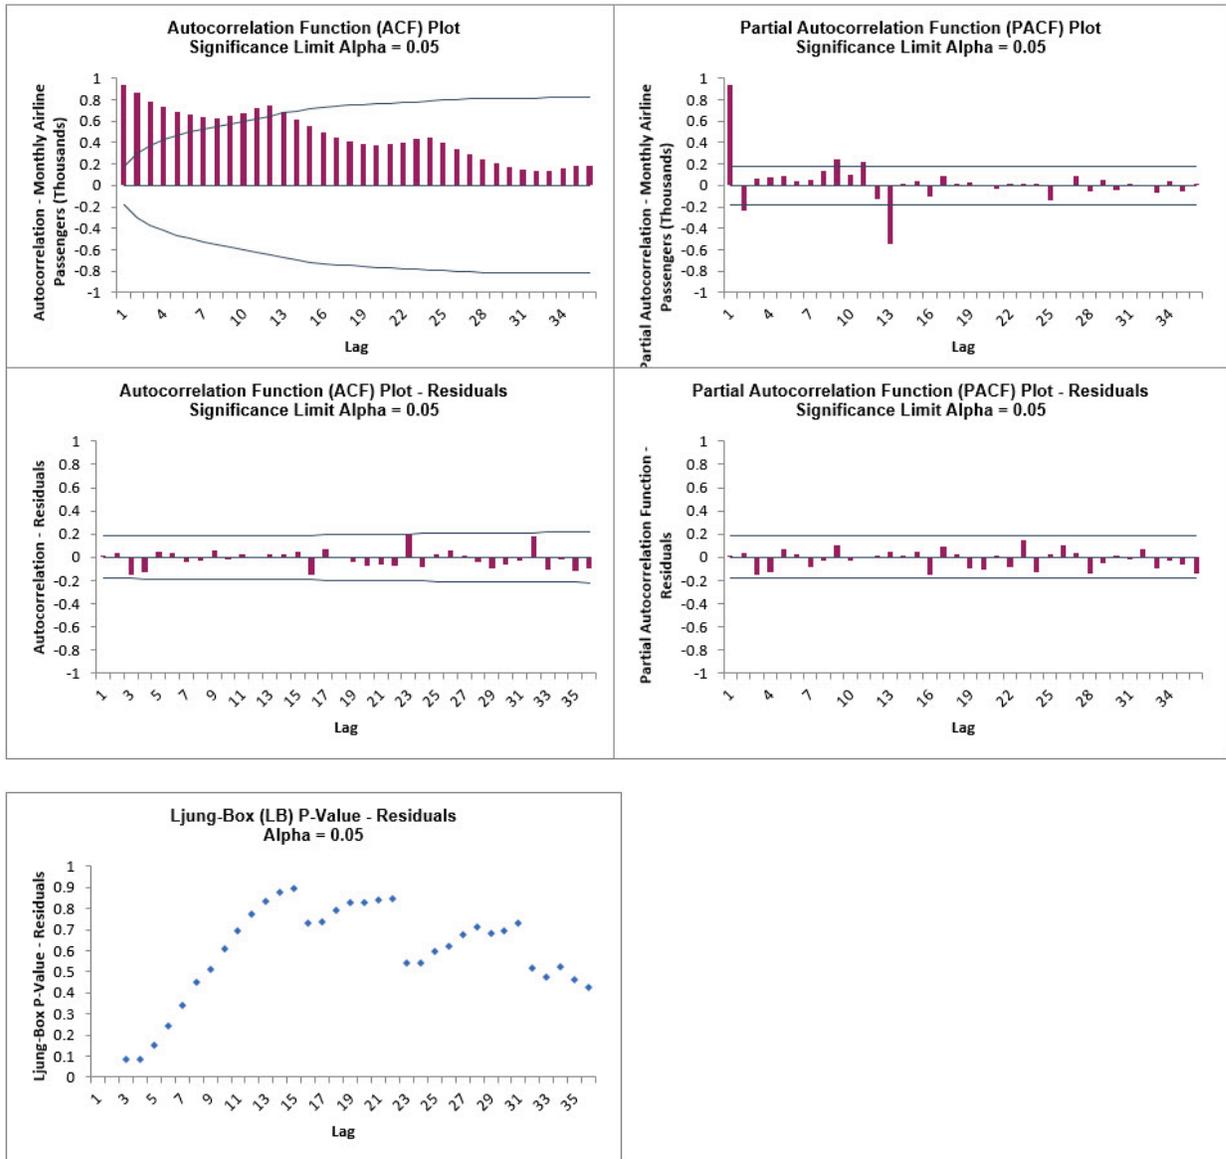
10. The Forecast Accuracy metrics are:

Forecast Accuracy		
Metric	In-Sample (Estimation) One-Step-Ahead Forecast	Out-of-Sample (Withhold) Multi-Step Ahead Forecast
N	119	24
RMSE	7.145744373	26.75002158
MAE	5.494721757	23.0424413
MAPE	2.364132406	4.85347701
MASE	0.192297456	0.806410778

Comparing to [Exponential Smoothing MSD Forecast Accuracy metrics](#), we see that the ARIMA MSD results are approximately the same.

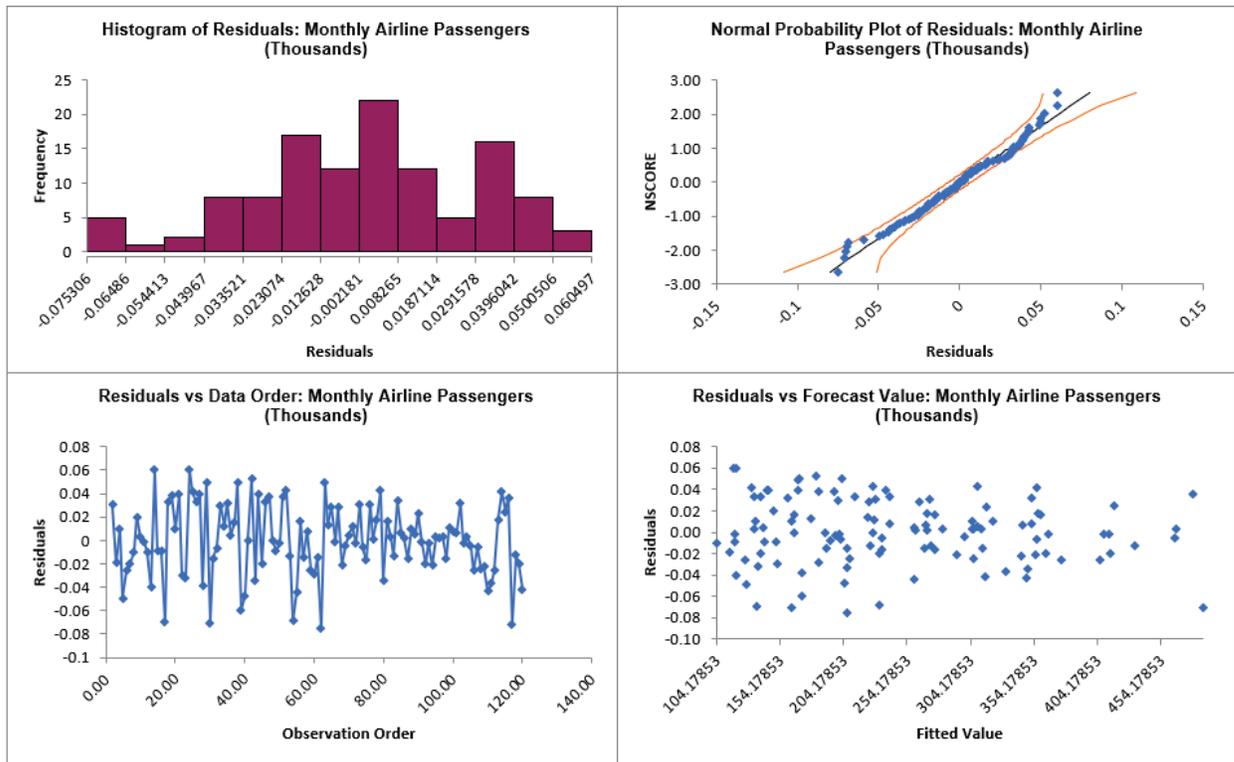
Forecast Accuracy metrics are calculated using the actual raw data versus inverse transformed forecast as displayed in the Forecast Chart and Table, so allow comparison across all model types and transformations.

11. Click on the **ARIMA MSD ACF PACF LB** sheet to view the ACF/PACF/LB Plots:



The ARIMA MSD ACF/PACF Residuals Plots are similar to the [Exponential Smoothing MSD ACF/PACF/LB Plots](#) and indicate that the autocorrelation has been accounted for in the model. All of the Ljung-Box P-Values are blue (i.e., not significant), with P-Values > .05.

12. Click on the **ARIMA MSD Residuals** sheet to view the Residual Plots:



The residuals are approximately normally distributed, with a roughly straight line on the normal probability plot. There are no obvious extreme outliers or patterns in the charts. The residual plots are approximately the same as those given previously in [Exponential Smoothing MSD Residual Plots](#).

Since a Box-Cox transformation was used, the residuals are in Ln transformed units.

Half-Hourly Multiple Seasonal Electricity Demand – Taylor

13. Open **Half-Hourly Multiple Seasonal Electricity Demand - Taylor.xlsx (Sheet 1 tab)**. This is half-hourly electricity demand (MW) in England and Wales from Monday, June 5, 2000 to Sunday, August 27, 2000 (taylor, R forecast). This data has multiple seasonality with frequency = 48 (observations per day) and 336 (observations per week), with a total of 4032 observations. See the [Run Chart](#), [ACF/PACF Plots](#), [Spectral Density Plot](#) and [Seasonal Trend Decomposition Plots](#) for this data.
14. Click **SigmaXL > Time Series Forecasting > ARIMA Forecast > Multiple Seasonal Decomposition Forecast**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
15. Select *Demand*, click **Numeric Time Series Data (Y) >>**. Check **Display ACF/PACF/LB Plots** and **Display Residual Plots**. Check **Specify Model Periods**. Set **Withhold Periods** = 96. Select **Withhold Forecast Type: Multi-Step-Ahead with Prediction Interval at Start of Withhold**. Check **Seasonal Frequency** with **Specify** = 48 336. Leave **Box-Cox Transformation** unchecked. We will use the default **Prediction Interval** = 95.0 %.

ARIMA MSD Forecast

Obs. No.

Numeric Time Series Data (Y) >> Demand

Optional Time Axis Labels >>

<< Remove

OK >>
Cancel
Help

No. of Forecast Periods 24

Prediction Interval 95.0 %

Specify Model Periods

Start Model at Period 1

Withhold Periods 96

End Model at Period

Model Options

Seasonal Frequency

Specify 48 336

Select 4-Quarterly

Automatically Detect

Display ACF/PACF/LB Plots

Display Residual Plots

Box-Cox Transformation

Rounded Lambda

Optimal Lambda

Lambda & Threshold (Shift)

Withhold Forecast Type:

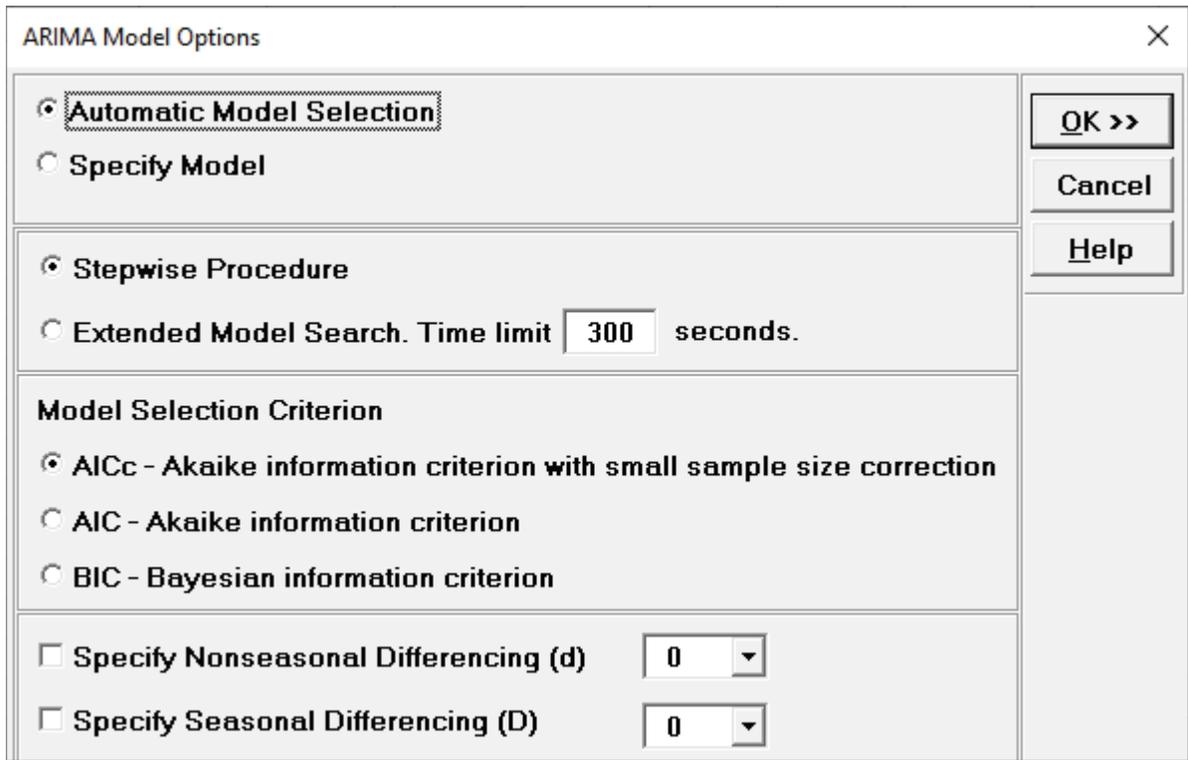
One-Step-Ahead with Prediction Interval at: Start of Withhold

Include in Residuals

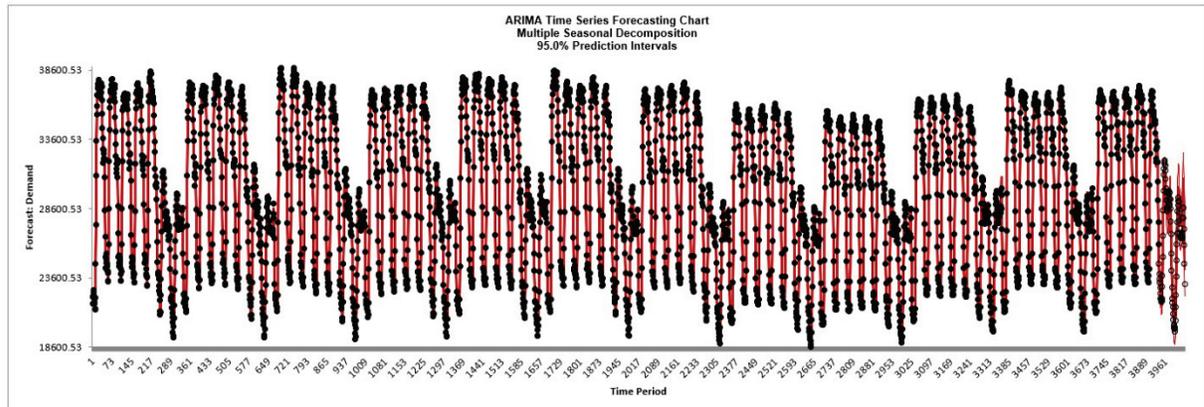
Multi-Step-Ahead with Prediction Interval at Start of Withhold.

Withhold Periods is 2*dominant seasonal frequency (48). Dominant frequency is obtained from the Spectral Density Plot. Start Model at Period = 1 is always greyed out for MSD.

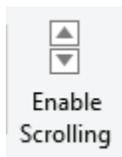
16. Click **Model Options**.



17. We will use the default **Automatic Model Selection** with **AICc** as the **Model Selection Criterion**. Click **OK** to return to the ARIMA MSD Forecast dialog. Click **OK**. The ARIMA forecast report is given:



18. We will want to zoom in on the last 3 days, i.e., 144 half-hourly time periods, using chart scrolling. Click **SigmaXL Chart Tools > Enable Scrolling**

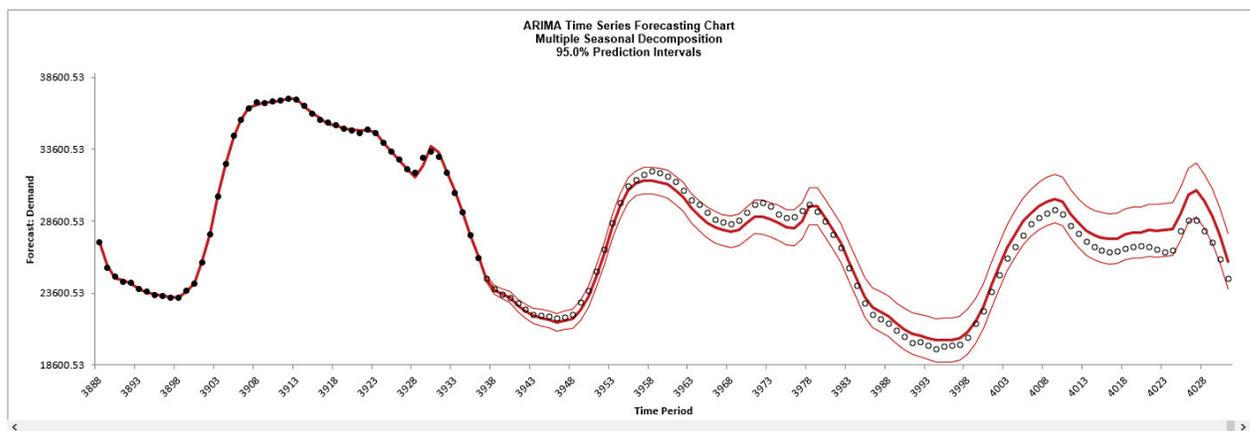


You may be prompted with a warning message that custom formatting on the chart will be cleared. You can avoid seeing this warning by checking **Save this choice as default and do not show this form again**.

19. Click **OK**. The scroll dialog appears allowing you to specify the **Start Period** and **Window Width**. Enter **Start Period** = 3888 and **Window Width** = 144:

At any point, you can click **Restore/Show All Data Points** or **Freeze Chart**. Freezing the chart will remove the scroll and unload the dialog. The scroll dialog will also unload if you change worksheets. To restore the dialog, click **SigmaXL Chart Tools > Enable Scrolling**.

20. Click **OK**. A scroll bar appears beneath the forecast chart. You can also change the **Start Subgroup** and **Window Width** and **Update**.



You can scroll through by clicking to the right or left, with the specified window width of 144.

The blank dots are the data values in the withhold sample with a multi-step forecast and prediction intervals displayed at the start of the withhold sample. The model does quite well at predicting the withhold 96 half-hour demand values. Note that the prediction error increases the further out we predict. The ARIMA MSD Forecast Chart looks approximately the same as the previous [Exponential Smoothing MSD Forecast Chart](#)

Click **Cancel** to exit the scroll dialog.

21. Scroll down to view the ARIMA Model header:

ARIMA Model (Multiple Seasonal Decomposition): Demand - Model Automatically Selected
Model Periods: Model parameter estimates calculated excluding 96 withhold periods.

22. The ARIMA Model Summary is given as:

ARIMA Model Summary	
AR Order (p)	2
I Order (d)	1
MA Order (q)	0
SAR Order (P)	0
SI Order (D)	0
SMA Order (Q)	0
Seasonal Frequency	48, 336 Decomposition
Include Constant	0
No. of Predictors	0
Model Selection Criterion	AICc
Box-Cox Transformation	N/A
Lambda	
Threshold	

This is a summary of the model information for the deseasonalized data: ARIMA (2,1,0) without a constant. Seasonal Frequency = 48, 336 using Decomposition and Model Selection Criterion = "AICc". There are no seasonal terms in the model.

23. The Parameter Estimates for the deseasonalized Demand data are:

Parameter Estimates				
Term	Coefficient	SE Coefficient	T	P
AR_1	0.033677486	0.010694482	3.14905	0.0017
AR_2	-0.022798133	0.0160335	1.42191	0.1551

ARIMA Parameter Estimates include significance tests; P-Values < .05 are significant and highlighted in red. This may be useful for model refinement with multiple predictors but for AR/MA model order selection, minimum AICc should be used, rather than significance tests (see [Kostenko, A.V. and Hyndman, R.J.](#)).

24. The ARIMA Model Statistics are:

ARIMA Model Statistics	
No. of Observations	3936
DF	3933
StDev	98.8439103
Variance	9770.11861
Log-Likelihood	-23647.086
AICc	47315.0387
AIC	47315.0326
BIC	47333.8656

- The number of observations, $n = 4032 - 96$ (*withhold*) = 3936
- Degrees of freedom (DF) = $3936 (n) - 3$ (2 terms in the model, 1 nonseasonal difference) = 3933
- Note that the model statistics are calculated using the deseasonalized data.

25. The Forecast Accuracy metrics are:

Forecast Accuracy		
Metric	In-Sample (Estimation) One-Step-Ahead Forecast	Out-of-Sample (Withhold) Multi-Step Ahead Forecast
N	3935	96
RMSE	101.4595732	775.9894896
MAE	73.7919471	649.1824792
MAPE	0.253330267	2.452663125
MASE	0.040188511	0.353557244

As expected, the **Out-of Sample (Withhold) Multi-Step-Ahead Forecast** errors are larger than the **In-Sample (Estimation) One-Step-Ahead Forecast** errors. Note, if we were primarily interested in a short term one-step ahead forecast, then we would have selected **Withhold Forecast Type: One-Step-Ahead** and the above table would show **Out-of Sample (Withhold) One-Step-Ahead Forecast** errors.

Comparing to the previous [Exponential Smoothing MSD Forecast Accuracy metrics](#), we see that the ARIMA MSD results are approximately the same.

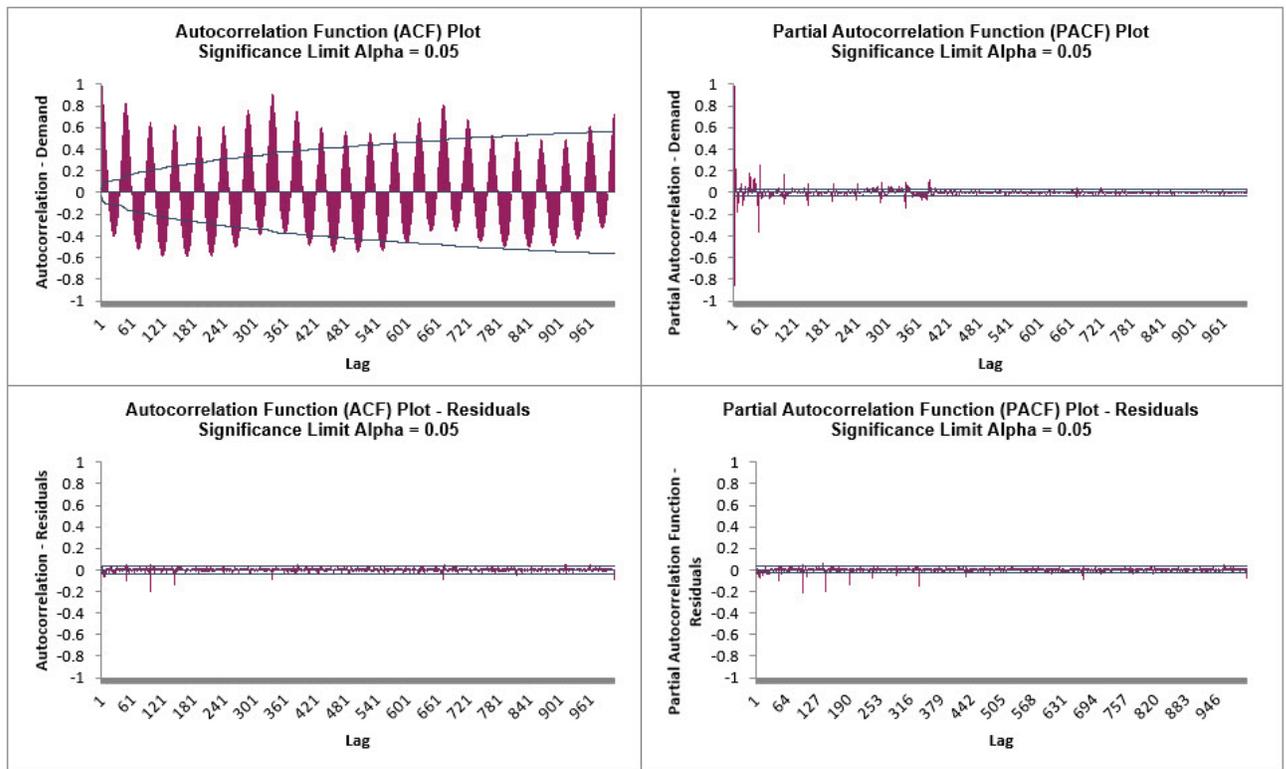
Forecast Accuracy metrics are calculated using the actual raw data versus forecast as displayed in the Forecast Chart and Table so, unlike the model statistics above, allows comparison across all forecast model types and transformations.

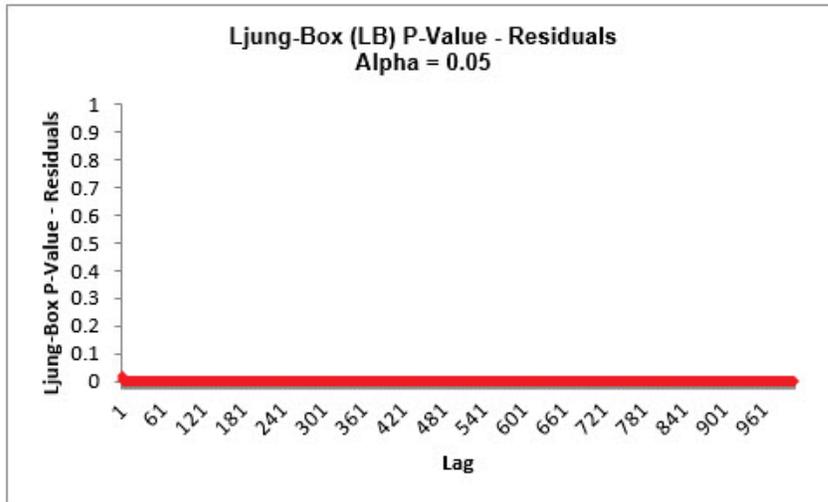
26. The Forecast Table is given as:

Forecast Table				
Period	Withhold Data	Multi-Step Ahead Forecast	Lower 95.0% PI	Upper 95.0% PI
3937	24653	24591.46494	24397.67483	24785.25505
3938	23879	23836.7087	23557.99505	24115.42235
3939	23508	23543.90813	23203.10958	23884.70668
3940	23275	23298.25208	22905.19703	23691.30713
3941	22890	22733.71045	22294.53559	23172.88531
3942	22462	22347.61	21866.71441	22828.5056

These are the same forecast and prediction interval values displayed in the Forecast Chart, but provided for further analysis or charting (e.g., a run chart of the forecast errors). The **Withhold Data** is also displayed.

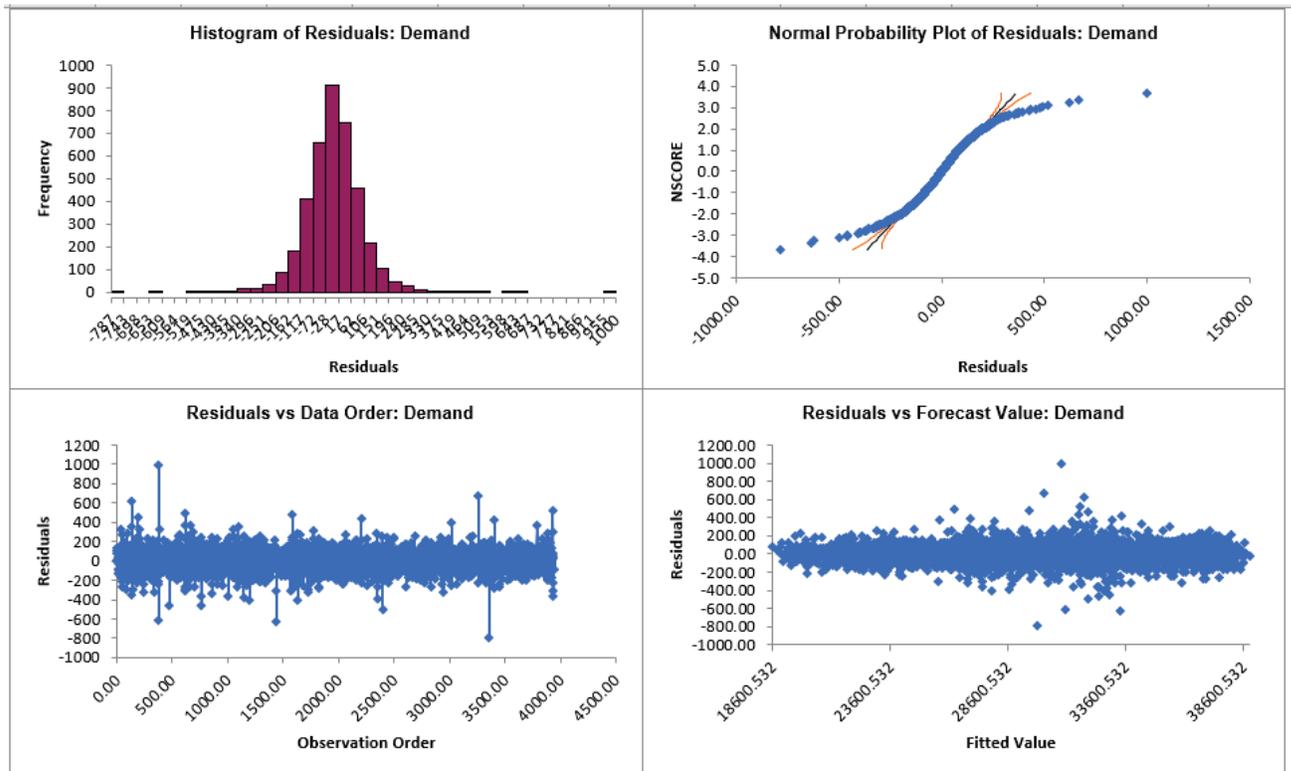
27. Click on the **ARIMA MSD ACF PACF LB** sheet to view the ACF/PACF/LB Plots:





The ACF/PACF Residuals Plots indicate that much of the autocorrelation has been accounted for in the model, but the Ljung-Box plot shows that some significant autocorrelation still remains (the red P-Values are significant at $\alpha=0.05$) - so the model can potentially be improved. This does not mean that the model is a bad model, it can still be very useful for prediction purposes, but the prediction intervals may not provide accurate coverage.

28. Click on the **ARIMA MSD Residuals** sheet to view the Residual Plots:



Similar to the [Exponential Smoothing MSD Residual Plots](#), the residuals are not normally distributed and there are extreme outliers. These should be investigated with a control chart on the residuals. Outliers in Electricity Demand are often explained by Temperature.

ARIMA Control Chart

Statistical Process Control (SPC) For Autocorrelated Data

An Individuals control chart is created using the residuals of the ARIMA forecast model.

The Moving Limits chart uses the one step prediction as the center line, so the control limits will move with the center line. If a Box-Cox transformation is used then an inverse transformation is applied to calculate the control limits.

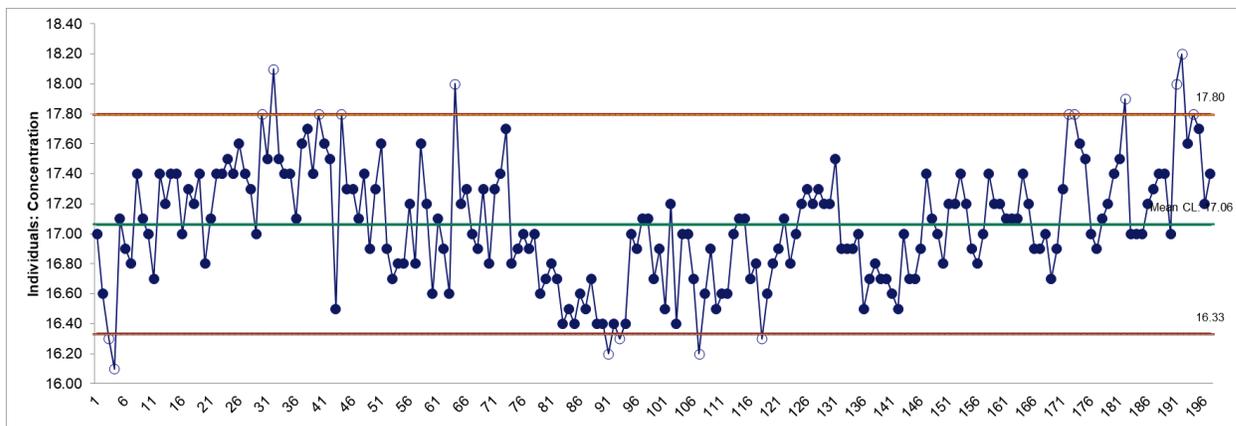
The popular “Add Data”, “Show Last 30” and “Scroll” features in SigmaXL Chart Tools are available for these control charts. The time series models are not refitted, but used to compute the residual values for the new data.

For further details and references, see the Appendix: [Control Charts for Autocorrelated Data](#).

Note that a Moving Range Chart and Tests for Special Causes are not available here, but the user can store and select Residuals, then create with **SigmaXL > Control Charts > Individuals & Moving Range**.

Chemical Process Concentration - Series A

1. Open **Chemical Process Concentration – Series A.xlsx (Sheet 1 tab)**. This is the Series A data from Box and Jenkins, a set of 197 concentration values from a chemical process taken at two-hour intervals. See the [Run Chart](#), [ACF/PACF Plots](#), [Spectral Density Plot](#) and [Seasonal Trend Decomposition Plots](#) for this data.
2. Earlier we saw that this process has significant autocorrelation. In order to see the impact on a control chart, we will construct an Individuals chart on the raw data. Click **SigmaXL > Control Charts > Individuals**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
3. Select *Concentration*, click **Numeric Data Variable (Y) >>**. Click **OK**. An Individuals Control Chart is produced:



There are 17 out-of-control data points, largely due to the autocorrelation. Searching for assignable causes using this chart as is, would be futile.

4. Now click **Sheet 1 tab** and **SigmaXL > Time Series Forecasting > ARIMA Control Chart > Control Chart**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
5. Select *Concentration*, click **Numeric Time Series Data (Y) >>**. Uncheck **Display ACF/PACF/LB Plots**. Leave **Display Residual Plots**, **Specify Model Periods**, **Seasonal Frequency** and **Box-Cox Transformation** unchecked.

ARIMA Control Chart

Observation No.

Numeric Time Series Data (Y) >> Concentration

Optional Time Axis Labels >>

<< Remove

OK >> Cancel Help

Specify Model Periods

Start Model/Control Limit Calculations at Period 1

Withhold Periods 0

End Model/Control Limit Calculations at Period

Model Options

Seasonal Frequency

Specify 12

Select 4-Quarterly

Automatically Detect

Display ACF/PACF/LB Plots

Display Residual Plots

Box-Cox Transformation

Rounded Lambda

Optimal Lambda

Lambda & Threshold (Shift)

- Click **Model Options**. Select **Specify Model**. Specify **Nonseasonal Order I – Integrated/Differencing (d) = 1** and **MA – Moving Average (q) = 1**. Leave **Include Constant** unchecked.

ARIMA Model Options

Automatic Model Selection

Specify Model

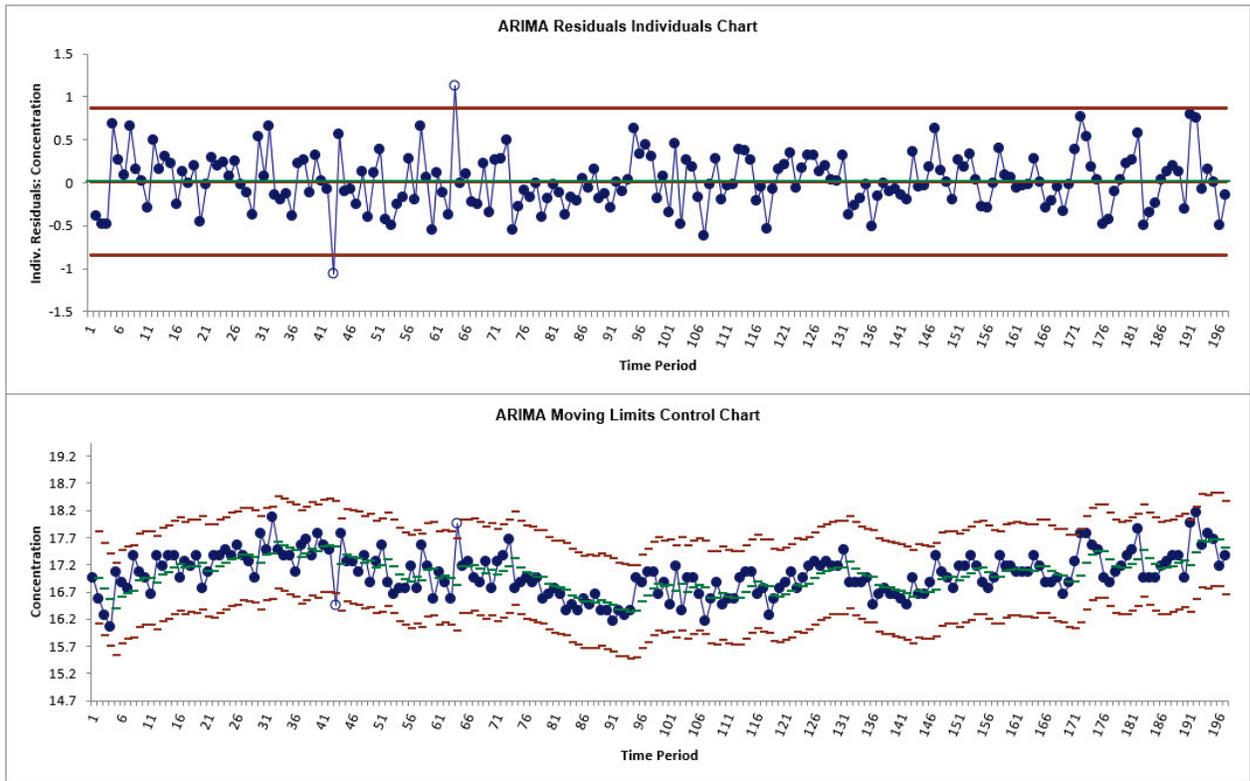
OK >> Cancel Help

Nonseasonal Order		Seasonal Order	
AR - Autoregressive (p)	0	SAR - Seasonal Autoregressive (P)	0
I - Integrated/Differencing (d)	1	SI - Seasonal Integrated/Differencing (D)	0
MA - Moving Average (q)	1	SMA - Seasonal Moving Average (Q)	0

Include Constant (Mean if d + D = 0; Trend/Drift if d or D = 1)

Typically, one would use **Automatic Model Selection**, but we want to demonstrate the ARIMA control chart using the ARIMA (0,1,1) model as used [above](#), and to compare to the equivalent [Exponential Smoothing Control Charts](#).

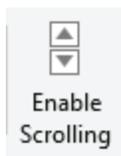
7. Click **OK** to return to the ARIMA Control Chart dialog. Click **OK**. The ARIMA control charts are produced:



Now we only have two out-of-control data points on the Individuals chart to investigate. The Moving Limits chart uses the one step prediction as the center line, so the control limits move with the center line.

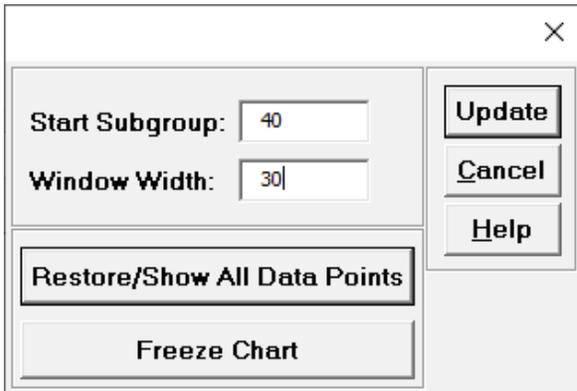
As expected, these are approximately the same as the [Exponential Smoothing Control Charts](#), with slight differences on the initial values. The ARIMA Residuals Individuals Chart does not show a data point at Time Period = 1 due to the differencing.

8. We will scroll through the chart data points as done in the Exponential Smoothing Control Charts. Click **SigmaXL Chart Tools > Enable Scrolling**

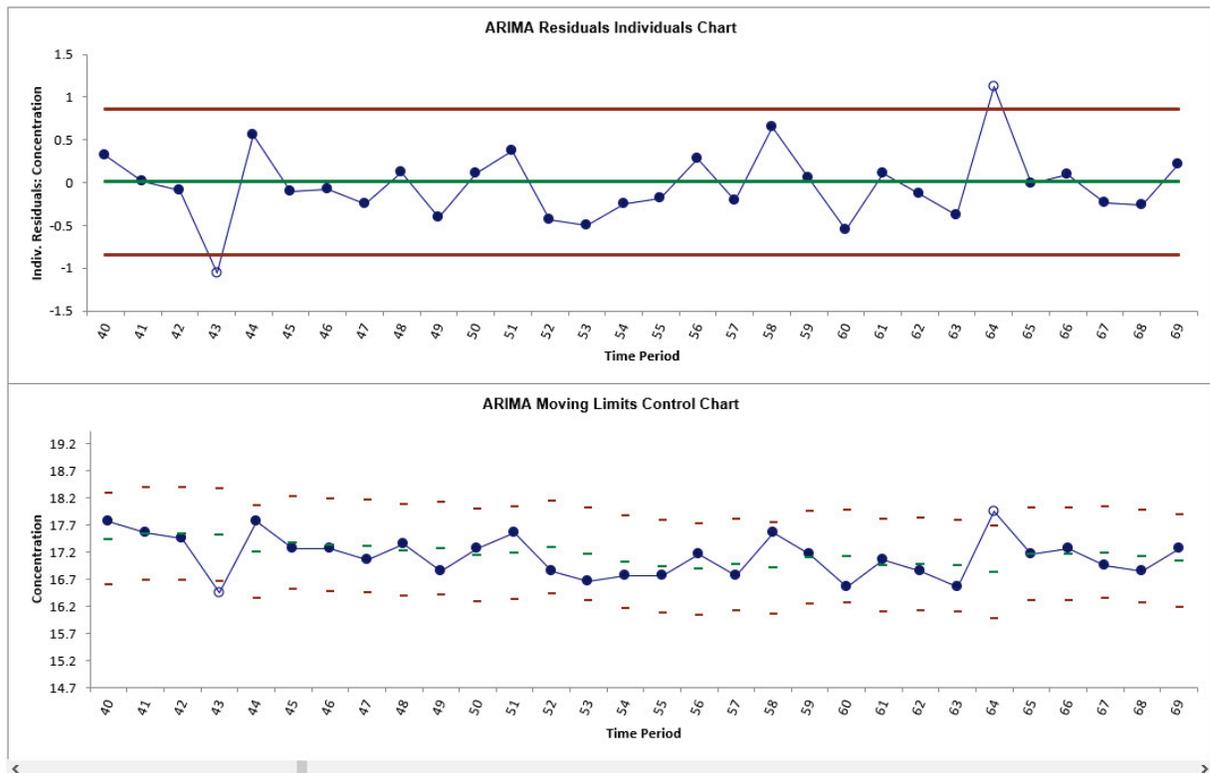


You may be prompted with a warning message that custom formatting on the chart will be cleared. You can avoid seeing this warning by checking **Save this choice as default and do not show this form again**.

- Click **OK**. The scroll dialog appears allowing you to specify the **Start Subgroup** and **Window Width**. Enter **Start Subgroup** = 40 and **Window Width** = 30 to view the two out-of-control data points.



- Click **OK**. This allows us to zoom in on the out-of-control points at 43 and 64.



Observation **43** is lower than expected from the ARIMA forecast model. Observation **64** is higher than expected.

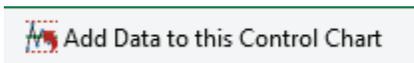
- Click **Cancel** to exit the scroll dialog.
- Now we will add a new data point to the Series A Concentration Data. The residuals will be computed using the same model as above without re-estimation of the model parameters or recalculation of the control limits. This is also known as the “Phase II” application of a Control Chart, where an out-of-control signal should lead to an investigation into the assignable cause

and corrective action or process adjustment applied. Click Sheet1, enter the value 19 as shown in cell **B199** (and optionally Observation number 198 in cell **A199**).

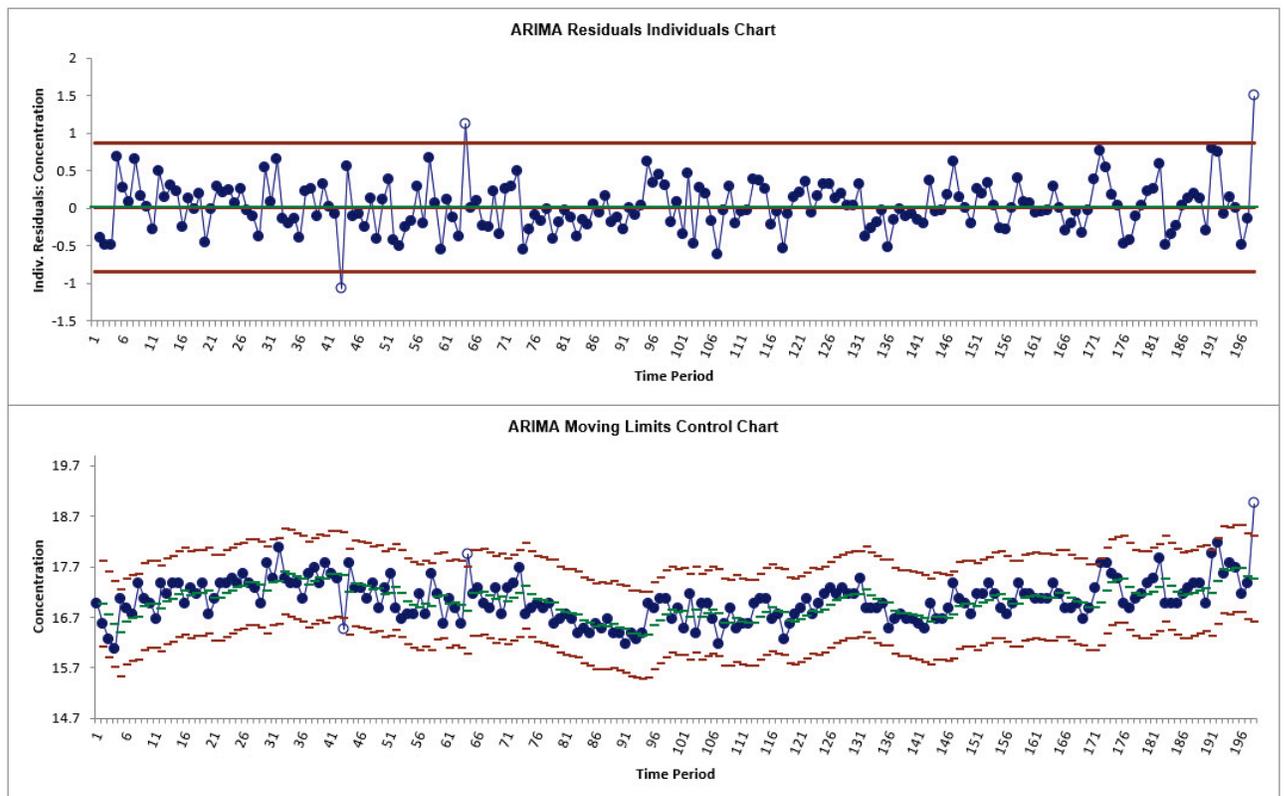
196	195	17.7
197	196	17.2
198	197	17.4
199	198	19

13. Click **ARIMA Control Charts** tab (if more than one control chart sheet exists in the workbook, please select the chart where the data will be added).

14. Click **SigmaXL Chart Tools > Add Data to this Control Chart**



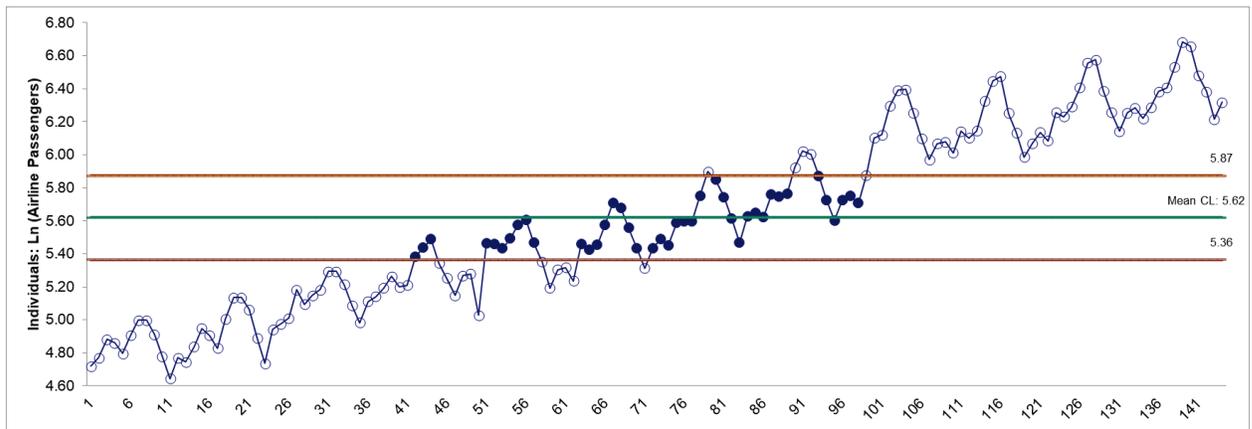
15. The Residuals Individuals Control Chart and Moving Limits Charts are now updated with the new data, showing this as an out-of-control data point:



We recommend renaming the workbook to **Chemical Process Concentration – Series A_AddData2.xlsx**, so that later use of the Concentration data does not include the added data point.

Monthly Airline Passengers – Modified for Control Charts

16. Open **Monthly Airline Passengers – Modified for Control Charts.xlsx (Sheet 1 tab)**. This is based on the Series G data from Box and Jenkins, monthly total international airline passengers for January 1949 to December 1960. A Ln transformation is applied (avoiding the need for a Box-Cox transformation), a negative outlier is added at 50 (-.25) and a level shift applied (+.25), starting at 100. Coded variables were added to help distinguish an outlier versus a shift and they will be analyzed later using ARIMA Forecast with Predictors.
17. Earlier we saw that this process has significant autocorrelation with a strong trend and seasonality. In order to see the impact on a control chart, we will construct an Individuals chart on the raw data. Click **SigmaXL > Control Charts > Individuals**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
18. Select *Ln (Airline Passengers)*, click **Numeric Data Variable (Y) >>**. Click **OK**. An Individuals Control Chart is produced:



With strong trend, seasonality and positive autocorrelation, this control chart is meaningless.

19. Now click **Sheet 1 tab** and **SigmaXL > Time Series Forecasting > ARIMA Control Chart > Control Chart**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
20. Select *Ln(Airline Passengers-Modified)*, click **Numeric Time Series Data (Y) >>**. Uncheck **Display ACF/PACF/LB Plots** and **Display Residual Plots**. Check **Seasonal Frequency** with **Specify = 12**. Leave **Specify Model Periods** and **Box-Cox Transformation** unchecked.

ARIMA Control Chart

Obs. No.
 Outlier 50
 Shift 50
 Outlier 100
 Shift 100

Numeric Time Series Data (Y) >> Ln (Airline Passengers-Mc)

Optional Time Axis Labels >>

<< Remove

OK >>
 Cancel
 Help

Specify Model Periods

Start Model/Control Limit Calculations at Period

Withhold Periods

End Model/Control Limit Calculations at Period

Model Options

Seasonal Frequency

Specify

Select

Automatically Detect

Display ACF/PACF/LB Plots

Display Residual Plots

Box-Cox Transformation

Rounded Lambda

Optimal Lambda

Lambda & Threshold (Shift)

21. Click **Model Options**.

ARIMA Model Options

Automatic Model Selection

Specify Model

Stepwise Procedure

Extended Model Search. Time limit seconds.

Model Selection Criterion

AICc - Akaike information criterion with small sample size correction

AIC - Akaike information criterion

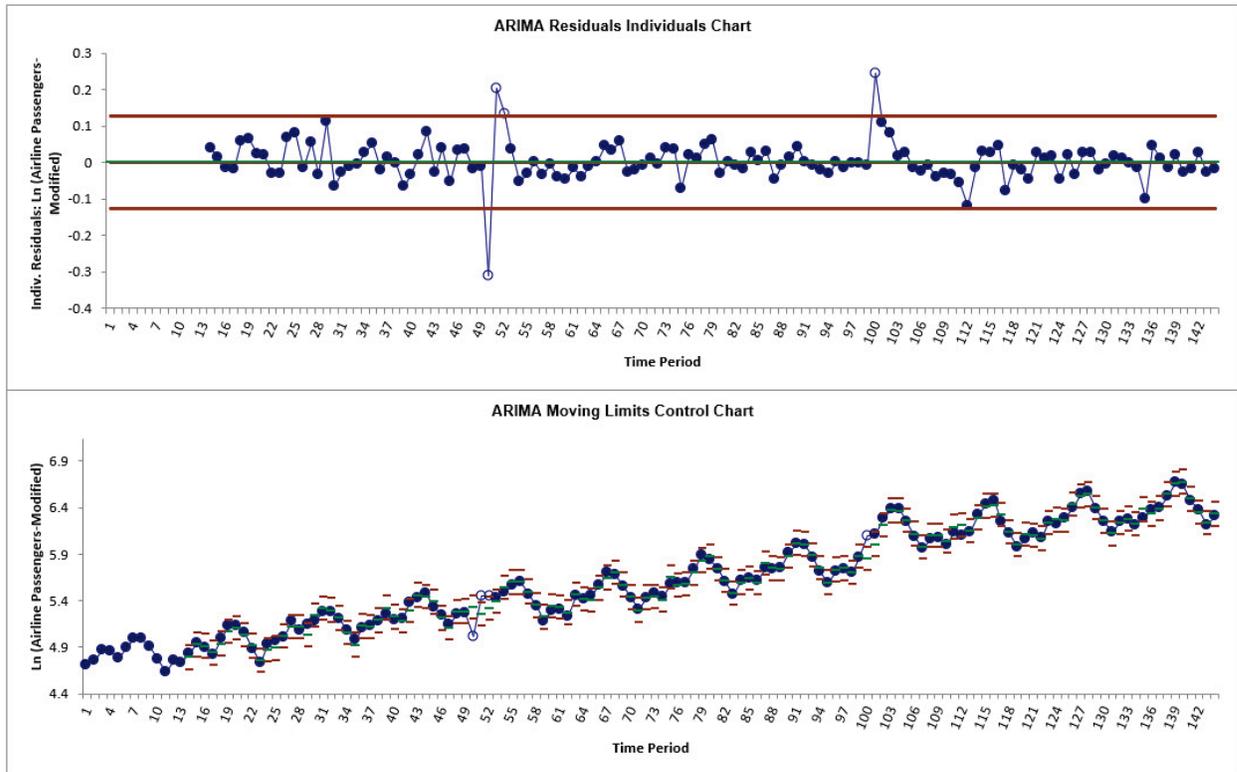
BIC - Bayesian information criterion

Specify Nonseasonal Differencing (d)

Specify Seasonal Differencing (D)

OK >>
 Cancel
 Help

22. We will use the default **Automatic Model Selection** with **AICc** as the **Model Selection Criterion**. Click **OK** to return to the ARIMA Control Chart dialog. Click **OK**. The ARIMA control charts are produced:



Now we can clearly see the out-of-control data points at 50, 51, 52 and 100 on the Individuals chart. These are approximately the same as those given in [Exponential Smoothing Control Charts](#) (52 was in-control). The ARIMA Residuals Individuals Chart does not show data points at Time Periods 1 to 13 due to the nonseasonal and seasonal differencing.

23. In order to view the points on the Moving Limits chart we will use scrolling. Click **SigmaXL Chart Tools > Enable Scrolling**



You may be prompted with a warning message that custom formatting on the chart will be cleared. You can avoid seeing this warning by checking **Save this choice as default and do not show this form again**.

24. Click **OK**. The scroll dialog appears allowing you to specify the **Start Subgroup** and **Window Width**. Enter **Start Subgroup** = 40 and **Window Width** = 20 to view the first two out-of-control data points.

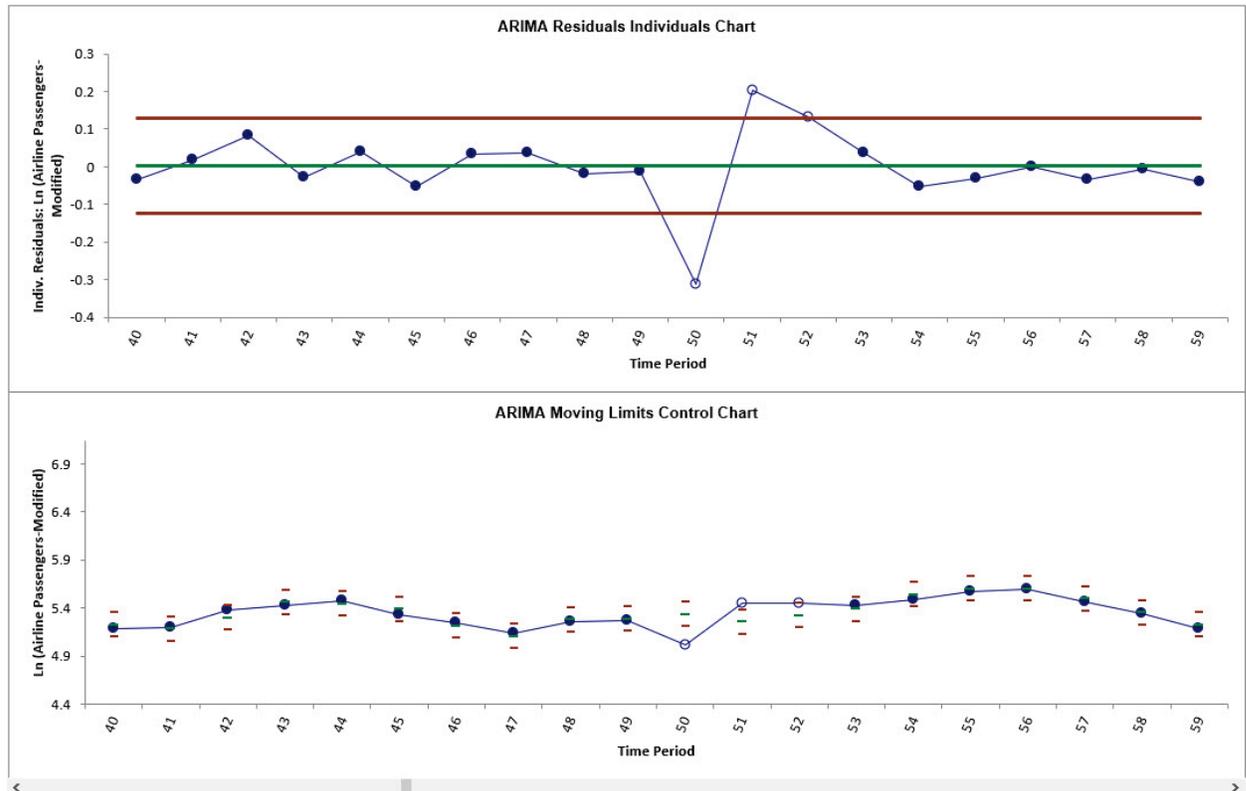
Start Subgroup:
OK >>

Window Width:
Cancel

Restore/Show All Data Points
Help

Freeze Chart

25. Click **OK**. This allows us to zoom in on the out-of-control points at 50, 51 and 52.



Observation **50** is lower than expected from the ARIMA forecast model. Observations **51** and **52** are higher than expected.

Tip: Scrolling keeps the original Y axis minimum and maximum setting. You may wish to change this to auto by clicking on the Y axis, right click **Format Axis**, click Bounds Minimum **Reset** and Bounds Maximum **Reset**. This changes the axis settings to Auto so when you scroll or Update the Y axis will automatically adjust as well.

26. Now enter **Start Subgroup** = 90 and **Window Width** = 20 to view the third out-of-control data point.

×

Start Subgroup:

Window Width:

Restore/Show All Data Points

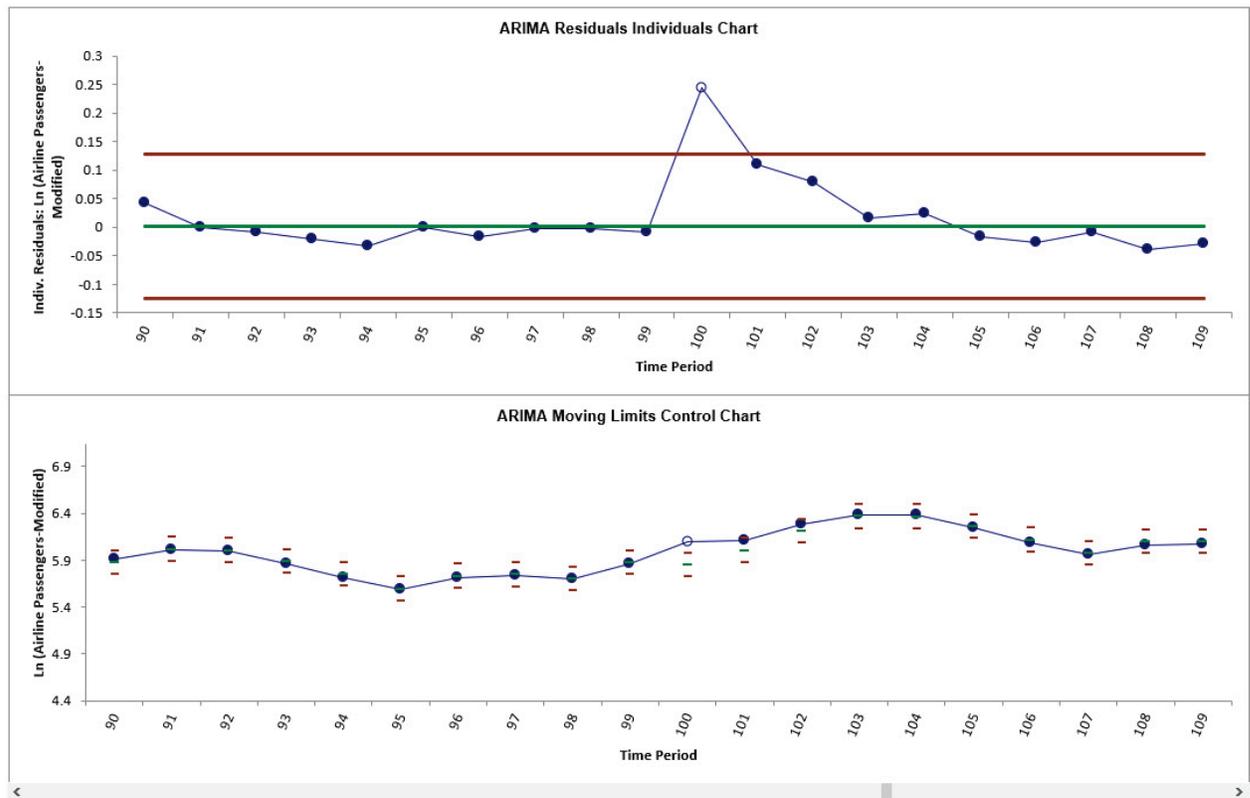
Freeze Chart

Update

Cancel

Help

27. Click **Update**.



Observation **100** is higher than expected from the ARIMA forecast model. Later investigation will reveal that this is a shift in the mean.

28. Click **Cancel** to exit the scroll dialog.

29. Scroll down to view the ARIMA Model header:

ARIMA Model: Ln (Airline Passengers-Modified) - Model Automatically Selected
Model Periods: All observations are used in the ARIMA model estimation. No withhold periods available for out-of-sample forecast accuracy evaluation.

30. The ARIMA Model Summary is:

ARIMA Model Summary	
AR Order (p)	0
I Order (d)	1
MA Order (q)	1
SAR Order (P)	0
SI Order (D)	1
SMA Order (Q)	1
Seasonal Frequency	12
Include Constant	0
No. of Predictors	0
Model Selection Criterion	AICc
Box-Cox Transformation	N/A
Lambda	
Threshold	

The ARIMA (0,1,1) (0,1,1) model (with no constant) was automatically selected based on the AICc criterion. This is the same as we obtained [earlier](#) when using a withhold sample and agrees with the results shown in the [ACF/PACF Plots](#) of the manually differenced data and is the model used in Box and Jenkins, Chapter 9., “Analysis of Seasonal Time Series”.

There are no predictors. Seasonal Frequency = 12; Model Selection Criterion = “AICc” and Box-Cox Transformation = “N/A”.

31. The Parameter Estimates and ARIMA Model Statistics are slightly different than our [earlier](#) analysis because we have introduced an outlier and a shift, as well here we are using all of the data, i.e., there are no withhold periods. Note that earlier we used a Box-Cox Transformation with Lambda=0 and here we are using Ln of the data.

32. The Forecast Accuracy metrics are given as:

Forecast Accuracy		
Metric	In-Sample (Estimation) One-Step-Ahead Forecast	Out-of-Sample (Withhold) One-Step-Ahead Forecast
N	131	
RMSE	0.055615558	
MAE	0.036285476	
MAPE	0.648152199	
MASE	0.249152676	

Note that these forecast errors are very different than our [earlier](#) analysis, where the forecast errors were calculated on the raw data versus final predicted values, but here we are using Ln of the Airline Passenger data.

33. Now we will reanalyze the data to determine whether the out-of-control signals are due to an outlier or process shift. Click the **Sheet 1** tab.

Obs. No.	Ln (Airline Passengers-Modified)	Outlier 50	Shift 50	Outlier 100	Shift 100
49	5.278114659	0	0	0	0
50	5.028114659	1	1	0	0
51	5.463831805	0	1	0	0
52	5.459585514	0	1	0	0
53	5.433722004	0	1	0	0
54	5.493061443	0	1	0	0
55	5.575949103	0	1	0	0
99	5.874930731	0	1	0	0
100	6.10220248	0	1	1	1
101	6.122117789	0	1	0	1
102	6.295005314	0	1	0	1
103	6.392037406	0	1	0	1
104	6.396329258	0	1	0	1
105	6.251414878	0	1	0	1

The first out-of-control signal occurred at Obs. No = 50, so two coded variables were created. *Outlier50* is coded so that Obs. No. 50 = 1 and all other values are 0. *Shift50* is coded so that Obs. No. 1 to 49 = 0, and Obs. No. 50 to 144 are 1.

To simplify the analysis, we will assume that the out-of-control signals at 51 and 52 are related to 50, so will ignore those. The next out-of-control signal occurs at 100, so similarly another two coded variables were created. *Outlier100* is coded so that Obs. No. 100 = 1 and all other values are 0. *Shift100* is coded so that Obs. No. 1 to 99 = 0, and Obs. No. 100 to 144 are 1.

Since these are coded as 0, 1 we will treat them as Continuous.

34. Click **SigmaXL > Time Series Forecasting > ARIMA Forecast > Forecast with Predictors**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
35. Select *Ln(Airline Passengers – Modified)*, click **Numeric Time Series Data (Y) >>**; select *Outlier50* to *Shift100*, click **Optional Continuous Pred. >>** (. Uncheck **Display ACF/PACF/LB Plots** and **Display Residual Plots**. Check **Seasonal Frequency** with **Specify = 12**. Leave **Specify Model Periods** and **Box-Cox Transformation** unchecked. We will use the default **Prediction Interval = 95.0 %**.

ARIMA with Predictors Forecast

Obs. No.

Numeric Time Series Data (Y) >> Ln (Airline Passengers-Mc)

Optional Time Axis Labels >>

Optional Continuous **Pred.** (X) >> Outlier 50
Shift 50
Outlier 100

Optional Categorical **Pred.** (X) >>

<< **Remove**

No. of Forecast Periods 24

Prediction Interval 95.0 %

Specify Model Periods

Start Model at Period 1

Withhold Periods 0

End Model at Period

Model Options

Seasonal Frequency

Specify 12

Select 4 - Quarterly

Automatically Detect

Display ACF/PACF/LB Plots

Display Residual Plots

Box-Cox Transformation

Rounded Lambda

Optimal Lambda

Lambda & Threshold (Shift)

OK >>
Cancel
Help

36. Click **Model Options**. With the addition of the coded predictors, we want to ensure that the ARIMA model is the same as that used in the control chart. Select **Specify Model**. Specify **Nonseasonal Order I – Integrated/Differencing (d) = 1** and **MA – Moving Average (q) = 1**. Specify **Seasonal Order SI – Seasonal Integrated/Differencing (D) = 1** and **SMA – Seasonal Moving Average (Q) = 1**. Leave **Include Constant** unchecked.

ARIMA Model Options

Automatic Model Selection

Specify Model

Nonseasonal Order

AR - Autoregressive (p) 0

I - Integrated/Differencing (d) 1

MA - Moving Average (q) 1

Seasonal Order

SAR - Seasonal Autoregressive (P) 0

SI - Seasonal Integrated/Differencing (D) 1

SMA - Seasonal Moving Average (Q) 1

Include Constant (Mean if d + D = 0; Trend/Drift if d or D = 1)

OK >>
Cancel
Help

37. Click **OK** to return to the ARIMA Forecast dialog. Click **OK**. The ARIMA forecast report is given.

38. Scroll down to view the Parameter Estimates.

Parameter Estimates				
Term	Coefficient	SE Coefficient	T	P
MA_1	0.395854258	0.086277876	4.588132	0.0000
SMA_1	0.556926317	0.073521579	7.575005	0.0000
Outlier 50	-0.296267312	0.032292885	9.174384	0.0000
Shift 50	0.044130277	0.035505653	1.242908	0.2162
Outlier 100	-0.000597974	0.031822469	0.018791	0.9850
Shift 100	0.249993132	0.035384072	7.065132	0.0000

Now we can see that *Outlier50* and *Shift100* are significant denoting Obs. No. 50 as an outlier and 100 as a shift. This is, of course, what we expected since that's how the Ln Airline Passenger data was modified.

This method to identify outlier versus shift is intended as a complement to process knowledge and the search for assignable causes used in classical SPC.

ARIMA Control Chart with Predictors

An Individuals control chart is created using the residuals of the ARIMA with Predictors forecast model.

The ARIMA with Predictors model supports continuous or categorical predictors. See [ARIMA Forecast with Predictors](#) for more information.

The Moving Limits chart uses the one step prediction as the center line, so the control limits will move with the center line. If a Box-Cox transformation is used then an inverse transformation is applied to calculate the control limits.

The popular “Show Last 30” and “Scroll” features in SigmaXL Chart Tools are available for these control charts. Currently, the “Add Data” feature is not supported in ARIMA Control Chart with Predictors.

For further details and references, see the Appendix: [Control Charts for Autocorrelated Data](#).

Note that a Moving Range Chart and Tests for Special Causes are not available here, but the user can store and select Residuals, then create with **SigmaXL > Control Charts > Individuals & Moving Range**.

Daily Electricity Demand with Predictors – ElecDaily

1. Open **Daily Electricity Demand with Predictors – ElecDaily.xlsx (Sheet 1 tab)**. This is daily electricity demand (GW) for the state of Victoria, Australia, every day during 2014 (see Hyndman, fpp2, Section 9.3, **Example: Forecasting electricity demand**, <https://otexts.com/fpp2/forecasting.html>). Temp (C) is the maximum daily temperature in degrees Celsius for the city of Melbourne. TempSq is Temperature squared. WorkDay takes on the value 1 on work days and 0 otherwise. This data has a seasonal frequency = 7 (observations per week). See the [Run Chart](#), [ACF/PACF Plots](#), [Spectral Density Plot](#) and [Seasonal Trend Decomposition Plots](#) for this data.
2. Click **SigmaXL > Time Series Forecasting > ARIMA Control Chart > Control Chart with Predictors**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
3. Select *Demand*, click **Numeric Time Series Data (Y) >>**; select *Date*, click **Optional X-Axis Labels >>**; select *Temp (C)* and *TempSq*, click **Optional Continuous Pred. >>**; select *WorkDay*, click **Optional Categorical Pred. >>**. Uncheck **Display ACF/PACF/LB Plots** and **Display Residual Plots**. Check **Seasonal Frequency with Select = 7 - Daily (or Specify = 7)**. Leave **Specify Model Periods** and **Box-Cox Transformation** unchecked.

ARIMA with Predictors Control Chart

Numeric Time Series Data (Y) >>	Demand	<input type="button" value="OK >>"/> <input type="button" value="Cancel"/> <input type="button" value="Help"/>
Optional Time Axis Labels >>	Date	
Optional Continuous Pred. (X) >>	Temp (C) TempSq	
Optional Categorical Pred. (X) >>	WorkDay	
<input type="button" value="<< Remove"/>		

<input type="checkbox"/> Specify Model Periods Start Model/Control Limit Calculations at Period <input type="text" value="1"/> <input checked="" type="radio"/> Withhold Periods <input type="text" value="0"/> End Model/Control Limit Calculations at Period <input type="text"/>	Model Options	<input type="checkbox"/> Display ACF/PACF/LB Plots <input type="checkbox"/> Display Residual Plots <input type="checkbox"/> Box-Cox Transformation <input checked="" type="radio"/> Rounded Lambda <input type="radio"/> Optimal Lambda <input type="radio"/> Lambda & Threshold (Shift)
	<input checked="" type="checkbox"/> Seasonal Frequency <input type="radio"/> Specify <input type="text" value="12"/> <input checked="" type="radio"/> Select <input type="text" value="7-Daily"/> <input type="button" value="v"/> <input type="radio"/> Automatically Detect	

- Click **Model Options**. Select **Automatic Model Selection**. We will use the defaults: **Stepwise Procedure** and **Model Selection Criterion: AICc – Akaike information criterion with small sample size correction**, leave **Specify Nonseasonal Differencing (d)** and **Specify Seasonal Differencing (D)** unchecked.

ARIMA Model Options

Automatic Model Selection

Specify Model

Stepwise Procedure

Extended Model Search. Time limit seconds.

Model Selection Criterion

AICc - Akaike information criterion with small sample size correction

AIC - Akaike information criterion

BIC - Bayesian information criterion

Specify Nonseasonal Differencing (d) ▾

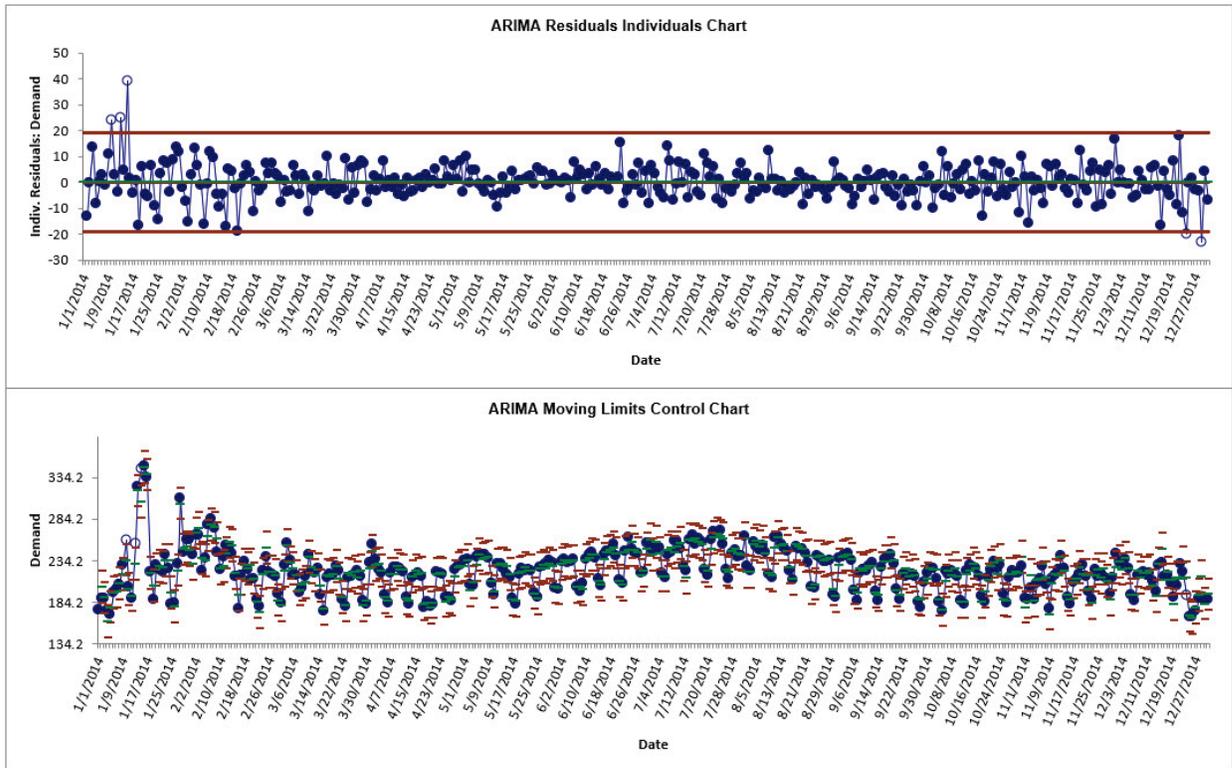
Specify Seasonal Differencing (D) ▾

OK >>

Cancel

Help

- Click **OK** to return to the ARIMA with Predictors Control Chart dialog. Click **OK**. This is a complex model, so computation time will be approximately one to two minutes. The ARIMA with Predictors control charts are produced:



We can see out of control signals in mid-January with higher demand than the model predicted and late December with lower demand than the model predicted.

6. Scroll down to view the ARIMA Model header:

ARIMA Model: Demand - Model Automatically Selected

Model Periods: All observations are used in the ARIMA model estimation. No withhold periods available for out-of-sample forecast accuracy evaluation.

7. The ARIMA Model Summary is given as:

ARIMA Model Summary	
AR Order (p)	2
I Order (d)	1
MA Order (q)	2
SAR Order (P)	2
SI Order (D)	0
SMA Order (Q)	0
Seasonal Frequency	7
Include Constant	0
No. of Predictors	3
Model Selection Criterion	AICc
Box-Cox Transformation	N/A
Lambda	
Threshold	

This is a summary of the model information: ARIMA (2,1,2) (2,0,0) with no constant and 3 predictors. Seasonal Frequency = 7; Model Selection Criterion = "AICc" and Box-Cox Transformation = "N/A" because Box-Cox Transformation was unchecked.

This is the same model used [previously](#) in ARIMA MSD Forecast, and agrees with the model given in Hyndman, fpp2, Section 9.3, **Example: Forecasting electricity demand**, <https://otexts.com/fpp2/forecasting.html>.

ARIMA Multiple Seasonal Decomposition (MSD) Control Chart

ARIMA does not have a theoretical frequency limit, but for computational efficiency and to minimize the potential loss of observations through differencing, we recommend using ARIMA – Multiple Seasonal Decomposition (MSD) for seasonal frequency greater than 52 (or with multiple frequencies). The seasonal component is first removed through decomposition, a nonseasonal ARIMA model fitted to the remainder (+trend), and then the seasonal component is added back in.

As the name implies, Multiple Seasonal Decomposition (MSD) also accommodates multiple seasonality, for example the half-hourly data with a seasonal frequency of 48 observations per day and 336 observations per week.

An Individuals control chart of the residuals is created for this forecast method. The Moving Limits chart uses the one step prediction as the center line, so the control limits will move with the center line. If a Box-Cox transformation is used then an inverse transformation is applied to calculate the control limits.

The popular “Add Data”, “Show Last 30” and “Scroll” features in SigmaXL Chart Tools are available for these control charts. For “Add Data”, the time series models are not refitted, but used to compute the residual values for the new data.

Monthly Airline Passengers – Modified for Control Charts

1. Open **Monthly Airline Passengers – Modified for Control Charts.xlsx (Sheet 1 tab)**. This is based on the Series G data from Box and Jenkins, monthly total international airline passengers for January 1949 to December 1960. A Ln transformation is applied (avoiding the need for a Box-Cox transformation), a negative outlier is added at 50 (-.25) and a level shift applied (+.25), starting at 100. The Multiple Seasonal Decomposition (MSD) option is not necessary for this data, but by way of introduction, we will use this to compare to the previous analysis.
2. Click **SigmaXL > Time Series Forecasting > ARIMA Control Chart > Multiple Seasonal Decomposition Control Chart**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
3. Select *Ln(Airline Passengers-Modified)*, click **Numeric Time Series Data (Y) >>**. Uncheck **Display ACF/PACF/LB Plots** and **Display Residual Plots**. Check **Seasonal Frequency** with **Specify = 12**. Leave **Specify Model Periods** and **Box-Cox Transformation** unchecked.

ARIMA MSD Control Chart

Obs. No. Outlier 50 Shift 50 Outlier 100 Shift 100	Numeric Time Series Data (Y) >>	Ln (Airline Passengers-Mc	OK >>
	Optional Time Axis Labels >>		Cancel
	<< Remove		Help

<input type="checkbox"/> Specify Model Periods	Model Options	<input type="checkbox"/> Display ACF/PACF/LB Plots
Start Model/Control Limit Calculations at Period <input type="text" value="1"/>	Seasonal Frequency	<input type="checkbox"/> Display Residual Plots
<input checked="" type="radio"/> Withhold Periods <input type="text" value="0"/>	<input checked="" type="radio"/> Specify <input type="text" value="12"/>	<input type="checkbox"/> Box-Cox Transformation
End Model/Control Limit Calculations at Period <input type="text"/>	<input type="radio"/> Select <input type="text" value="4 - Quarterly"/>	<input checked="" type="radio"/> Rounded Lambda
	<input type="radio"/> Automatically Detect	<input type="radio"/> Optimal Lambda
		<input type="radio"/> Lambda & Threshold (Shift)

4. Click **Model Options**.

ARIMA Model Options ✕

Automatic Model Selection

Specify Model

Stepwise Procedure

Extended Model Search. Time limit seconds.

Model Selection Criterion

AICc - Akaike information criterion with small sample size correction

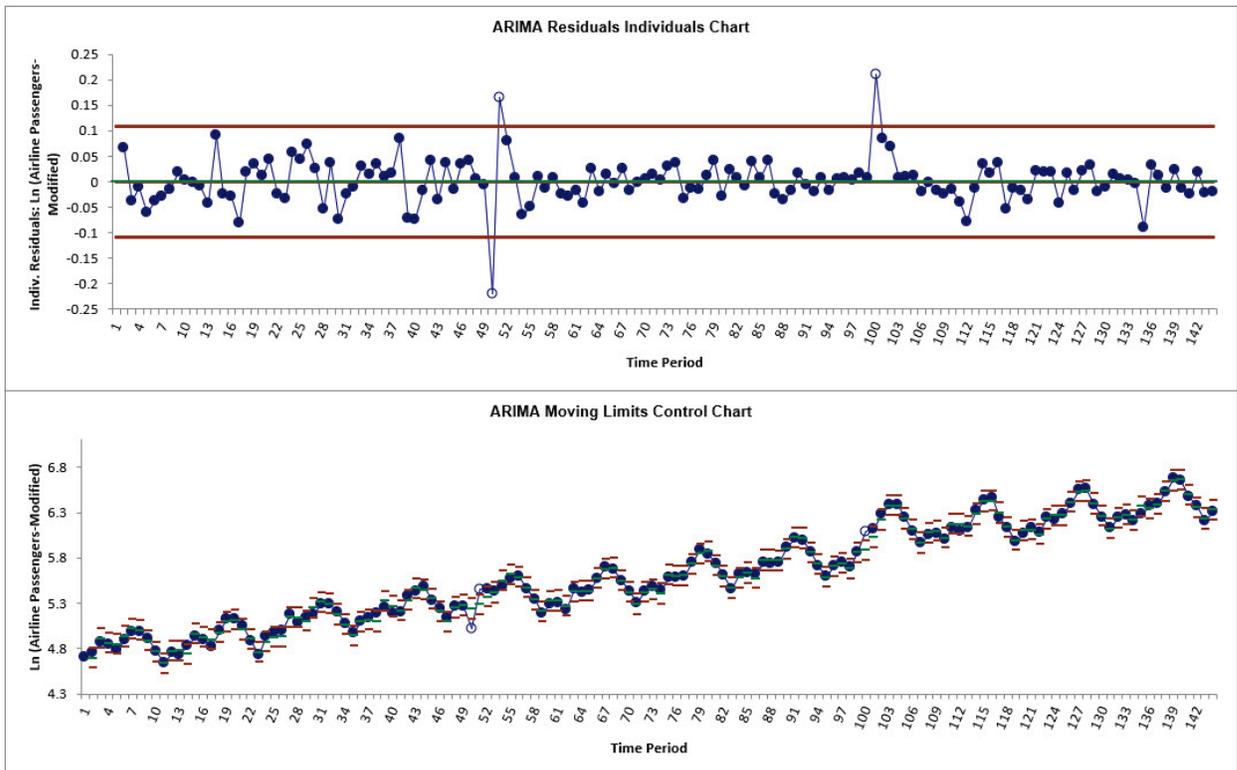
AIC - Akaike information criterion

BIC - Bayesian information criterion

Specify Nonseasonal Differencing (d) ▼

Specify Seasonal Differencing (D) ▼

5. We will use the default **Automatic Model Selection** with **AICc** as the **Model Selection Criterion**. Click **OK** to return to the ARIMA Control Chart dialog. Click **OK**. The ARIMA (MSD) control charts are produced:



We can clearly see the out-of-control data points at 50, 51 and 100 on the Residuals Individuals chart. This is similar to what we observed [previously](#) with regular ARIMA Control Charts.

6. Scroll down to view the ARIMA MSD Model header:

ARIMA Model (Multiple Seasonal Decomposition): Ln (Airline Passengers-Modified) - Model Automatically Selected
Model Periods: All observations are used in the ARIMA model estimation. No withhold periods available for out-of-sample forecast accuracy evaluation.

7. The ARIMA Model Summary is given as:

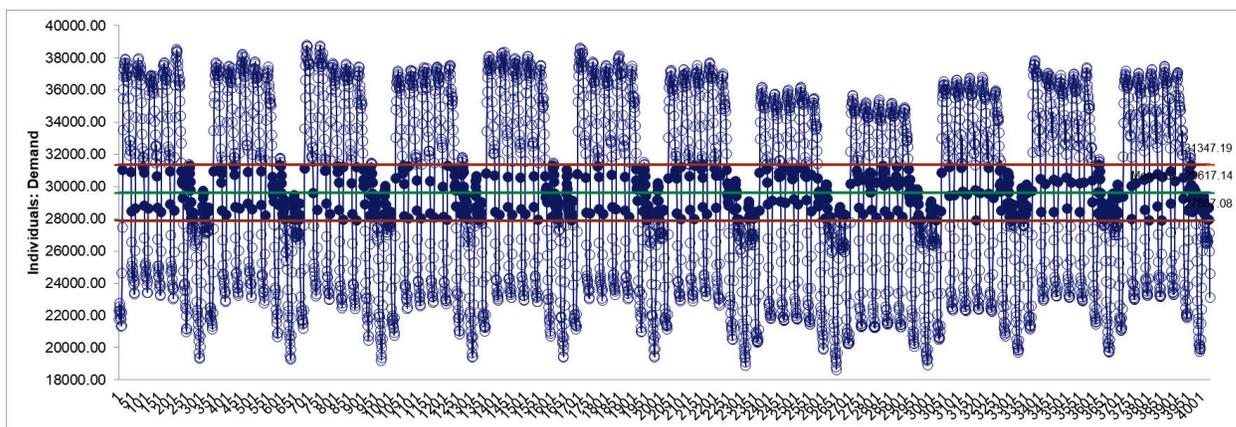
ARIMA Model Summary	
AR Order (p)	0
I Order (d)	1
MA Order (q)	1
SAR Order (P)	0
SI Order (D)	0
SMA Order (Q)	0
Seasonal Frequency	12 Decomposition
Include Constant	1
No. of Predictors	0
Model Selection Criterion	AICc
Box-Cox Transformation	N/A
Lambda	
Threshold	

This is a summary of the model information for the deseasonalized data: ARIMA (0,1,1) with a constant. Seasonal Frequency = 12 using Decomposition and Model Selection Criterion = "AICc". There are no seasonal terms in the model. The Box-Cox Transformation is "N/A".

8. We will not review the Parameter Estimates, Model Statistics and Forecast Accuracy as they are close to the ARIMA MSD values given [earlier](#), although note that slight differences are due to the introduction of an outlier and a shift, as well now we are using all of the data, i.e., there are no withhold periods. Earlier we used a Box-Cox Transformation with Lambda=0 and here we are using Ln of the data.

Half-Hourly Multiple Seasonal Electricity Demand – Taylor

9. Open **Half-Hourly Multiple Seasonal Electricity Demand - Taylor.xlsx (Sheet 1 tab)**. This is half-hourly electricity demand (MW) in England and Wales from Monday, June 5, 2000 to Sunday, August 27, 2000 (taylor, R forecast). This data has multiple seasonality with frequency = 48 (observations per day) and 336 (observations per week), with a total of 4032 observations. See the [Run Chart](#), [ACF/PACF Plots](#), [Spectral Density Plot](#) and [Seasonal Trend Decomposition Plots](#) for this data.
10. We will first construct a classical Individuals Control Chart on the raw data. Click **SigmaXL > Control Charts > Individuals**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
11. Select *Demand*, click **Numeric Data Variable (Y) >>**. Click **OK**. An Individuals Control Chart is produced:



With the high frequency seasonality, this control chart is meaningless.

12. Click **SigmaXL > Time Series Forecasting > ARIMA Control Chart > Multiple Seasonal Decomposition Control Chart**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
13. Select *Demand*, click **Numeric Time Series Data (Y) >>**. Uncheck **Display ACF/PACF/LB Plots** and **Display Residual Plots**. Check **Seasonal Frequency** with **Specify = 48 336**. Leave **Specify Model Periods** and **Box-Cox Transformation** unchecked.

ARIMA MSD Control Chart

Obs. No.

Numeric Time Series Data (Y) >> Demand

Optional Time Axis Labels >>

<< Remove

OK >>

Cancel

Help

Specify Model Periods

Start Model/Control Limit Calculations at Period 1

Withhold Periods 0

End Model/Control Limit Calculations at Period

Display ACF/PACF/LB Plots

Display Residual Plots

Box-Cox Transformation

Rounded Lambda

Optimal Lambda

Lambda & Threshold (Shift)

Model Options

Seasonal Frequency

Specify 48 336

Select 4-Quarterly

Automatically Detect

14. Click **Model Options**.

ARIMA Model Options

Automatic Model Selection

Specify Model

Stepwise Procedure

Extended Model Search. Time limit 300 seconds.

Model Selection Criterion

AICc - Akaike information criterion with small sample size correction

AIC - Akaike information criterion

BIC - Bayesian information criterion

Specify Nonseasonal Differencing (d) 0

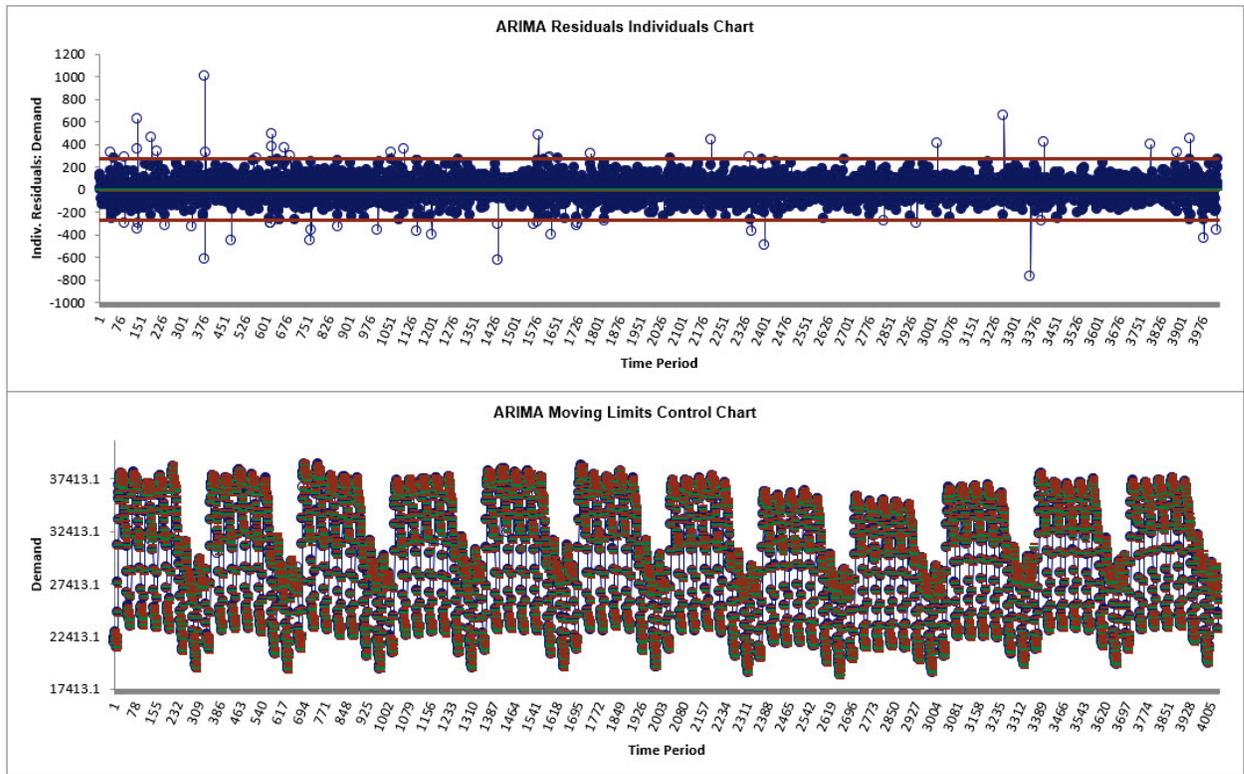
Specify Seasonal Differencing (D) 0

OK >>

Cancel

Help

15. We will use the default **Automatic Model Selection** with **AICc** as the **Model Selection Criterion**. Click **OK** to return to the ARIMA MSD Control Chart dialog. Click **OK**. The ARIMA MSD control charts are produced:



Now we have a chart that can be used to identify assignable causes. The number of out-of-control signals have been dramatically reduced. This is similar to what we observed [previously](#) with Exponential Smoothing MSD Control Charts.

16. Scroll down to view the ARIMA MSD Model header:

ARIMA Model (Multiple Seasonal Decomposition): Demand - Model Automatically Selected
Model Periods: All observations are used in the ARIMA model estimation. No withhold periods available for out-of-sample forecast accuracy evaluation.

17. The ARIMA Model Summary is given as:

ARIMA Model Summary	
AR Order (p)	2
I Order (d)	1
MA Order (q)	0
SAR Order (P)	0
SI Order (D)	0
SMA Order (Q)	0
Seasonal Frequency	48, 336 Decomposition
Include Constant	0
No. of Predictors	0
Model Selection Criterion	AICc
Box-Cox Transformation	N/A
Lambda	
Threshold	

This is a summary of the model information for the deseasonalized data: ARIMA (2,1,0) without a constant. Seasonal Frequency = 48, 336 using Decomposition and Model Selection Criterion = "AICc". There are no seasonal terms in the model. This is the same model used [previously](#) in ARIMA MSD Forecast, but that used a withhold sample of 96.

18. We will not review the Parameter Estimates, Model Statistics and Forecast Accuracy as they are close to the ARIMA MSD values given [earlier](#), although note that now we are using all of the data, i.e., there are no withhold periods.

Utilities – Difference Data

Nonseasonal and Seasonal Differencing are used to make a process stationary. This is done automatically in ARIMA, but this utility makes it easy to do manually. This was demonstrated previously in the section [Run Chart](#).

Chemical Process Concentration - Series A

1. Open **Chemical Process Concentration – Series A.xlsx (Sheet 1 tab)**. This is the Series A data from Box and Jenkins, a set of 197 concentration values from a chemical process taken at two-hour intervals.
2. Click the **Sheet 1** tab. Click **SigmaXL > Time Series Forecasting > Utilities > Difference Data**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
3. Select *Concentration*, click **Numeric Time Series Data (Y) >>**. Enter 1 for **Nonseasonal Differencing (d)**.

4. Click **OK**. A new sheet is created with the order $d = 1$ differenced data ($Y_2 - Y_1, Y_3 - Y_2, \dots$).

Concentration	Concentration: d = 1, D = 0
17	
16.6	-0.4
16.3	-0.3
16.1	-0.2
17.1	1
16.9	-0.2
16.8	-0.1
17.4	0.6
17.1	-0.3
17	-0.1
16.7	-0.3
17.4	0.7

Monthly Airline Passengers - Series G

- Open **Monthly Airline Passengers - Series G.xlsx** (Sheet 1 tab). This is the Series G data from Box and Jenkins, monthly total international airline passengers for January 1949 to December 1960 and is one of the most popular datasets used in introductory time series forecasting.
- Click the **Sheet 1** tab. Click **SigmaXL > Time Series Forecasting > Utilities > Difference Data**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
- Select *Ln(Airline Passengers)*, click **Numeric Time Series Data (Y) >>**. Enter 1 for **Nonseasonal Differencing (d)**; enter 1 for **Seasonal Differencing (D)**; **Seasonal Frequency** is specified as 12.

Difference Data

Obs. No.
Date
Monthly Airline Passenger

Numeric Time Series Data (Y) >> Ln (Airline Passengers)

<< Remove

Nonseasonal Differencing (d): 1

Seasonal Differencing (D): 1 Seasonal Frequency: 12

OK >>
Cancel
Help

- Click **OK**. A new sheet is created with the seasonal order $D = 1$ differenced data ($Y_{13} - Y_1, Y_{14} - Y_2, \dots$) and nonseasonal order $d=1$ ($D_2 - D_1, D_3 - D_2, \dots$).

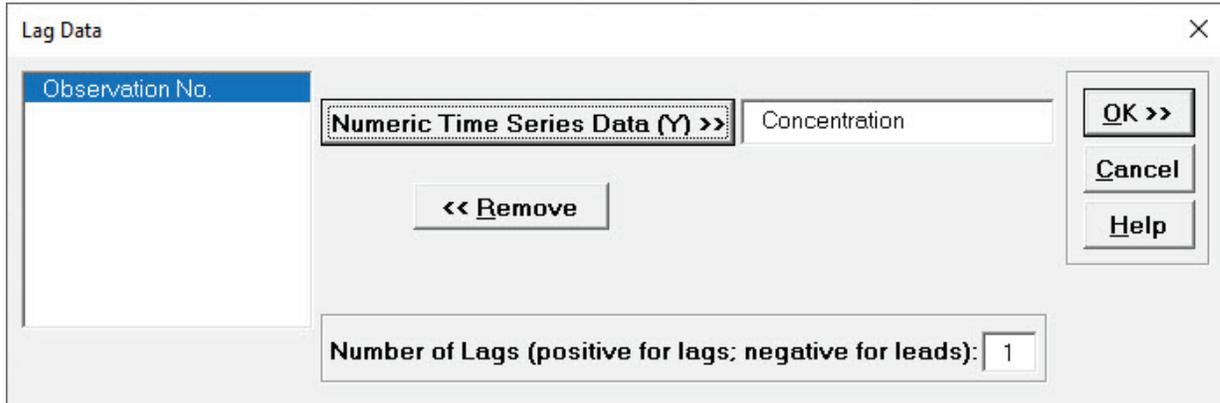
Ln (Airline Passengers)	Ln (Airline Passengers): d = 1, D = 1
4.718498871	
4.770684624	
4.882801923	
4.859812404	
4.795790546	
4.905274778	
4.997212274	
4.997212274	
4.912654886	
4.779123493	
4.644390899	
4.770684624	
4.744932128	
4.836281907	0.039164025
4.94875989	0.000360685
4.905274778	-0.020495594
4.828313737	-0.012939182

Utilities – Lag Data

The Lag Data utility is a useful utility to create Lag (or Lead) columns of data. This was demonstrated previously in the section [Autocorrelation \(ACF/PACF\) Plots](#).

Chemical Process Concentration - Series A

1. Open **Chemical Process Concentration – Series A.xlsx (Sheet 1 tab)**. This is the Series A data from Box and Jenkins, a set of 197 concentration values from a chemical process taken at two-hour intervals.
2. Click **SigmaXL > Time Series Forecasting > Utilities > Lag Data**. Ensure that the entire data table is selected. If not, check Use Entire Data Table. Click Next.
3. Select *Concentration*, click **Numeric Time Series Data (Y) >>**. Use the default **Number of Lags = 1**.



4. Click **OK**. A new sheet is created with the Lag 1 data.

Concentration	Concentration: Lag = 1
17	
16.6	17
16.3	16.6
16.1	16.3
17.1	16.1
16.9	17.1

5. Click **Recall SigmaXL Dialog** menu or press **F3** to recall last dialog. Enter **Lag = 2**. Click **OK**. Repeat for **Lag = 3**. The Lag columns are appended as shown:

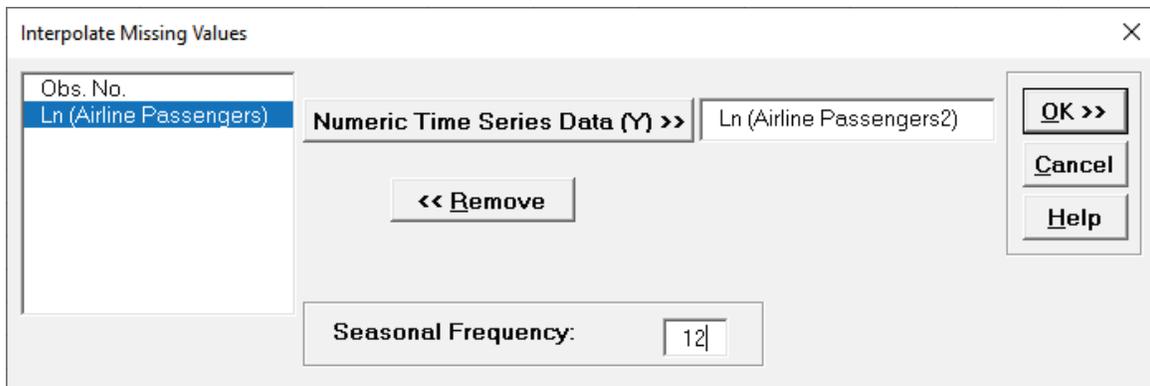
Concentration	Concentration: Lag = 1	Concentration: Lag = 2	Concentration: Lag = 3
17			
16.6	17		
16.3	16.6	17	
16.1	16.3	16.6	17
17.1	16.1	16.3	16.6
16.9	17.1	16.1	16.3
16.8	16.9	17.1	16.1

Utilities – Interpolate Missing Values

The Interpolate Missing Values utility was demonstrated [previously](#). See the Appendix: [Seasonally Adjusted Linear Interpolation of Missing Values](#) for further details. It is used automatically in Exponential Smoothing when there are missing values. While there is robustness against some missing values, if the number of missing values is large, then model estimation and forecast accuracy will be degraded

Monthly Airline Passengers – Missing Values

1. Open **Monthly Airline Passengers – Missing Values.xlsx (Sheet 1 tab)**. Ln(Airline Passengers2) have missing values at observations 50 and 100.
2. Click the **Sheet 1** tab. Click **SigmaXL > Time Series Forecasting > Utilities > Interpolate Missing Values**. Ensure that the entire data table is selected. If not, check **Use Entire Data Table**. Click **Next**.
3. Select Ln(Airline Passengers2), click **Numeric Time Series Data (Y) >>**. **Seasonal Frequency** is specified as 12.



4. Click **OK**. A new sheet is created with the seasonally adjusted linear interpolated values highlighted in yellow.

Ln (Airline Passengers2)	Y: Interpolated (12)
4.718498871	4.718498871
4.770684624	4.770684624
4.882801923	4.882801923
4.859812404	4.859812404
	5.254047791
	5.846836235

Y50 is estimated as 5.254 (versus original value of 5.278). Y100 is estimated as 5.847 (versus original value of 5.852).

SigmaXL Appendix

SigmaXL[®] Version 9.0 Feature List Summary

Compatible with Excel 2019 for Windows and Mac

Menu Layout: Classical or DMAIC

Recall Last Dialog

Data Manipulation:

- Subset by Category, Number, Date or Random
- Stack Subgroups Across Rows
- Stack and Unstack Columns
- Standardize Data
- Random Number Generator
 - Normal
 - Uniform (Continuous & Integer)
 - Lognormal
 - Weibull
 - Triangular
- Box-Cox Transformation

Templates & Calculators:

- DMAIC & DFSS Templates
 - Team/Project Charter
 - SIPOC Diagram
 - Flowchart Toolbar
 - Data Measurement Plan
 - Cause & Effect (Fishbone) Diagram and Quick Template
 - Cause & Effect (XY) Matrix with Pareto
 - Failure Mode & Effects Analysis (FMEA) with RPN Sort
 - Quality Function Deployment (QFD)
 - Pugh Concept Selection Matrix
 - Control Plan
- Lean Templates
 - Takt Time Calculator
 - Value Analysis/Process Load Balance
 - Value Stream Mapping
- Graphical Templates
 - Pareto Chart, Histogram, Run Chart
- Statistical Templates
 - Sample Size – Discrete and Continuous
 - Minimum Sample Size for Robust t-Tests and ANOVA
 - 1 Sample Z-test & Confidence Interval for Mean
 - 1 Sample t-Test & Confidence Interval for Mean
 - 2 Sample t-Test & Confidence Interval (Compare 2 Means)
 - 1 Sample Equivalence Test for Mean
 - 2 Sample Equivalence Test (Compare 2 Means)
 - 1 Sample Chi-Square Test and CI for Standard Deviation
 - 2 Sample F-Test and CI (Compare 2 Standard Deviations)
 - 1 Proportion Test and Confidence Interval
 - 2 Proportions Test and Confidence Interval (with Fisher Exact P-Value)
 - 2 Proportions Equivalence Test
 - 1 Poisson Rate Test and Confidence Interval
 - 2 Poisson Rates Test and Confidence Interval
 - 2 Poisson Rates Equivalence Test
 - One-Way Chi-Square Goodness-of-Fit Test (with Exact and Monte Carlo P-Value)
- Probability Distribution Calculators
 - Normal, Lognormal, Exponential, Weibull
 - Binomial, Poisson, Hypergeometric
- Measurement Systems Analysis (MSA) Templates
 - Type 1 Gage Study
 - Gage Bias and Linearity Study
 - Gage R&R Study – with Multi-Vari Analysis
 - Attribute Gage R&R (Attribute Agreement Analysis)
- Process Sigma Level – Discrete and Continuous
- Process Capability & Confidence Intervals
- Tolerance Interval Calculator (Normal Exact)
- DOE Templates
 - 2 to 5 Factors
 - 2-Level Full and Fractional-Factorial designs
 - Main Effects & Interaction Plots
- Taguchi DOE Templates
 - L4, L8, L9, L12, L16, L18, L27
 - Signal-to-Noise Ratios: Nominal is Best, Nominal is Best (Variance Only), Nominal is Best (Mean Square Deviation with Target), Larger is Better, Smaller is Better
 - Pareto of Deltas (Effects) and ANOVA SS (Sum-of-Squares) % Contribution (for Main Effects and Two-Way Interactions)
 - Main Effects & Interaction Plots
- Control Chart Templates
 - Individuals, C (Count)
 - Rare Events: T, G and Probability Based G**
 - Exponentially Weighted Moving Average (EWMA)**
 - Tabular Cumulative Sum (CUSUM)**
 - Trend/Tool Wear**
 - Average Run Length (ARL) Calculators: Shewhart with Tests for Special Causes, Attribute C & P, EWMA & CUSUM**

SigmaXL[®] Version 9.0 Feature List Summary

Graphical Tools:

- Basic and Advanced (Multiple) Pareto Charts
- EZ-Pivot/Pivot Charts: Easily create Pivot Tables and Charts
- Basic Histogram
- Multiple Histograms and Descriptive Statistics (includes Confidence Interval for Mean and StDev., and Anderson-Darling Normality Test)
- Multiple Histograms and Process Capability (Pp, Ppk, Cpm, ppm, %)
- Multiple Boxplots, Dotplots
- Run Charts (with Nonparametric Runs Test allowing you to test for Clustering, Mixtures, Lack of Randomness, Trends and Oscillation)
- Overlay Run Chart
- Multiple Normal Probability Plots (with 95% confidence intervals to ease interpretation of normality/non-normality)
- Multi-Vari Charts
- Scatter Plots (with linear regression and optional 95% confidence intervals and prediction intervals)
- Scatter Plot Matrix
- Analysis of Means (ANOM) Charts
 - All charts support balanced & unbalanced data
 - Normal, Binomial Proportions and Poisson Rates: One-Way & Two-Way with Main Effects and Slice Charts
 - Slice Charts are a modified ANOM chart developed by Dr. Peter Wludyka that enables one to easily interpret the effects in the presence of an interaction. In collaboration, Peter Wludyka and John Noguera of SigmaXL extended the Slice Charts to Binomial and Poisson. Yellow highlight automatically recommends Main Effects (if interaction is not significant) or Slice Chart (if interaction is significant).
 - Nonparametric Transformed Ranks
 - Variances & Levene Robust Variances

Measurement Systems Analysis:

- Create Gage R&R (Crossed) Worksheet:
 - Generate worksheet with user specified number of parts, operators, replicates
- Analyze Gage R&R (Crossed)
 - ANOVA, %Total, %Tolerance (with upper and/or lower specifications), %Process, Variance Components, Number of Distinct Categories
 - Gage R&R Multi-Vari and X-bar R Charts
 - Confidence Intervals for %Total, %Tolerance, %Process and Standard Deviations
 - Handles unbalanced data
- Attribute MSA (Binary, Ordinal, Nominal)
 - Any number of samples, appraisers and replicates
 - Within Appraiser Agreement, Each Appraiser vs Standard Agreement, Each Appraiser vs Standard Disagreement, Between Appraiser Agreement, All Appraisers vs Standard Agreement; Fleiss' Kappa; Kendall's Coefficients (Ordinal Concordance and Correlation)
 - Confidence Interval Graphs for Percent Agreement, Fleiss' Kappa and Kendall's Coefficients
 - Kappa and Kendall color highlight to aid interpretation (Green/Yellow/Red)
 - Effectiveness Report (treats each appraisal trial as an opportunity) and Misclassification Summary

Process Capability:

- Multiple Histograms and Process Capability
- Capability Combination Report for Individuals/Subgroups:
 - Histogram, Normal Probability Plot and Normality Test
 - Capability Report (Cp, Cpk, Pp, Ppk, Cpm, ppm, %)
 - Control Charts
- Capability Combination Report for Nonnormal Data (Individuals)
 - Box-Cox Transformation (includes an automatic threshold option so that data with negative values can be transformed)
 - Johnson Transformation
 - Distributions supported: Half-Normal, Lognormal (2 & 3 parameter), Exponential (1 & 2), Weibull (2 & 3), Beta (2 & 4), Gamma (2 & 3), Logistic, Loglogistic (2 & 3), Largest Extreme Value, Smallest Extreme Value
 - Automatic Best Fit based on AD p-value
 - Nonnormal Process Capability Indices: Z-Score (Cp, Cpk, Pp, Ppk) and Percentile (ISO) Method (Pp, Ppk)
- Distribution Fitting Report
 - All valid distributions and transformations reported with histograms, curve fit and probability plots
 - Sorted by AD p-value

SigmaXL[®] Version 9.0 Feature List Summary

Statistical Tools:

- P-values turn red when results are significant: (p-value < alpha)
- Descriptive Statistics including Anderson-Darling Normality test, Skewness and Kurtosis with p-values
- Descriptive Statistics options:
 - Percentile Report and Percentile Ranges
 - Percentile Confidence and Tolerance Intervals
 - Additional Descriptive Statistics and Normality Tests (Shapiro-Wilk and Doornik-Hansen Normality)
 - Outlier (Boxplot and Grubbs) and Randomness (Nonparametric Runs) Tests
- 1 Sample t-test and confidence intervals
- Paired t-test, 2 Sample t-test
- 2 Sample comparison tests:
 - Reports AD Normality, F-test and Levene's for variance, t-test assuming equal and unequal variance, Mann-Whitney test for medians.
 - Recommended tests are highlighted based on sample size, normality, and variance
- One-Way ANOVA and Means Matrix
 - Multiple Comparison of Means Probability Methods (Post-Hoc): Fisher, Tukey, Dunnett with Control
- Automatic Assumptions Check for One Sample, Two-Sample, Paired T-tests and One-Way ANOVA
 - Text report with color highlight: Green (OK), Yellow (Warning) and Red (Serious Violation)
 - Test each sample for Normality. If not, check minimum sample size for robustness of test
 - Check each sample for Outliers: Potential (Tukey's Boxplot 1.5*IQR); Likely (2.2*IQR); Extreme (3*IQR)
 - Randomness (Nonparametric Runs Test)
 - Equal Variance (for 2 or more samples)
- Two-Way ANOVA (Balanced and Unbalanced)
- Equal Variance Tests (Bartlett, Levene and Welch's ANOVA)
 - Multiple Comparison of Variances Probability Methods (Post-Hoc): F-Test (with Bonferroni Correction), Levene, Tukey ADM (Absolute Deviations from Median)
 - Welch Multiple Comparison of Means Probability Methods (Post-Hoc): Welch Pairwise, Games-Howell
- Correlation Matrix (Pearson and Spearman's Rank Correlation)
 - Automatic normality check for correlation utilizing the powerful Doornik-Hansen bivariate normality test
 - Yellow highlight to recommend significant Pearson or Spearman correlation - Pearson is highlighted if the data are bivariate normal, otherwise Spearman is highlighted
- Multiple Linear Regression:
 - Accepts continuous and/or categorical (discrete) predictors
 - Interactive Predicted Response Calculator with 95% Confidence Interval and 95% Prediction Interval
 - Residual Plots: histogram, normal probability plot, residuals vs. time, residuals vs. predicted and residuals vs. X factors
 - Residual types include Regular, Standardized, Studentized (Deleted t) and Cook's Distance (Influence), Leverage and DFITS
 - Highlight of significant outliers in residuals
 - Durbin-Watson Test for Autocorrelation in Residuals with p-value
 - ANOVA report for categorical predictors
 - Pure Error and Lack-of-Fit report
 - Collinearity Variance Inflation Factor (VIF) and Tolerance report
 - Fit Intercept is optional
- Binary and Ordinal Logistic Regression
 - Powerful and user-friendly logistic regression.
 - Report includes a calculator to predict the response event probability for a given set of input X values.
 - Categorical (discrete) predictors can be included in the model in addition to continuous predictors.
 - Model summary and goodness of fit tests include Likelihood Ratio Chi-Square, Pseudo R-Square, Pearson Residuals Chi-Square, Deviance Residuals Chi-Square, Observed and Predicted Outcomes – Percent Correctly Predicted.
 - Stored data includes Event Probabilities, Predicted Outcome, Observed-Predicted, Pearson Residuals, Standardized Pearson Residuals, and Deviance Residuals.
- Chi-Square Test (Stacked Column data and Two-Way Table data)
 - With Fisher Exact (utilizing permutations and fast network algorithms) and Monte Carlo P-Values
 - Options: Advanced Tests and Measures of Association for Nominal & Ordinal Categories
- Nonparametric Tests:
 - 1 Sample Sign and 1 Sample Wilcoxon
 - 2 Sample Mann-Whitney
 - Kruskal-Wallis and Mood's Median Test (with graph of Group Medians and 95% Median Confidence Intervals)
 - Runs Test
 - With Exact and Monte Carlo P-Value
- Power and Sample Size Calculators for:
 - 1 and 2 Sample t-Test
 - One-Way ANOVA
 - 1 Proportion Test, 2 Proportions Test
 - The Power and Sample Size Calculators allow you to solve for Power (1 – Beta), Sample Size, or Difference (specify two, solve for the third).
- Power and Sample Size Chart. Quickly create a graph showing the relationship between Power, Sample Size and Difference.

SigmaXL® Version 9.0 Feature List Summary

Design of Experiments:

- Generate 2-Level Factorial and Plackett-Burman Screening Designs
 - User-friendly dialog box
 - 2 to 19 Factors; 4,8,12,16,20 Runs
 - Unique "view power analysis as you design"
 - Randomization, Replication, Blocking and Center Points
- Basic DOE Templates
 - 2 to 5 Factors, 2-Level Full and Fractional-Factorial designs
 - Automatic update to Pareto of Coefficients
 - Easy to use, ideal for training
- Main Effects & Interaction Plots
- Analyze 2-Level Factorial and Plackett-Burman Screening Designs
 - Used in conjunction with Recall Last Dialog, it is very easy to iteratively remove terms from the model
 - Interactive Predicted Response Calculator with 95% Confidence Interval and 95% Prediction Interval.
 - ANOVA report for Blocks, Pure Error, Lack-of-Fit and Curvature
 - Collinearity Variance Inflation Factor (VIF) and Tolerance report
 - Residual plots: histogram, normal probability plot, residuals vs. time, residuals vs. predicted and residuals vs. X factors
 - Highlight of significant outliers in residuals
 - Durbin-Watson Test for Autocorrelation
- Contour & 3D Surface Plots
- Response Surface Designs
 - 2 to 5 Factors
 - Central Composite and Box-Behnken Designs
 - Easy to use design selection sorted by number of runs

Control Charts:

- Control Chart Selection Tool
- Individuals, Individuals & Moving Range
- X-Bar & R, X-Bar & S
- I-MR-R, I-MR-S (Between/Within)
- P, NP, C, U
- P' and U' (Laney) to handle overdispersion
- Control charts include a report on tests for special causes. Special causes are also labeled on the control chart data point. Set defaults to apply any or all of Tests 1-8.
- Tests for Special Causes: Support for Varying Subgroup Sizes (Moving Limits), Historical Groups and MR/Range/StDev Charts (Tests 1-4)**
- Process Capability report (Pp, Ppk, Cp, Cpk) is available for I, I-MR, X-Bar & R, X-bar & S charts.
- Add data to existing charts for operator ease of use!
- Scroll through charts with user defined window size
- Advanced Control Limit options: Subgroup Start and End; Historical Groups (e.g. split control limits to demonstrate before and after improvement)
- Exclude data points for control limit calculation
- Add comment to data point for assignable cause
- $\pm 1, 2$ Sigma Zone Lines
- Control charts for Nonnormal data (Individuals)
 - Box-Cox and Johnson Transformations
 - 16 Nonnormal distributions supported (see Process Capability)
 - Individuals chart of original data with percentile based control limits
 - Individuals/Moving Range chart for normalized data with optional tests for special causes
- Control chart templates:
 - Individuals, C (Count)
 - Rare Events: T, G and Probability Based G**
 - Time Weighted: Exponentially Weighted Moving Average (EWMA) and Tabular Cumulative Sum (CUSUM)**
 - Trend/Tool Wear**
 - Average Run Length (ARL) Calculators: Shewhart with Tests for Special Causes, Attribute C & P, EWMA & CUSUM**

Reliability/Weibull Analysis:

- Weibull Analysis
 - Complete and Right Censored data
 - Least Squares and Maximum Likelihood
 - Output includes percentiles with confidence intervals, survival probabilities, and Weibull probability plot.

SigmaXL[®] Version 9.0 Feature List Summary

Time Series Forecasting and Control Charts for Autocorrelated Data:**

- Run Chart
- Autocorrelation Function (ACF)/Partial Autocorrelation (PACF) Plots
- Cross Correlation (CCF) Plots with Pre-Whiten Data option
- Seasonal Trend Decomposition Plots
- Spectral Density Plot with Detection of Seasonal Frequency
- Exponential Smoothing:
 - Forecast with Prediction Intervals
 - Exponential Smoothing models use Rob Hyndman's taxonomy:
 - Additive/Multiplicative Error
 - Additive/Additive Damped Trend
 - Additive/Multiplicative Seasonal
 - This includes all of the classical exponential smoothing models such as Simple/Single/EWMA, Double and Holt-Winters
 - Multiple Seasonal Decomposition (MSD) option: useful for high frequency and/or multiple frequency data, such as Monthly with frequency = 12, Daily with frequency = 7 and Hourly with frequency = 24
- Exponential Smoothing Residuals Control Chart for autocorrelated data:
 - Individuals and Moving Limits (with One-Step Ahead Forecast) Charts
 - Add Data, Show Last 30 Data Points, Enable Scroll options
 - MSD option
- Autoregressive Integrated Moving Average (ARIMA):
 - Forecast with Prediction Intervals
 - ARIMA Forecast with Predictors (Continuous and/or Categorical)
 - MSD option
- ARIMA Residuals Control Chart for autocorrelated data:
 - Individuals and Moving Limits (with One-Step Ahead Forecast) Charts
 - ARIMA Control Chart with Predictors
 - Add Data, Show Last 30 Data Points, Enable Scroll options
 - MSD option
- Utilities:
 - Difference Data
 - Lag Data
 - Interpolate Missing Values (seasonally adjusted linear interpolation)
- Time Series Forecasting Model Features:
 - ARIMA and Exponential Smoothing models are fully automatic or user specified
 - Utilizes modern State Space and Kalman Filter models for accurate parameter estimation
 - ARIMA estimates missing values with Kalman Filter; Exponential Smoothing uses seasonally adjusted linear interpolation
 - Automatic Box-Cox Transformation
 - Automatic seasonal frequency detection
- Model diagnostics:
 - ACF/PACF Plots
 - Ljung-Box p-values
 - Log-Likelihood, AIC, AICc, BIC, Residual StDev
 - Residual plots (histogram, normal probability, residual versus fits, residuals versus order)
- Forecast Accuracy:
 - In-Sample (Estimation) one-step-ahead forecast errors (RMSE, MAE, MASE, MAPE)
 - Out-of-Sample (Withhold) one-step-ahead forecast errors
 - Out-of-Sample (Withhold) multi-step-ahead forecast errors
 - Evaluated using the benchmark standard M4 forecast competition data, a total of 100,000 data sets with Yearly, Quarterly, Monthly, Weekly, Daily and Hourly data. Using a hybrid average of automatic Exponential Smoothing and ARIMA, SigmaXL (unofficially) ranked 10th out of 60 in the Overall Weighted Average forecast accuracy score, ahead of three well known commercial forecast software packages.

Descriptive Statistics Options and Tolerance Interval Calculator (Normal Exact)

Shapiro-Wilk (SW) and Kolmogorov-Smirnov-Lilliefors (KSL) Normality Test

Shapiro-Wilk [1] is a popular and powerful alternative to the Anderson Darling normality test. It was originally restricted to sample sizes less than 50. Royston [2] extended the sample size to 2000. Rahman and Govidarajulu [3] extended the sample size further to 5000 and this has been implemented in SigmaXL.

For sample sizes greater than 5000, Kolmogorov-Smirnov-Lilliefors (KSL) is used. Lilliefors [4] modified the Kolmogorov-Smirnov test for the case of unknown mean and standard deviation.

Doornik-Hansen (DH) Normality Test

The Doornik-Hansen Normality Test [5] is a univariate or multivariate variate omnibus test based on Skewness and Kurtosis. In descriptive statistics, the univariate version is used, which is based on the method of Shenton and Bowman [6]. The bivariate version of DH is used in the Correlation Matrix to test for bivariate normality.

This test is best suited for data with ties, i.e. “chunky” data. Anderson-Darling and Shapiro-Wilk are severely affected by ties in the data and will trigger a low P-Value even if the data are normal.

Confidence and Tolerance Intervals

A confidence interval is described by the percentile of the population and the confidence probability. Thus, the 95% confidence interval for the 80th percentile of the sample will consist of a lower bound and an upper bound, where the true population 80th percentile will lie 95% of the time.

A tolerance interval is the interval, described by a lower and upper bound, that a specified minimum proportion of the sample’s values will lie with a given confidence probability. Thus, as an example, the 95% tolerance interval for the range in which 80 percent of the sample will fall will consist of a lower bound and an upper bound.

Notation:

n	is the number of observations in the sample.
v	are the degrees of freedom in the sample ($n-1$).
\bar{X}	is the mean of the sample.

s	is the standard deviation of the sample.
p	is the quantile (CI), or the percentage of the population in the interval (TI).
Z_p	is the p th quantile of the standard normal distribution.
α	is the alpha probability - typically 0.05.

Tolerance Interval Calculator (Normal Exact)

When the sample can be assumed to be derived from a population that has a normal distribution, then parametric assumptions can be used to ascertain the Tolerance Interval. For each of the following, the lower L and upper U bounds are given by:

$$L = \bar{X} - ks$$

$$U = \bar{X} + ks$$

The value of k is shown below:

Normal Tolerance Interval - one sided exact

$$k = \frac{t_{1-\alpha, v, \lambda}}{\sqrt{n}}$$

where $t_{1-\alpha, v, \lambda}$ is the $1 - \alpha$ percentile of the non-central t distribution with v degrees of freedom and non-centrality parameter $\lambda = \sqrt{n}Z_p$.

Normal Tolerance Interval - two sided exact

$$G(z, n, p, k) = 1 - F_v \left[\frac{v \chi_{p, 1, \lambda}^2}{k^2} \right]$$

where: $\chi_{p, 1, \lambda}^2$ is the p percentile of the non-central χ^2 distribution with 1 degree of freedom and non-centrality parameter $\lambda = z^2/n$, and where F_v is the cdf of the χ^2 distribution with v degrees of freedom. Then k is the root such that:

$$2 \int_0^{\infty} G(z, n, p, k) \phi(z) dz = 1 - \alpha$$

where $\phi(z)$ is the pdf of the standard normal distribution. SigmaXL uses a 40-point Gauss-Legendre quadrature to evaluate the integral. See Krishnamoorthy and Mathew [7]. The approximation method by Howe [8] is used to obtain initial values.

Percentile (Nonparametric) Confidence and Tolerance Intervals

This methodology is used when no *a priori* assumptions can be made regarding the distribution of the data. Let x be the sorted data. The exact lower L and upper U bounds are given by:

$$\begin{aligned} L &= x_r \\ U &= x_s \end{aligned}$$

The interpolated lower L and upper U bounds are given by:

$$\begin{aligned} L &= \lambda_l x_{r+1} + (1 - \lambda_l) x_r \\ U &= \lambda_u x_{s+1} + (1 - \lambda_u) x_s \end{aligned}$$

Confidence Interval – two sided exact

$$\begin{aligned} r &= B_{\alpha/2, n, p} \\ s &= B_{1-\alpha/2, n, p} + 1 \end{aligned}$$

where $B_{\alpha, n, p}$ is the α percentile of the binomial distribution with parameters n and p .

Confidence Interval – two sided interpolated

Let r and s be the exact order statistics for confidence interval. Then

$$\begin{aligned} \lambda_l &= \left(1 + \frac{r(1-p)(B_{n,p}(r) - \alpha/2)}{(n-r)p(\alpha/2 - B_{n,p}(r-1))} \right)^{-1} \\ \lambda_u &= \left(1 + \frac{s(1-p)(B_{n,p}(s) - (1-\alpha/2))}{(n-s)p(1-\alpha/2 - B_{n,p}(s-1))} \right)^{-1} \end{aligned}$$

where $B_{n,p}(k)$ is the cdf of the binomial distribution with parameters n and p . See Nyblom [9] which is an extension of the interpolation method of Hettmansperger and Sheather [10] for median using a double exponential distribution.

Tolerance Interval - two sided exact

$$\begin{aligned} k &= B_{1-\alpha, n, p} + 1 \\ r &= \lfloor (n - k + 1)/2 \rfloor \\ s &= n - r + 1 \end{aligned}$$

where $B_{1-\alpha, n, p}$ is the $1 - \alpha$ percentile of the binomial distribution with parameters n and p .

Tolerance Interval - two sided interpolated

Let k , r and s be the two sided exact order statistics for tolerance interval.

$$\lambda = \frac{1 - \alpha/2 - B_{n,p}(k-2)}{B_{n,p}(k-1) - B_{n,p}(k-2)}$$

$$v_1 = \lambda x_r + (1 - \lambda)x_{r+1}$$

$$v_2 = \lambda x_s + (1 - \lambda)x_{s-1}$$

If $x_s - v_1 < v_2 - x_r$, then $L = v_1$ and $U = x_s$; else $L = x_r$ and $U = v_2$. See Young and Mathew [11].

References for Descriptive Statistics Options and Tolerance Interval Calculator (Normal Exact)

- [1] “Shapiro-Wilk Test” in *Wikipedia: The Free Encyclopedia, Wikimedia Foundation Inc.*, https://en.wikipedia.org/wiki/Shapiro%E2%80%93Wilk_test.
- [2] Royston, Patrick (September 1992). “Approximating the Shapiro–Wilk W-test for non-normality”, *Statistics and Computing*. 2 (3): 117–119.
- [3] Rahman and Govidarajulu (1997). “A modification of the test of Shapiro and Wilk for normality”, *Journal of Applied Statistics*. 24 (2): 219–236
- [4] Lilliefors, H. (June 1967), "On the Kolmogorov–Smirnov test for normality with mean and variance unknown", *Journal of the American Statistical Association*, Vol. 62. pp. 399–402.
- [5] Doornik, J.A. and Hansen, H. (2008). “An Omnibus test for univariate and multivariate normality”, *Oxford Bulletin of Economics and Statistics*, 70, 927-939.
- [6] Shenton, L. R. and Bowman, K. O. (1977). “A bivariate model for the distribution of v_{b1} and b_2 ”, *Journal of the American Statistical Association*, Vol. 72, pp. 206–211.
- [7] Krishnamoorthy, K. and Mathew, T. (2009), *Statistical Tolerance Regions: Theory, Applications, and Computation*, Hoboken, NJ:Wiley.
- [8] Howe, W. G. (1969), “Two-Sided Tolerance Limits for Normal Populations - Some Improvements”, *Journal of the American Statistical Association* 64, pp. 610-620.
- [9] Nyblom, J. (1992). “Note on interpolated order statistics”, *Statistics and Probability Letters* 14, pp. 129-131.
- [10] Hettmansperger, T.P. and S.J. Sheather (1986). “Confidence Intervals Based on Interpolated Order Statistics”, *Statistics and Probability Letters* 4, pp. 75-79.
- [11] Young, D.S. and Mathew, T. (2014). “Improved nonparametric tolerance intervals based on interpolated and extrapolated order statistics”, *Journal of Nonparametric Statistics* 26 (3), pp. 415-432,

Statistical Details for Nonnormal Distributions and Transformations

Maximum Likelihood Estimation (MLE)

Maximum likelihood estimates of the parameters are calculated by maximizing the likelihood function with respect to the parameters. The likelihood function is simply the sum of the log of the probability density function (PDF) for each uncensored observation, and the log of the complement of the cumulative density function (CDF) for each right censored observation (in Reliability/Weibull Analysis). Initial estimates are derived using a branch and bound algorithm.

The maximum likelihood estimates are then calculated using the Newton-Raphson method. This is an iterative process that uses both the first and second derivatives to move to a point at which no further improvement in the likelihood is possible.

The standard errors of the parameter estimates are derived from the Hessian matrix. This matrix, which describes the curvature of a function, is the square matrix of second-order partial derivatives of the function.

For some data sets, the likelihood function for threshold models is unbounded, and the maximum likelihood methodology fails. In this context, a threshold is estimated using a bias correction method. This is an iterative methodology that evaluates the threshold based on the difference between the minimum value of the variate and the prediction for the minimum value, conditional on the current values of the parameters.

References for MLE and Distributions:

Greene, W.H. *Econometric Analysis* 4th Ed Prentice Hall, New Jersey.

Johnson, N. L., and Kotz, S. (1990). "Use of moments in deriving distributions and some characterizations", *Mathematical Scientist*, Vol. 15, pp. 42-52.

Johnson, N.L., Kotz, S. and Balakrishnan, N. (1994). *Continuous Univariate Distributions-Volume 1*, Second Ed., Wiley, New York.

Johnson, N.L., Kotz, S. and Balakrishnan, N. (1995). *Continuous Univariate Distributions-Volume 2*, Second Ed., Wiley, New York.

Nocedal, J. and Wright, S.J. (1999). *Numerical Optimization*, Springer-Verlag, New York.

Sleeper, A. (2006). *Six Sigma Distribution Modeling*, McGraw-Hill, New York.

Distributions

Beta Distribution

PDF

$$\frac{1}{B(\alpha, \beta)} x^{\alpha-1} (1-x)^{\beta-1}$$

CDF

$$\int_{-\infty}^x \frac{1}{B(\alpha, \beta)} t^{\alpha-1} (1-t)^{\beta-1} dt$$

where B is the Beta function:

$$B(\alpha, \beta) = \frac{\Gamma(\alpha)\Gamma(\beta)}{\Gamma(\alpha + \beta)}$$

Range $0 \leq x \leq 1$.

Shape1 parameter $\alpha > 0$.

Shape2 parameter $\beta > 0$.

$\Gamma(\beta)$ is the Gamma function and is described below under Gamma distribution.

Beta Distribution with Lower/Upper Threshold

PDF

$$\frac{1}{B(\alpha, \beta)} \left(\frac{x - \theta_1}{\theta_2 - \theta_1} \right)^{\alpha-1} \left(\frac{\theta_2 - x}{\theta_2 - \theta_1} \right)^{\beta-1}$$

CDF

$$\int_{-\infty}^{\frac{x - \theta_1}{\theta_2 - \theta_1}} \frac{1}{B(\alpha, \beta)} t^{\alpha-1} (1 - t)^{\beta-1} dt$$

where B is the Beta function.

Range $0 \leq x \leq 1$.

Shape1 parameter $\alpha > 0$.

Shape2 parameter $\beta > 0$.

Lower threshold θ_1

Upper threshold θ_2

Notes: Estimation of the 4 parameter Beta distribution is undertaken in two parts. In the first part, initial parameter estimates are derived using the method of moments. The threshold parameters are then held using these values, and the shape parameters are then estimated using maximum likelihood.

Reference for Beta Distribution with Lower/Upper Threshold – Parameter Estimation:

Wang, J.Z. (2005). "A note on Estimation in the Four Parameter Beta Distribution", *Comm in Stats Simulation and computation*, Vol. 34 pp. 495- 501.

Box-Cox Distribution with Threshold

PDF

$$\frac{(x - \theta)^{\lambda-1}}{\sigma\sqrt{2\pi}} e^{-(((x-\theta)^\lambda - 1)/\lambda - \mu)^2 / 2\sigma^2}$$

CDF

$$\int_{-\infty}^{[(x-\theta)^\lambda - 1]/\lambda} \frac{1}{\sigma\sqrt{2\pi}} e^{-(t-\mu)^2 / 2\sigma^2} dt$$

Range $0 < x - \theta < \infty$.

Location parameter, μ , the mean.

Scale parameter, $\sigma > 0$, the standard deviation.

Shape parameter λ .

Threshold parameter $\theta < \min(x)$.

Notes: The concentrated likelihood is used in the ML estimation. This implies that the location and scale parameters are not estimated freely, but are derived as the mean and standard deviation of the BoxCox transformed variate. The estimated parameters λ and θ are then used in the Box-Cox (Power) transformation. See **Transformations** below.

Exponential Distribution

PDF

$$\frac{e^{-x/\alpha}}{\alpha}$$

CDF

$$1 - e^{-x/\alpha}$$

Range $0 \leq x < \infty$.

Scale parameter, $\alpha > 0$, the mean.

Exponential Distribution with Threshold

PDF

$$\frac{e^{-(x-\theta)/\alpha}}{\alpha}$$

CDF

$$1 - e^{-(x-\theta)/\alpha}$$

Range $0 \leq x < \infty$.

Scale parameter, $\alpha > 0$, the mean.

Threshold parameter $\theta < \min(x)$.

Gamma Distribution

PDF

$$\frac{(x/\alpha)^{\beta-1} e^{-x/\alpha}}{\alpha\Gamma(\beta)}$$

CDF

$$\int_0^{x/\alpha} \frac{e^{-t} t^{\beta-1}}{\Gamma(\beta)} dt$$

where $\Gamma(\beta)$ is the Gamma function:

$$\Gamma(\beta) = \int_0^{\infty} e^{-t} t^{\beta-1} dt$$

Range $0 \leq x < \infty$.

Scale parameter $\alpha > 0$.

Shape parameter $\beta > 0$.

Gamma Distribution with Threshold

PDF

$$\frac{[(x - \theta)/\alpha]^{\beta-1} e^{-(x-\theta)/\alpha}}{\alpha\Gamma(\beta)}$$

CDF

$$\int_0^{(x-\theta)/\alpha} \frac{e^{-t} t^{\beta-1}}{\Gamma(\beta)} dt$$

Range $0 \leq x - \theta < \infty$.

Scale parameter $\alpha > 0$.

Shape parameter $\beta > 0$.

Threshold parameter $\theta < \min(x)$.

Half Normal Distribution

PDF

$$\frac{\sqrt{2/\pi}}{\sigma} e^{-x^2/2\sigma^2}$$

CDF

$$\int_{-\infty}^x \frac{\sqrt{2/\pi}}{\sigma} e^{-t^2/2\sigma^2} dt - 1$$

Range $0 \leq x < \infty$.

Scale parameter, $\sigma > 0$, the standard deviation

Reference for Application of Half Normal Distribution:

Chou, C., & H. Liu, (1998). "Properties of the half-normal distribution and its application to quality control", *Journal of Industrial Technology* Vol. 14(3) pp. 4-7

Largest Extreme Value Distribution

PDF

$$\frac{1}{\sigma} e^{-e^{-(x-\mu)/\sigma}} e^{-e^{-(x-\mu)/\sigma}}$$

CDF

$$e^{-e^{-(x-\mu)/\sigma}}$$

Range $-\infty < x < \infty$.

Location parameter, μ , the mode.

Scale parameter $\sigma > 0$.

Logistic Distribution

PDF

$$\frac{e^{(x-\mu)/\sigma}}{\sigma(1 + e^{(x-\mu)/\sigma})^2}$$

CDF

$$\frac{1}{1 + e^{-(x-\mu)/\sigma}}$$

Range $-\infty < x < \infty$.

Location parameter, μ , the mean.

Scale parameter $\sigma > 0$.

Loglogistic Distribution

PDF

$$\frac{e^{(\ln(x)-\mu)/\sigma}}{x\sigma(1 + e^{(\ln(x)-\mu)/\sigma})^2}$$

CDF

$$\frac{1}{1 + e^{-(\ln(x)-\mu)/\sigma}}$$

Range $0 < x < \infty$.

Location parameter, μ , the mean.

Scale parameter $\sigma > 0$.

Loglogistic Distribution with Threshold

PDF

$$\frac{e^{(\ln(x-\theta)-\mu)/\sigma}}{(x-\theta)\sigma(1 + e^{(\ln(x-\theta)-\mu)/\sigma})^2}$$

CDF

$$\frac{1}{1 + e^{-(\ln(x-\theta)-\mu)/\sigma}}$$

Range $0 < x < \infty$.

Location parameter, μ , the mean.

Scale parameter $\sigma > 0$.

Threshold parameter $\theta < \min(x)$.

Lognormal Distribution

PDF

$$\frac{1}{x\sqrt{2\pi\sigma^2}} e^{-.5(\ln(x)-\mu)^2/\sigma^2}$$

CDF

$$\int_{-\infty}^x \frac{1}{t\sigma\sqrt{2\pi}} e^{-.5(\ln(t)-\mu)^2/\sigma^2} dt$$

Range $0 < x < \infty$.

Scale parameter, μ , the mean of $\ln(x)$.

Shape parameter, $\sigma > 0$, the standard deviation of $\ln(x)$.

Lognormal Distribution with Threshold

PDF

$$\frac{1}{(x-\theta)\sqrt{2\pi\sigma^2}} e^{-.5(\ln(x-\theta)-\mu)^2/\sigma^2}$$

CDF

$$\int_{-\infty}^{x-\theta} \frac{1}{t\sigma\sqrt{2\pi}} e^{-.5(\ln(t)-\mu)^2/\sigma^2} dt$$

Range $0 < x < \infty$.

Scale parameter, μ , the mean of $\ln(x)$.

Shape parameter, $\sigma > 0$, the standard deviation of $\ln(x)$.

Threshold parameter $\theta < \min(x)$.

Reference for Lognormal Distribution with Threshold – Parameter Estimation:

Giesbrecht, F. and A.H. Kempthorne (1966). "Maximum Likelihood Estimation in the Three-parameter Lognormal Distribution", *Journal of the Royal Statistical Society*, B 38, pp. 257-264.

Normal Distribution

PDF

$$\frac{1}{\sqrt{2\pi\sigma^2}} e^{-.5(x-\mu)^2/\sigma^2}$$

CDF

$$\int_{-\infty}^x \frac{1}{\sqrt{2\pi\sigma^2}} e^{-.5(t-\mu)^2/\sigma^2} dt$$

Range $-\infty < x < \infty$.

Location parameter, μ , the mean.

Scale parameter, $\sigma > 0$, the standard deviation.

Note: For consistency with other reports in SigmaXL such as Descriptive Statistics, the standard deviation is estimated as the sample standard deviation using n-1 (rather than n).

Smallest Extreme Value Distribution

PDF

$$\frac{1}{\sigma} e^{(x-\mu)/\sigma} e^{-e^{(x-\mu)/\sigma}} \quad \text{CDF}$$

$$1 - e^{-e^{(x-\mu)/\sigma}}$$

Range $-\infty < x < \infty$.

Location parameter, μ , the mode.

Scale parameter $\sigma > 0$.

Weibull Distribution

PDF

$$\frac{\beta x^{\beta-1}}{\alpha^\beta} e^{-(x/\alpha)^\beta}$$

CDF

$$1 - e^{-(x/\alpha)^\beta}$$

Range $0 \leq x < \infty$.

Scale parameter, $\alpha > 0$, the characteristic life.

Shape parameter $\beta > 0$.

Weibull Distribution with Threshold

PDF

$$\frac{\beta x^{\beta-1}}{\alpha^{\beta}} e^{-[(x-\theta)/\alpha]^{\beta}}$$

CDF

$$1 - e^{-[x-\theta)/\alpha]^{\beta}}$$

Range $0 \leq x < \infty$.

Scale parameter, $\alpha > 0$, the characteristic life.

Shape parameter $\beta > 0$.

Threshold parameter $\theta < \min(x)$.

Reference for Weibull Distribution with Threshold – Parameter Estimation:

Lockhart, R.A. and M.A. Stephens (1994). "Estimation and Tests of Fit for the Three-parameter Weibull Distribution", *Journal of the Royal Statistical Society*, Vol.56(3), pp. 491-500.

Transformations

Box-Cox (Power) Transformation

$\lambda \neq 0$

$$\begin{aligned} z &= x^\lambda \\ x &= z^{1/\lambda} \end{aligned}$$

$\lambda = 0$

$$\begin{aligned} z &= \ln(x) \\ x &= e^z \end{aligned}$$

Range $0 < x < \infty$.

Shape parameter λ .

Note: The optimum shape parameter, λ , is derived using a grid search in which the criteria function is the standard deviation of the standardized transformed variable.

Box-Cox (Power) Transformation with Threshold

$\lambda \neq 0$

$$\begin{aligned} z &= (x - \theta)^\lambda \\ x &= z^{1/\lambda} + \theta \end{aligned}$$

$\lambda = 0$

$$\begin{aligned} z &= \ln(x - \theta) \\ x &= e^z + \theta \end{aligned}$$

Range $0 < x - \theta < \infty$.

Shape parameter λ .

Threshold parameter $\theta < \min(x)$.

Note: The parameters λ and θ are estimated using MLE as described above in **Box-Cox Distribution with Threshold**.

Johnson Transformation

The Johnson Transformation selects one of the three families of distribution: SB (bounded), SL (lognormal), and SU (unbounded) and the associated parameters so as to transform the data to be normally distributed. The methodology follows Chou et al (1998) and uses the Anderson Darling P-Value as the normality criteria.

Johnson Transformation - SB

$$z = \gamma + \eta \ln \left(\frac{x - \epsilon}{\lambda + \epsilon - x} \right)$$

$$x = \epsilon + \frac{1}{1 + e^{-(z-\gamma)/\eta}}$$

z is $N(0, 1)$

Range $\epsilon < x < \epsilon + \lambda$.

Location parameter $\eta > 0$.

Scale parameter $\lambda > 0$.

Shape parameter γ unbounded.

Shape parameter ϵ unbounded.

Johnson Transformation - SL

$$z = \gamma + \eta \ln(x - \epsilon)$$
$$x = \epsilon + e^{(z-\gamma)/\eta}$$

z is $N(0, 1)$

Range $x > \epsilon$

Location parameter $\eta > 0$.

Shape parameter γ unbounded.

Shape parameter ϵ unbounded.

Johnson Transformation - SU

$$z = \gamma + \eta \sinh^{-1}[(x - \epsilon)/\lambda]$$
$$x = \epsilon + \lambda \sinh[(z - \gamma)/\eta]$$

z is $N(0, 1)$

Range $-\infty < x < \infty$.

Location parameter $\eta > 0$.

Scale parameter $\lambda > 0$.

Shape parameter γ unbounded.

Shape parameter ϵ unbounded.

References for Johnson Transformation:

Chou, Y., A.M. Polansky, and R.L. Mason (1998). "Transforming Nonnormal Data to Normality in Statistical Process Control," *Journal of Quality Technology*, Vol. 30(2), pp. 133-141.

David, H.A. (1981). *Order Statistics*, John Wiley & Sons, New York

Tadikamalla, P., R. (1980). "Notes and Comments: On Simulating Nonnormal Distributions", *Psychometrika*, Vol. 45(2), pp. 273-279.

Monte Carlo Random Number Generation

A random number generator (RNG) is required when using Monte Carlo simulations in order to draw random samples from a specified distribution. The creation of a random number involves an algorithm that can automatically create long runs of numbers with good random properties. These algorithms are called pseudo-random number generators. The state of the RNG after each iteration is used as an input for the generation of the next random number. The initial state is called a seed; this seed is derived using the system clock.

A good RNG will produce a long run of numbers which are independent, so that there is no correlation between successive numbers. The RNG used in SigmaXL and DiscoverSim is based on a recur-with-carry “KISS+Monster” random number generator developed by George Marsaglia. This algorithm produces random integers between 0 and $2^{32} - 1$ and has a period of 10^{8859} . This implementation of the KISS+Monster algorithm has been tested and passes all of the Diehard tests (Marsaglia, 2003).

Pearson Family of Distributions

The Pearson family consists of eight distributions:

Type 0 Normal distribution.

Type 1 Beta distribution with minimum and maximum.

Type 2 Symmetric Beta distribution with minimum and maximum. The symmetry is imposed by requiring the two shape parameters to have the same value.

Type 3 Gamma distribution with threshold.

Type 4 Pearson IV distribution.

Type 5 Inverse Gamma distribution with threshold.

Type 6 F distribution with scale and threshold.

Type 7 Student's T distribution with location and scale

The Pearson methodology selects one member of the family using the methodology described in Johnson, Kotz and Balakrishnan, Volume 1, pp. 15-33. Specified Mean (0), Standard Deviation (1), Skewness and Excess Kurtosis are converted to the selected distribution's parameters.

Types 0 to 3 have been described above.

Type 4 Pearson Type IV Distribution

PDF

$$\frac{k(\nu, \omega)}{\lambda} \left(\frac{1}{1+z^2} \right)^\omega e^{-\nu \operatorname{atan}(z)}$$

CDF

$$\int_{-\infty}^z \frac{k(\nu, \omega)}{\lambda} \left(\frac{1}{1+t^2} \right)^\omega e^{-\nu \operatorname{atan}(t)} dt$$

where

$$\begin{aligned} b &= 2(\omega - 1) \\ a &= \sqrt{\frac{b^3 - b^2}{b^2 + \nu^2}} \\ z &= \frac{x - \mu}{a\lambda} - \frac{\nu}{b} \end{aligned}$$

and $k(\nu, \omega)$ is a scaling constant.

Range $0 \leq x < \infty$.

Location parameter μ

Scale parameter, $\lambda > 0$.

Shape parameter ν .

Shape parameter $\omega > 1.5$.

Type 5 Inverse Gamma Distribution with Threshold

PDF

$$\frac{\alpha^\beta}{\Gamma(\beta)} (x - \theta)^{-\beta-1} e^{-\alpha/(x-\theta)}$$

CDF

$$\frac{\gamma(\beta, \alpha/(x - \theta))}{\Gamma(\beta)}$$

where $\Gamma(s)$ is the the gamma function, and $\gamma(s, x)$ is the incomplete gamma function.

Range $0 \leq x - \theta < \infty$.

Scale parameter $\alpha > 0$.

Shape parameter $\beta > 0$.

Threshold parameter $\theta < \min(x)$.

Type 6 F Distribution with Scale and Threshold

PDF

$$\frac{1}{(x - \theta) B(.5\nu, .5\omega)} \sqrt{\frac{(\nu(x - \theta)/\alpha)^\nu \omega^\omega}{(\nu(x - \theta)/\alpha + \omega)^{\nu+\omega}}}$$

CDF

$$\int_{-\infty}^z \frac{1}{B(.5\nu, .5\omega)} t^{.5\nu-1} (1-t)^{.5\omega-1} dt$$

where

$$z = \frac{(\nu(x - \theta))}{(\nu x + \alpha\omega)}$$

and where B is the Beta function.

Range $0 \leq x - \theta < \infty$.

Scale parameter $\alpha > 0$.

Shape parameter $\nu > 0$.

Shape parameter $\omega > 0$. Threshold parameter $\theta < \min(x)$.

Type 7 Student's T distribution with Location and Scale

PDF

$$\frac{1}{\alpha \sqrt{\pi \nu}} \frac{\Gamma(.5(\nu + 1))}{\Gamma(.5\nu)} \left(\frac{\nu \alpha^2}{(x - \mu)^2 + \nu \alpha^2} \right)^{.5(\nu+1)}$$

CDF

$$.5 + .5 I_z(.5, .5\nu) \quad x \geq 0$$

$$.5 - .5 I_z(.5, .5\nu) \quad x < 0$$

where

$$z = \frac{(x - \mu)^2}{(x - \mu)^2 + \nu\alpha^2}$$

and where Γ is the gamma function, and $I_z(a, b)$ is the regularized incomplete beta function.

Range $-\infty \leq x < \infty$.

Location parameter μ .

Scale parameter $\alpha > 0$.

Shape parameter $\nu > 0$.

References for Pearson Family Monte Carlo Simulation:

Johnson, N.L., Kotz, S. and Balakrishnan, N. (1994). *Continuous Univariate Distributions-Volume 1*, Second Ed., Wiley, New York.

Johnson, N.L., Kotz, S. and Balakrishnan, N. (1995). *Continuous Univariate Distributions-Volume 2*, Second Ed., Wiley, New York.

Marsaglia, G. (2003). "Random Number Generators," *Journal Of Modern Applied Statistical Methods*, Vol 2(1), pp. 2-13.

Automatic Best Fit

SigmaXL uses the Anderson Darling P-Value as the criteria to determine best fit. All distributions and transformations are considered and the model with the highest AD P-Value is initially selected (denoted as *adpvalmax*). A search is then carried out for models that are close, having an AD P-Value greater than *adpvalmax* - 0.1 (with an added criteria that AD P-Value be > 0.2), but having fewer parameters than the initial best fit model. If a simpler model is identified, then this is selected as the best fit.

Since AD P-Values are not available for distributions with thresholds (other than Weibull), an estimate is obtained by transforming the data to normality and then using a modified Anderson Darling Normality test on the transformed data. The transformed z-values are obtained by using the inverse cdf of the normal distribution on the cdf of the nonnormal distribution. The Anderson Darling Normality test assumes a mean = 0 and standard deviation = 1.

This approach, unique to SigmaXL, is an extension of the Chou methodology used in Johnson Transformations and allows a goodness of fit comparison across all distributions and transformations.

Another approach to comparing models is Akaike's information criterion (AIC) developed by Hirotugu Akaike. AIC is the MLE log-likelihood with a penalty for the number of terms in the model (where the penalty factor also depends on the sample size, n). The AD P-Value method used by SigmaXL has the advantage that it is not limited to models with maximum likelihood parameter estimation.

References for Automatic Best Fit:

Chou, Y., A.M. Polansky, and R.L. Mason (1998). "Transforming Nonnormal Data to Normality in Statistical Process Control," *Journal of Quality Technology*, Vol. 30(2), pp. 133-141.

D'Agostino, R.B. and Stephens, M.A. (1986). *Goodness-of-Fit Techniques*, Marcel Dekker.

http://en.wikipedia.org/wiki/Akaike_information_criterion

Process Capability Indices (Nonnormal)

Z-Score Method (Default)

The Z-Score method for computing process capability indices utilizes the same z-value transformation to normality described above in **Automatic Best Fit**. The transformed z-values are obtained by using the inverse cdf of the normal distribution on the cdf of the nonnormal distribution. Normal based capability indices are then applied to the transformed z-values. This approach offers two key advantages: the relationship between the capability indices and calculated

defects per million is consistent across the normal and all nonnormal distributions, and short term capability indices Cp and Cpk can be estimated using the standard deviation from control chart methods on the transformed z-values. The Z-Score method was initially developed by Davis Bothe and expanded on by Andrew Sleeper. For further details, see Sleeper, *Six Sigma Distribution Modeling*.

Percentile (ISO) Method

The Percentile method to calculate process capability indices uses the following formulas:

$$Ppu = (USL - 50^{\text{th}} \text{ percentile}) / (99.865 \text{ percentile} - 50^{\text{th}} \text{ percentile})$$

$$Ppl = (50^{\text{th}} \text{ percentile} - LSL) / (50^{\text{th}} \text{ percentile} - 0.135 \text{ percentile})$$

$$Ppk = \min(Ppu, Ppl)$$

$$Pp = (USL - LSL) / (99.865 \text{ percentile} - 0.135 \text{ percentile})$$

References for Process Capability Indices:

Bothe, D.R. (1997). *Measuring Process Capability*, McGraw-Hill, New York.

Sleeper, A. (2006). *Six Sigma Distribution Modeling*, McGraw-Hill, New York.

Control Charts (Nonnormal)

Individuals – Original Data

The **Individuals – Original Data** chart displays the untransformed data with control limits calculated as:

UCL = 99.865 percentile

CL = 50th percentile

LCL = 0.135 percentile

The benefit of displaying this chart is that one can observe the original untransformed data. Since the control limits are based on percentiles, this represents the overall, long term variation rather than the typical short term variation. The limits will likely be nonsymmetrical.

Individuals/Moving Range – Normalized Data

The **Individuals/Moving Range – Normalized Data** chart displays the transformed z-values with control limits calculated using the standard Shewhart formulas for Individuals and Moving Range charts. The benefit of using this chart is that tests for special causes can be applied and the control limits are based on short term variation. The disadvantage is that one is observing transformed data on the chart rather than the original data.

SigmaXL's default setting is to display the two charts: Individuals – Original Data, Individuals & Moving Range – Normalized Data (with tests for special causes unchecked).

Exact and Monte Carlo P-Values for Nonparametric and Contingency Tests

Introduction

Nonparametric tests are popular because they do not assume that the sample data are normally distributed, but they do, however, assume that the test statistic follows a normal or chi-square distribution when computing the P-Value using “large sample” or “asymptotic” theory. Similarly, the Chi-Square test for independence used in $r \times c$ contingency tables assumes that the test statistic is distributed as chi-square. When the sample size is small or sparse, these assumptions may be invalid and can result in substantial error in the computed P-Value. For example, Cochran’s rule [5] for $r \times c$ contingency tables is a commonly used guideline: no more than 20% of the cells should have an expected value less than 5 and none should have an expected value less than 1.

The solution to the small sample problem is to compute exact P-Values using the true permutation distribution of the test statistic. However, explicit full enumeration of all possible outcomes quickly becomes computationally infeasible. For example, a 5x6 contingency table results in 1.6 billion possible outcomes.

Network algorithms by Mehta & Patel [6, 7] use techniques from operations research such as backward induction and forward probing to implicitly and efficiently solve the exact P-Value for contingency tables. These algorithms have also been extended to work with continuous data. For example, the nonparametric Mann-Whitney can be converted to a weighted singly ordered contingency table and the exact P-Value solved using the network algorithm.

Some data sets are too large even for the network algorithm to solve in a reasonable time period, but still do not satisfy Cochran’s rule. In these cases, Monte Carlo simulation is used to sample from the full enumeration of possible outcomes. 10,000 replications are typically used and 99% confidence intervals provided for the Monte Carlo P-Values. If the P-Value is low, 1 million replications are recommended.

Full Enumeration

To deal with a situation where the size of the data set does not justify the large sample asymptotic assumption, we will first consider full enumeration. To illustrate, consider the problem of evaluating the probability of getting a total of 3 or less when tossing two dice.

Evaluated for each of the $n = 36$ possible outcomes the sum of those situations (m) for which the sum of the two dice is 3 or less. Since $m = 3$ {1,1}, {1,2}, {2,1}, the probability is $m/n = 3/36 = 1/12$.

Full enumeration is 100% accurate, and always produces the exact P-Value. It works fine provided n is not too large. However, for 20 dice, n , the number of possible outcomes is 3656 trillion, which makes full enumeration infeasible. As mentioned above, even a relatively small problem such as a 5x6 contingency table results in 1.6 billion possible outcomes, again making full enumeration infeasible.

Exact

Instead of evaluating each of the n possible outcomes, shortcuts can be derived which reduces the number of outcomes that need to be evaluated. Starting with the first dice showing unity, once the second dice exceeds 2 the remaining four outcomes can be counted as a block, since the statistic for each case will exceed the test statistic. Thus rather than evaluating each as one would with full enumeration, the block is evaluated. When the first dice shows two, and once the second dice exceeds 1, the remaining 5 outcomes can again be counted as a block. Once the first dice shows 3, all the remaining outcomes can be evaluated as a block, since again every possible outcome will then exceed 3. Hence instead of enumerating 36 outcomes, one enumerates 6 outcomes {1,1}, {1,2}, {1, 3+}, {2,1}, {2,2+}, {3+,1+}. This type of model is referred to as Enhanced Enumeration. It uses full enumeration of all possible outcomes node by node, and evaluates the statistic as a block if the largest statistic possible from the current position exceeds the critical value (in which case the probability of this path is added to the total probability), or if the smallest statistic possible from the current position is less than the critical value, implying no augmentation of the total probability [15].

A second methodology involves creating a network of a set of nodes in a number of stages. The distances between nodes are defined so that the total distance of a path through the network is the corresponding value of the test statistic. At each node, the algorithm computes the shortest and longest path distances for all the paths that pass through that node. This works easily for those statistics that can be evaluated as the linear sum of partial statistics evaluated at each node. This methodology is commonly referred to in the literature as the Network Model, after the pioneering work of Mehta and Patel [1, 6, 7].

The Network Model has advantages over enhanced enumeration in that the problem is broken down to nodes. For example, if we had 4 dice, then there would be 5 nodes. The advantage of the network model is that the length (and hence probability) at each node for the remaining nodes can be estimated using a backward step through the network. Since networks tend to be sparse, the set of all paths through the network which have the same probability up to the current node can be retained. Thus, for any set of paths which have the same length up to the current node, one can quickly read off the longest and shortest remaining sub-paths to the terminal node. Specifically, as in enhanced enumeration, if every possible outcome does not contribute to the P-Value, the entire

set can be eliminated, while if every possible outcome contributes, then the probability of the entire set can be evaluated as a closed form.

Thus an exact evaluation of a probability is identical to full enumeration, but significantly faster. Again, as the size of the problem increases, the exact method also becomes infeasible since the number of outcomes that do have to be evaluated becomes too large. The methodology used to evaluate each of the exact nonparametric tests is described below.

Monte Carlo

Simulate the throw of the two dice $nrep$ times, and count the number of times (m) that the sum of the two dice is 3 or less. The Monte Carlo P-Value is then defined as $m/nrep$; this is an estimate of the exact P-Value, and can be made more accurate by increasing $nrep$. Thus, with 20 dice, one could choose a value of $nrep = 100,000$, and carry out the simulation to ascertain an unbiased estimate of the P-Value. Thus the beauty of Monte Carlo is that it can be used on any problem, no matter how large the total possible outcomes; however it is only asymptotically exact.

Two Tailed (Sided) Tests

A two tailed test is a statistical test in which the critical area of a distribution is two sided and tests whether a sample is either greater than or less than a certain range of values. A one tailed test, on the other hand, tests a one sided hypothesis, such as whether a sample is less than a certain range of values. Thus a two tailed test will be the sum of a tail on the left hand side of a distribution, and the corresponding tail on the right hand side of the distribution. If a distribution is symmetric, a two tailed test will simply be twice the value of the one tailed test.

Formally, the two tailed probability is

$$\Pr(|stat - \mu| \geq |statcrit - \mu|)$$

summed over all possible combinations, where $statcrit$ is the statistic evaluated for each possibility, and μ is the expected value of the statistic. For symmetric distributions, 2*1-Sided and 2-Sided will be identical. For non-symmetric distributions, such as the Runs Test and the Mann-Whitney Test, there will be a difference between 2*1-Sided and 2-Sided.

Sample Sizes for Exact

Nonparametric Tests

The following sample size guidelines are based on our recommendations and those given in Siegel [2] and Gibbons [3] for when exact nonparametric tests should be used rather than asymptotic (Normal or Chi-Square approximation):

Sign Test: $N \leq 50$ (exact is automatically utilized in the regular Sign Test)

Wilcoxon Signed Rank: $N \leq 15$

Mann-Whitney: Each sample $N \leq 10$

Kruskal-Wallis: Each sample $N \leq 5$

Mood's Median: Each sample $N \leq 10$

Runs Test (Above/Below) or Runs Test (Up/Down) Test: $N \leq 50$

These are minimum sample size guidelines. It is acceptable to use an exact test for larger sample sizes, but computation time can become an issue especially for tests with two or more samples. In those cases, one can always use a Monte Carlo P-Value with 99% confidence interval.

The maximum total sample size (total of all samples) that can be used for Nonparametric Exact or Monte Carlo is 1000. The Runs Test for Randomness and Wilcoxon Test automatically switch to large sample asymptotic if $N > 1000$.

One Way Chi-Square and Contingency Tables

Use Chi-Square Exact/Monte Carlo or Fisher's Exact/Monte Carlo if more than 20% of the cells have expected counts less than 5 or if any of the cells have an expected count less than 1 [4,5].

The maximum number of cells in a table is limited to 1000. For a One-Way Chi-Square the maximum sum of cell counts is 1000. For a Contingency Table, the maximum sum of cell counts is 5000.

Power

It is important to note that while exact P-Values are "correct," they do not increase (or decrease) the power of a small sample test, so they are not a solution to the problem of failure to detect a change due to inadequate sample size.

Validation of SigmaXL Exact P-Values

SigmaXL Exact P-Values are validated by comparison to textbook examples [2, 3], published paper examples [1, 6, 7, 8, 9] and other exact software such as StatXact, SPSS Exact, SAS and Matlab.

Monte Carlo P-Values are validated using 1e6 replications and compared against exact. Repeated simulations are used to validate the confidence intervals.

1 Sample Tests

Sign Test: The number of sign changes in one direction is distributed binomial, and so the binomial distribution function can be used directly to evaluate the tail probability [10].

Runs Test (Above/Below): Since the number of runs cannot exceed the length of the sample, it is easy to evaluate the probability of r runs when there is a binary variable with n_1 and n_2 elements [11].

Runs Test (Up/Down): Given the sample size and the number of up/down runs, the probability is provided by published tables [12].

Wilcoxon Test: Given the signed rank statistic, look at all combinations of ranks that could contribute to the observed value of the Wilcoxon critical value. For each of these cases, the number of combinations that actually contribute is evaluated. These counts are then converted to probabilities [13].

2 Sample Tests

Mann-Whitney Test: The continuous data is first converted to a weighted singly ordered contingency table and the exact P-Value is solved using the network algorithm.

Each possible combination of ranks is systematically evaluated, but the process is cut short for any given path if the largest statistic possible from the current position is less than the critical value (in which case all paths from the present are added to the feasible combinations), or if the smallest statistic possible from the current position is greater than the critical value (in which none of the paths from the present are added to the feasible combinations).

K Sample Tests

Kruskal-Wallis Test: Full enumeration.

Mood's Median Test: This test is evaluated using the Network Model. In each case, a chi-square statistic is evaluated. The number of combinations at the end of each path that has a chi-square statistic greater than the critical value is summed over all the enumerations, and is then converted to a probability.

Contingency Tests

One Way Chi-Square Test: For a one dimension vector of frequencies, and expected probabilities, the Chi Square test is evaluated using the Enhanced Enumeration model.

Two Way Fisher Exact Test: The exact probability for a two dimension $r \times c$ contingency table is evaluated using the Network model [1]. This is a one tailed test, and includes the probabilities of all the equivalent tables (those with same row and column counts) that have a lower or equal probability as the sample table. As in the one dimensional case, tests are made to see if the shortest path from the current position exceeds the path length of the observed table (in which case the probability of all paths from the current position are added to the probability, or whether the longest path from the current position is less than the path length of the current table, implying no augmentation of the probability.

References for Exact and Monte Carlo P-Values

- [1] Cyrus R. Mehta & Nitin R. Patel. (1983). "A network algorithm for performing Fisher's exact test in $r \times c$ contingency tables." *Journal of the American Statistical Association*, Vol. 78, pp. 427-434.
- [2] Siegel, S., & Castellan, N.J. (1988). *Nonparametric Statistics for the Behavioral Sciences* (2nd Ed.). New York, NY: McGraw-Hill.
- [3] Gibbons, J.D. and Chakraborti, S. (2010). *Nonparametric Statistical Inference* (5th Edition). New York: Chapman & Hall.
- [4] Yates, D., Moore, D., McCabe, G. (1999). *The Practice of Statistics* (1st Ed.). New York: W.H. Freeman.
- [5] Cochran WG. "Some methods for strengthening the common [chi-squared] tests." *Biometrics* 1954; 10:417–451.
- [6] Mehta, C. R.; Patel, N. R. (1997) "Exact inference for categorical data," [unpublished preprint](#). See Table 7 (p. 33) for validation of Fisher's Exact P-Value.
- [7] Mehta, C.R. ; Patel, N.R. (1998). "Exact Inference for Categorical Data." In P. Armitage and T. Colton, eds., *Encyclopedia of Biostatistics*, Chichester: John Wiley, pp. 1411–1422.
- [8] Mehta, C.R. ; Patel, N.R. (1986). "A Hybrid Algorithm for Fisher's Exact Test in Unordered $r \times c$ Contingency Tables," *Communications in Statistics - Theory and Methods*, 15:2, 387-403.

- [9] Narayanan, A. and Watts, D. "Exact Methods in the NPAR1WAY Procedure," SAS Institute Inc., Cary, NC.
- [10] Myles Hollander and Douglas A. Wolfe (1973), *Nonparametric Statistical Methods*. New York: John Wiley & Sons.
- [11] <https://onlinecourses.science.psu.edu/stat414/node/330>
- [12] Eugene S. Edgington, (1961). "Probability Table for Number of Runs of Signs of First Differences in Ordered Series," *Journal of the American Statistical Association*, Vol. 56, pp. 156-159.
- [13] David F. Bauer (1972), "Constructing confidence sets using rank statistics," *Journal of the American Statistical Association* Vol 67, pp. 687–690.
- [14] Hongsuk Jorn & Jerome Klotz (2002). "Exact Distribution of the K Sample Mood and Brown Median Test," *Journal of Nonparametric Statistics*, 14:3, 249-257.
- [15] Valz, P. D. and Thompson, M. E. (1994), "Exact Inference for Kendall's S and Spearman's Rho with Extensions to Fisher's Exact Test in rxc Contingency Tables," *Journal of Computational and Graphical Statistics*, Vol 3(4), pp. 459- 472.

Hypothesis Test Assumptions Report

The Hypothesis Test Assumptions Report is an optional text report for the 1 Sample t-Test, Paired t-Test, 2 Sample t-Test, One-Way ANOVA and Welch's ANOVA. The cells are color highlighted: Green (OK), Yellow (Warning) and Red (Serious Violation).

Normality

Each sample is tested for Normality using the Anderson-Darling (AD) test. If the AD P-Value is less than 0.05, the cell is highlighted as yellow (i.e., warning – proceed with caution). The Skewness and Kurtosis are reported and a note added, "See robustness and outliers." An example report is shown:

Anderson Darling P-Value = 0.030. Reject null hypothesis: "data are sampled from a normal distribution," so conclude that the assumption of normality is violated (at 95% confidence level). Skewness value = -0.9680 and Kurtosis value = 0.6796. See robustness and outliers.

If the AD P-Value is greater than or equal to 0.05, the cell is highlighted as green.

Anderson Darling p-value = 0.531. Fail to reject null hypothesis: "data are sampled from a normal distribution," so conclude that the assumption of normality is not violated.

A red highlight is not used for the Normality report.

Robustness

A minimum sample size for robustness to nonnormality is determined using the regression equations in [Basic Statistical Templates – Minimum Sample Size for Robust t-Tests and ANOVA](#). These were derived from extensive Monte Carlo simulations to determine a minimum sample size required for a test to be robust, given a specified sample Skewness and Kurtosis.

When there are 2 or more samples, a minimum sample size is calculated for each sample, and the largest value is (conservatively) considered as the minimum sample size required for the test.

If the AD Normality test for every sample is greater than or equal to 0.05, the following message is shown for robustness:

Not applicable for normal data.

If each sample size is greater than or equal to the minimum for robustness, the minimum sample size value is reported and the test is considered to be robust to the degree of nonnormality present in the sample data:

**Minimum sample size for a robust t-test = 22.
Since the sample size is greater than this, the t-test is robust to nonnormality.**

If any sample size is less than the minimum for robustness, the minimum sample size value is reported and a suitable Nonparametric Test is recommended. The cell is highlighted in red:

Minimum sample size for a robust t-test = 16. It is recommended that the Exact Nonparametric One Sample Wilcoxon Test be used (SigmaXL > Statistical Tools > Nonparametric Tests - Exact > 1 Sample Wilcoxon - Exact).

Recommended Nonparametric Tests – One Sample

The following Nonparametric Tests may be recommended for one sample:

1 Sample Sign

1 Sample Wilcoxon

1 Sample Wilcoxon – Exact

The 1 Sample Wilcoxon is more powerful than the 1 Sample Sign, but assumes that the data are symmetrical. Wilcoxon is therefore the default recommendation, but if the sample is moderately skewed then the Sign test is recommended.

Robustness to symmetry equations were developed in a manner similar to those used for robustness to nonnormality above. Monte Carlo simulations were run using sample sizes (N) of 20 to 2000. Observed alpha values were determined empirically from the P-Values of 100,000 replicate Wilcoxon tests. For a given sample size, nonnormal data with Skew = 0.01 to +1.0 was generated using the DiscoverSim Pearson Family function (see SigmaXL's DiscoverSim Workbook for details). The population test median (H0) used for the Wilcoxon test was estimated empirically by averaging the 100,000 sample medians. The test was considered robust when the simulated

observed alpha was between .04 and .06, .i.e. within +/- 20% of the specified alpha of .05. “Maximum Skewness” occurred when the observed alpha exceeded .04 or .06.

Other alpha values were not considered. Power was not considered.

This simulation was performed for H_a = “Not Equal To” (NE), “Greater Than” (GT), and “Less Than” (LT). A regression analysis was performed on the Log10 (Maximum Skewness) versus Log10(N), resulting in the following equations (with R-Square reported):

1. $\log_{10}(\text{Maximum Skewness NE}) = 0.5391 - 0.5028 \log_{10}(N)$ (R-Square = 99.5%)
2. $\log_{10}(\text{Maximum Skewness GT}) = 0.2380 - 0.5725 \log_{10}(N)$ (R-Square = 98.1%)
3. $\log_{10}(\text{Maximum Skewness LT}) = 0.2796 - 0.5459 \log_{10}(N)$ (R-Square = 99.7%)

Note that if Skewness is negative, equation 2 is used for LT and equation 3 is used for GT.

For example, given a sample size of 30 and H_a = “Not Equal To,” solving for **Maximum Skewness NE** in equation 1, the value is 0.626. Therefore, if the sample Skewness is less than or equal to 0.626, Wilcoxon is recommended. If the sample Skewness is greater than 0.626, the Sign Test is recommended.

If the Sign Test is recommended, Exact is not explicitly recommended since it is already “built in” the Sign Test for $N \leq 50$.

If the Wilcoxon Test is appropriate and $N \leq 15$, then Wilcoxon – Exact is recommended (using the rules of thumb in [Sample Sizes for Exact](#)).

Recommended Nonparametric Tests – Two Samples

The following Nonparametric Tests may be recommended for two samples:

2 Sample Mann-Whitney

2 Sample Mann-Whitney – Exact

If each sample $N \leq 10$, then Mann-Whitney – Exact is recommended (using the rules of thumb in [Sample Sizes for Exact](#)).

Recommended Nonparametric Tests – Three or More Samples (ANOVA)

The following Nonparametric Tests may be recommended for three or more samples:

Kruskal-Wallis

Kruskal-Wallis - Exact

Mood's Median Test

Mood's Median Test – Exact

Kruskal-Wallis is more powerful than Mood's Median Test, but Mood's Median is robust to outliers. Kruskal-Wallis is therefore the default recommendation, but if any sample has outliers that are Likely ($2.2 \cdot \text{IQR}$) or Extreme ($3 \cdot \text{IQR}$), then Mood's Median is recommended. See **Outliers (Boxplot Rules)** below for further details.

If Kruskal-Wallis is appropriate and each sample $N \leq 5$, then Kruskal-Wallis – Exact is recommended (using the rules of thumb in [Sample Sizes for Exact](#)).

If Mood's Median is appropriate and each sample $N \leq 10$, then Mood's Median – Exact is recommended.

Outliers (Boxplot Rules)

Each sample is tested for outliers using Tukey's Boxplot Rules: Potential ($> Q3 + 1.5 \cdot \text{IQR}$ or $< Q1 - 1.5 \cdot \text{IQR}$); Likely: $2.2 \cdot \text{IQR}$; Extreme: $3 \cdot \text{IQR}$. If outliers are present, a warning is given and recommendation to review the data with a Boxplot and Normal Probability Plot and to consider using a Nonparametric Test.

If no outliers are found, the cell is highlighted as green:



If a Potential or Likely outlier is found, the cell is highlighted as yellow:

Potential ($1.5 \cdot \text{IQR}$) outlier lower count = 1. It is recommended to review the data with graphical tools: Boxplot, Normal Probability Plot, Histogram and Run Chart / Control Chart. Consider using a Nonparametric Test.

Likely ($2 \cdot \text{IQR}$) outlier lower count = 1. It is recommended to review the data with graphical tools: Boxplot, Normal Probability Plot, Histogram and Run Chart / Control Chart. Consider using a Nonparametric Test.

Note that upper and lower outliers are distinguished.

If an Extreme outlier is found, the cell is highlighted as red:

Extreme (3*IQR) outlier lower count = 1. It is recommended to review the data with graphical tools: Boxplot, Normal Probability Plot, Histogram and Run Chart / Control Chart. Consider using a Nonparametric Test.

The Anderson Darling normality test is applied to the sample data with outliers excluded. If this results in an AD P-Value that is greater than 0.1, a notice is given, "Excluding the outliers, data are inherently normal." The cell remains highlighted as yellow or red.

Extreme (3*IQR) outlier upper count = 1. It is recommended to review the data with graphical tools: Boxplot, Normal Probability Plot, Histogram and Run Chart / Control Chart. Consider using a Nonparametric Test. Note that if the outlier is removed, Anderson Darling P-Value = 0.961. Excluding the outlier, data are inherently normal.

Randomness (Independence)

Each sample is tested for randomness (serial independence) using the Exact Nonparametric Runs Test. If the sample data is not random, a warning is given and recommendation to review the data with a Run Chart or Control Chart. See [Nonparametric Runs Test for Randomness - Exact](#) for further details. Note that for this report, the option **Values Equal to Median** is set to both **Counted as Below** and **Counted as Above**, with the reported P-Value being the larger of the two. This results in a more conservative test to minimize false alarms.

If the Exact Nonparametric Runs Test P-Value is greater than or equal to 0.05, the cell is highlighted as green:

Nonparametric Runs Test (Exact) P-Value = 0.396. Fail to reject null hypothesis: "data are random," so conclude that the assumption of randomness (independence) is not violated.

If the Exact Nonparametric Runs Test P-Value is less than 0.05, but greater than or equal to 0.01, the cell is highlighted as yellow:

Nonparametric Runs Test (Exact) P-Value = 0.017. Reject null hypothesis: "data are random," so conclude that the assumption of randomness (independence) is violated (at 95% confidence level). It is recommended to review the data with a Run Chart or Control Chart.

If the Exact Nonparametric Runs Test P-Value is less than 0.01, the cell is highlighted as red:

Nonparametric Runs Test (Exact) P-Value = 0.004. Reject null hypothesis: "data are random," so conclude that the assumption of randomness (independence) is violated (at 99% confidence level). It is recommended to review the data with a Run Chart or Control Chart.

Equal Variance

The test for Equal Variances is applicable for two or more samples. If all sample data are normal, the F-Test (2 sample) or Bartlett's Test (3 or more samples) is utilized. If any samples are not normal, i.e., have an AD P-Value < .05, Levene's test is used. If the variances are unequal and the test being used is the equal variance option, then a warning is given and Unequal Variance (2 sample) or Welch's Test (3 or more samples) is recommended.

If the test for Equal Variances P-Value is $\geq .05$, the cell is highlighted as green:

F-test for Equal Variances P-Value = 0.479. Fail to reject null hypothesis: "variances are equal," so conclude that the assumption of equal variances (or standard deviations) is not violated.

If the test for Equal Variances P-Value is $\geq .05$, but the Assume Equal Variances is unchecked (2 sample) or Welch's ANOVA (3 or more samples) is used, the cell is highlighted as yellow:

F-test for Equal Variances P-Value = 0.479. Fail to reject null hypothesis: "variances are equal," so conclude that the assumption of equal variances (or standard deviations) is not violated. The Assume Equal Variances option should be checked.

If the test for Equal Variances P-Value is < .05, and the Assume Equal Variances is checked (2 sample) or regular One-Way ANOVA (3 or more samples) is used, the cell is highlighted as red:

Levene's Test for Equal Variances P-Value = 0.014. Reject null hypothesis: "variances are equal," so conclude that the assumption of equal variances (or standard deviations) is violated (at 95% confidence level). It is recommended that Welch's ANOVA be used (SigmaXL > Statistical Tools > Equal Variance Tests > Welch's ANOVA) or alternatively use a nonparametric test.

If the test for Equal Variances P-Value is $< .05$, and the Assume Equal Variances is unchecked (2 sample) or Welch's ANOVA (3 or more samples) is used, the cell is highlighted as green:

Levene's Test for Equal Variances P-Value = 0.014. Reject null hypothesis: "variances are equal," so conclude that the assumption of equal variances (or standard deviations) is violated (at 95% confidence level). Welch's ANOVA is appropriate.

Attribute Measurement Systems Analysis

Percent Confidence Intervals (Exact Versus Wilson Score)

Confidence intervals for binomial proportions have an "oscillation" phenomenon where the coverage probability varies with n and p . Exact (Clopper-Pearson) is strictly conservative and will guarantee the specified confidence level as a minimum coverage probability, but results in wide intervals. This is recommended only for applications requiring strictly conservative intervals. Wilson Score has mean coverage probability matching the specified confidence interval. Since the Wilson Score intervals are narrower and thereby more powerful, they are recommended for use in Attribute MSA studies due to the small sample sizes typically used [1, 2, 3].

Kappa

Kappa (κ) measures the degree of agreement between appraisers when they assess the same sample. Cohen's Kappa, used in the Attribute MSA template, is only defined for two raters. Fleiss' Kappa can be used for multiple assessors [4]. Thus Kappa provides a mechanism for evaluating the consistency of assessors.

Let \bar{P} be the measure of agreement between the assessors, and \bar{P}_e be the measure of agreement that would occur by chance. Then:

$$\kappa = \frac{\bar{P} - \bar{P}_e}{1 - \bar{P}_e}$$

The denominator provides the measure of agreement above chance, while the numerator provides the actual measure of agreement above chance. Thus κ is constrained between -1 and +1, with +1 implying complete consistency or perfect agreement between assessors, zero implying no more consistency between assessors than would be expected by chance and -1 implying perfect disagreement.

Rule-of-Thumb Interpretation for Kappa

Fleiss [4, p. 609] gives the following rule of thumb for interpretation of Kappa:

Kappa: ≥ 0.75 or so signifies excellent agreement, for most purposes, and ≤ 0.40 or so signifies poor agreement.

AIAG recommends the Fleiss guidelines [5].

In Six Sigma process improvement applications, a more rigorous level of agreement is commonly used. Futrell [6] recommends:

The lower limit for an acceptable Kappa value (or any other reliability coefficient) varies depending on many factors, but as a general rule, if it is lower than 0.7, the measurement system needs attention. The problems are almost always caused by either an ambiguous operational definition or a poorly trained rater. Reliability coefficients above 0.9 are considered excellent, and there is rarely a need to try to improve beyond this level.

SigmaXL uses the guidelines given by Futrell and color codes Kappa as follows: ≥ 0.9 is green, 0.7 to < 0.9 is yellow and < 0.7 is red. This is supported by the following relationship to Spearman Rank correlation and Percent Effectiveness/Agreement (applicable when the response is binary with an equal proportion of good and bad parts):

Kappa = 0.7; Spearman Rank Correlation = 0.7; Percent Effectiveness = 85%; Percent Agreement = 85% (two trials)

Kappa = 0.9; Spearman Rank Correlation = 0.9; Percent Effectiveness = 95%; Percent Agreement = 95% (two trials)

(Note that these relationships do not hold if there are more than two response levels or the reference proportion is different than 0.5).

Concordance

Let two assessors rank n subjects, such that x is the rankings of assessor 1, and y is the rankings assessor 2. If for any pair of observations (i, j) , the ranking of the two assessors agree - that is $x_i > x_j$ and $y_i > y_j$, or $x_i < x_j$ and $y_i < y_j$, then the pair of observations are said to be concordant. Pairs that are not concordant are said to be discordant. If x were a measurement of exercise, and y a measurement of fitness, then we would expect high levels of concordance.

Kendall's Coefficient of Concordance

Kendall's Coefficient of Concordance (Kendall's W) is a measure of association for discrete ordinal data, used for assessments that do not include a known reference standard. Kendall's coefficient of concordance ranges from 0 to 1: A coefficient value of 1 indicates perfect agreement. If the coefficient = 0, then agreement is random, i.e., the same as would be expected by chance.

There is a close relationship between ρ_s , Spearman's correlation coefficient and W , Kendall's coefficient of concordance. W can be calculated directly from the mean ($\bar{\rho}_s$) of the pairwise Spearman correlations using the relationship [7]:

$$W = \frac{(k-1)\bar{\rho}_s + 1}{k}$$

where k is the number of trials (within) or trials*appraisers (between) among which Spearman's correlation coefficients are computed.

Confidence limits for Kendall's Concordance cannot be solved analytically, so are estimated using bootstrapping. Ruscio [8] demonstrates the bootstrap for Spearman's correlation and we apply this method to Kendall's Concordance. The data are row wise randomly sampled with replacement to provide the bootstrap sample ($N = 2000$). W can be derived immediately from the mean value of the Spearman's correlation matrix from the bootstrap sample. The bias corrected values of α and $1 - \alpha$ are then used to determine the confidence interval from the sorted bootstrap vector W . In Efron's bias-corrected and accelerated bootstrap, the value of α is adjusted to account for skew and/or bias. This has been determined to be a superior bootstrap method compared to the simple percentile method [9]. A fixed seed is used so that the report results are consistent.

Rule-of-Thumb Interpretation for Kendall's Coefficient of Concordance

SigmaXL uses the following "rule-of-thumb" interpretation guidelines: ≥ 0.9 very good agreement (color coded green); 0.7 to < 0.9 marginally acceptable, improvement should be considered (yellow); and < 0.7 unacceptable (red).

This is consistent with Kappa and is supported by the relationship to Spearman's correlation.

Note however that in the case of Within Appraiser agreement with only two trials, then the rules should be adjusted: very good agreement is ≥ 0.95 and unacceptable agreement is < 0.85 .

Kendall's Correlation Coefficient

Kendall's Correlation Coefficient (Kendall's tau-b) is a measure of association for discrete ordinal data, used for assessments that include a known reference standard. Kendall's correlation coefficient ranges from -1 to 1: A coefficient value of 1 indicates perfect agreement. If coefficient = 0, then agreement is random, i.e., the same as would be expected by chance. A coefficient value of -1 indicates perfect disagreement. Kendall's Correlation Coefficient is a measure of rank correlation, similar to the Spearman rank coefficient, but uses concordant and discordant pairs and corrects for ties.

Let two assessors rank n subjects, such that x is the rankings of assessor 1, and y is the rankings of assessor 2. If for any pair of observations (i, j) , the ranking of the two assessors agree - that is $x_i > x_j$ and $y_i > y_j$, or $x_i < x_j$ and $y_i < y_j$, then the pair of observations are said to be concordant. Pairs that are not concordant are said to be discordant. If x were a measurement of exercise, and y a measurement of fitness, then we would expect high levels of concordance.

Define:

$$C_{ij} = \sum_{h < i} \sum_{k < j} x_{hk} + \sum_{h > i} \sum_{k > j} x_{hk}$$

$$D_{ij} = \sum_{h < i} \sum_{k > j} x_{hk} + \sum_{h > i} \sum_{k < j} x_{hk}$$

Then the number of concordant (n_c) and discordant (n_d) pairs are derived as:

$$n_c = 0.5P \text{ where } P = \sum_{i,j} x_{ij} C_{ij}$$

$$n_d = 0.5Q \text{ where } Q = \sum_{i,j} x_{ij} D_{ij}$$

and the Kendall Correlation τ_b is:

$$\tau_b = \frac{2(n_c - n_d)}{\sqrt{D_r D_c}}$$

where:

$$D_r = n^2 - \sum_i R_i^2$$

$$D_c = n^2 - \sum_j C_j^2$$

Rule-of-Thumb Interpretation for Kendall's Correlation Coefficient

SigmaXL uses the following "rule-of-thumb" interpretation guidelines:

≥ 0.8 very good agreement (color coded green);

0.6 to < 0.8 marginally acceptable, improvement should be considered (yellow);

< 0.6 unacceptable (red).

These values were determined using Monte Carlo simulation with correlated integer uniform distributions. They correspond approximately to Spearman 0.7 and 0.9 when there are 5 ordinal response levels (1 to 5). With 3 response levels, the rule should be modified to 0.65 and 0.9.

References for Attribute Measurement Systems Analysis

- [1] Agresti, A. and Coull, B.A.(1998). "Approximate is Better than "Exact" for Interval Estimation of Binomial Proportions." *The American Statistician*, 52, 119–126.
- [2] Clopper, C.J., and Pearson, E.S. (1934). "The Use of Confidence or Fiducial Limits Illustrated in the Case of the Binomial." *Biometrika*, 26, 404–413.
- [3] Newcombe, R. (1998a). "Two-sided confidence intervals for the single proportion: Comparison of seven methods." *Statistics in Medicine*, 17, 857-872.
- [4] Fleiss, J.L. (2003). *Statistical Methods for Rates and Proportions*, 3rd Edition, Wiley & Sons, NY.
- [5] Automotive Industry Action Group AIAG (2010). *Measurement Systems Analysis MSA Reference Manual*, 4th Edition, p. 137.
- [6] Futrell, D. (May 1995). "When Quality Is a Matter of Taste, Use Reliability Indexes," *Quality Progress*, 81-86.
- [7] Siegel, S., & Castellan, N.J. (1988). *Nonparametric statistics for the behavioral sciences* (2nd Ed.). New York, NY: McGraw-Hill. p. 262.
- [8] Ruscio, J. (2008). "Constructing Confidence Intervals for Spearman's Rank Correlation with Ordinal Data: A Simulation Study Comparing Analytic and Bootstrap Methods," *Journal of Modern Applied Statistical Methods*, Vol. 7, No. 2, 416-434.
- [9] Bradley Efron, (1987). "Better bootstrap confidence intervals (with discussion)," *J. Amer. Statist. Assoc.* Vol. 82, 171-200.
- [10] http://en.wikipedia.org/wiki/Kendall_tau_rank_correlation_coefficient

Chi-Square Tests and (Contingency) Table Associations

This section describes a number of tests and measures of association that are derived from a contingency table. These tables depict the observed counts of two variables, shown in r rows and c columns, with a row sum vector R and a column sum vector C . The total number of cells is n . In a contingency table, x_{ij} represents the number of elements that correspond to row attribute i and column attribute j .

The contingency table is described using the following summary statistics:

n_{ij} , the observed counts or number of elements in each cell.

n_i , the sum of the elements in each row.

n_j , the sum of the elements in each column.

$e_{ij} = n_i n_j / n$, the expected count or row-column product for that cell as a proportion of the total count.

$r_{ij} = n_{ij} - e_{ij}$, the raw residual or difference between the actual and expected cell count.

$s_{ij} = r_{ij} / \sqrt{e_{ij}}$, the standardized residual.

$a_{ij} = s_{ij} / \sqrt{(1 - n_i/n)(1 - n_j/n)}$, the adjusted (standardized) residual expressed as a z score.

s_{ij}^2 , the cell's contribution to the chi-squared statistic.

Chi-Square Tests for Nominal Categories

Pearson Chi-squared

Pearson's chi-squared test [3] is applied to a contingency table to evaluate how likely it is that this observed table occurs by chance. Given the observed marginals (row and column counts), the expected cell observation $e_{ij} = R_i C_j / n$. The statistic:

$$\chi^2 = \sum_{i,j} (x_{ij} - e_{ij})^2 / e_{ij}$$

is distributed χ^2 with $df = (r-1)(c-1)$ degrees of freedom.

Likelihood Ratio Chi-squared

The likelihood ratio chi-squared test [3] is applied to a contingency table to evaluate the likelihood of the observed table under the null hypothesis relative to the maximum likelihood. The statistic:

$$\chi^2 = 2 \sum_{i,j} e_{ij} \ln(x_{ij}/e_{ij})$$

is distributed χ^2 with df degrees of freedom.

McNemar-Bowker Symmetry

For square tables, the McNemar statistic [11] tests the hypothesis that the contingency table is symmetric. It is defined as:

$$M = \sum_{i < j} \frac{(x_{ij} - x_{ji})^2}{x_{ij} + x_{ji}}$$

for all $x_{ij} + x_{ji} > 0$. Under the null, M is distributed χ^2 with $n(n-1)/2$ degrees of freedom.

Measures of Association for Nominal Categories

Confidence bounds are evaluated as $t \pm z_{\alpha} s$, where z is the inverse of the standard normal cdf, t is the statistic and s is the asymptotic standard error (ASE₁) of the statistic. See [9,12] for ASE₁ formulas.

Perason's Phi

The Phi (ϕ) statistic [1,9] is defined as:

$$\phi = \sqrt{\chi^2/n}$$

where χ^2 is the Pearson Chi-squared statistic. For a 2x2 table, ϕ is a measure of association for two binary variables, and equivalent to Pearson Correlation, measures the percent of concentration of cases on the diagonal. See Cohen [15] for interpretation rules-of-thumb.

Cramer's V

The Cramer's V statistic [3, 9] measures the association between two variables as a percentage of their maximum possible variation. It has a range of $0 \leq V \leq 1$. It is defined as:

$$V = \sqrt{\chi^2/(n(m-1))}$$

where χ^2 is the Pearson Chi-squared statistic, and $m = \min(r, c)$

Contingency Coefficient

The contingency statistic [3] is defined as:

$$z = \sqrt{\chi^2/(n + \chi^2)}$$

where χ^2 is the Pearson Chi-squared statistic. Unlike the other statistics, it does not have a specified range.

Cohen's Kappa

For square tables, Cohen's κ statistic [3,5] tests the null hypothesis that there is no association between the row and column variables. It is calculated from the observed and expected frequencies on the diagonal of a square contingency table. It is frequently used to provide a mechanism for evaluating the consistency of assessors. It is defined as:

$$\kappa = \frac{n \sum_i x_{ii} - \sum_i R_i C_i}{n^2 - \sum_i R_i C_i}$$

A value of $\kappa = 0$ implies no association, while $\kappa = 1$ implies perfect association.

Goodman-Kruskal Lambda

Lambda [3] is defined as the proportional reduction in error of prediction of the dependent variable as a consequence of using the information in the independent variable.

Column as dependent, row as independent:

$$\lambda_r = \frac{\sum_i x_{im} - C_m}{n - C_m}$$

Row as dependent, column as independent:

$$\lambda_c = \frac{\sum_j x_{mj} - R_m}{n - R_m}$$

Symmetric version:

$$\lambda_s = \frac{\sum_i x_{im} + \sum_j x_{mj} - C_m - R_m}{2n - C_m - R_m}$$

where x_{im} and x_{mj} are the largest cell count in row i and column j respectively, and R_m and C_m are the largest row and column subtotals respectively.

Goodman-Kruskal Tau

Tau [9] is the same as lambda, except the probabilities are specified by marginal proportions.

Column as dependent, row as independent:

$$\tau_r = \frac{n \sum_{i,j} x_{ij}^2 / R_i - \sum_j C_j^2}{n^2 - \sum_j C_j^2}$$

Row as dependent, column as independent:

$$\tau_c = \frac{n \sum_{i,j} x_{ij}^2 / C_j - \sum_i R_i^2}{n^2 - \sum_i R_i^2}$$

Thiel's Uncertainty

Uncertainty [12] is defined as the proportional reduction in the uncertainty (entropy) of the dependent variable as a consequence of using the information in the independent variable.

Column as dependent, row as independent:

$$U_{c|r} = \frac{U_r + U_c - U_{rc}}{U_c}$$

Row as dependent, column as independent:

$$U_{r|c} = \frac{U_r + U_c - U_{rc}}{U_r}$$

Symmetric version:

$$U_s = \frac{2(U_r + U_c - U_{rc})}{U_r + U_c}$$

where:

$$U_r = -\sum_i R_i/n \ln(R_i/n)$$

$$U_c = -\sum_j C_j/n \ln(C_j/n)$$

$$U_{rc} = -\sum_{i,j} x_{ij}/n \ln(x_{ij}/n)$$

Tests for Ordinal Categories

Concordant - Discordant

Let two assessors rank n subjects, such that x is the rankings of assessor 1, and y is the rankings of assessor 2. If for any pair of observations (i, j) , the ranking of the two assessors agree - that is $x_i > x_j$ and $y_i > y_j$, or $x_i < x_j$ and $y_i < y_j$, then the pair of observations are said to be concordant [2]. Pairs that are not concordant are said to be discordant. If x were a measurement of exercise, and y a measurement of fitness, then we would expect high levels of concordance.

Define:

$$C_{ij} = \sum_{h<i} \sum_{k<j} x_{hk} + \sum_{h>i} \sum_{k>j} x_{hk}$$

$$D_{ij} = \sum_{h<i} \sum_{k>j} x_{hk} + \sum_{h>i} \sum_{k<j} x_{hk}$$

Then the number of concordant (n_c) and discordant (n_d) pairs are derived as:

$$n_c = 0.5P \text{ where } P = \sum_{i,j} x_{ij} C_{ij}$$

$$n_d = 0.5Q \text{ where } Q = \sum_{i,j} x_{ij} D_{ij}$$

The probability value returned for this test is given in **Kendall's Tau-b** below.

Spearman Rank Correlation Coefficient

The Spearman rank correlation coefficient (ρ) [3] provides a non-parametric measure of the association between the two variables. It has a range of $-1 \leq \rho \leq 1$. Let RI_i and CI_j be the row and column rank scores. Then ρ is evaluated as:

$$\rho = \frac{\sum_{i,j} x_{ij} (RI_i - n/2)(CI_j - n/2)}{\sqrt{(n^3 - \sum R^3)(n^3 - \sum C^3)}/12}$$

Under the null hypothesis that there is no correlation,

$$t = \frac{\rho\sqrt{n-2}}{\sqrt{1-\rho^2}}$$

has a t distribution with $n - 2$ degrees of freedom.

Measures of Association for Ordinal Categories

Confidence bounds are evaluated as $t \pm z_{\alpha} s$, where z is the inverse of the standard normal cdf, t is the statistic and s is the asymptotic standard error (ASE₁) of the statistic. See [9,12] for ASE₁ formulas.

Spearman Rank Correlation Coefficient

Since the Spearman rank correlation coefficient is not distributed normally, the Fisher transform is used to evaluate the confidence bounds. See Bonnett & Wright [4] for details.

Kendall's Tau-b

The Kendall τ_b coefficient [10], like gamma, is a non-parametric measure of association using rank correlation, but makes a correction for ties. τ_b is defined as:

$$\tau_b = \frac{2(n_c - n_d)}{\sqrt{D_r D_c}}$$

where:

$$D_r = n^2 - \sum_i R_i^2$$

$$D_c = n^2 - \sum_j C_j^2$$

τ_b ranges from +1 to -1. A value of zero implies independence. Given s :

$$s = 2\sqrt{\frac{\sum_{i,j} x_{ij}(C_{ij} - D_{ij})^2 - (P - Q)^2/n}{D_r D_c}} \quad \text{under the null hypothesis,}$$

then τ_b/s is distributed standard normal. This is used to provide a P-Value for the **Concordant – Discordant test**.

Kendall-Stuart Tau-c

Stuart's τ_c coefficient [1] is a measure of association, but makes a correction for ties as well as table size. τ_c is defined as:

$$\tau_c = \frac{2m(n_c - n_d)}{n^2(m - 1)}$$

where $m = \min(r, c)$. τ_c ranges from +1 to -1. A value of zero implies independence.

Goodman-Kruskal Gamma

The Goodman-Kruskal's gamma statistic [8] provides a measure of association using rank correlation, and is based on the number of concordant and discordant pairs. It is defined as:

$$\gamma = \frac{n_c - n_d}{n_c + n_d}$$

Somer's D

Somer's D statistic [1,14] is also a measure of association, similar to τ_b , but differs in respect to a consideration as to whether the row variable is considered the independent variable (and the column variable dependent), or *vice versa*.

Column as dependent, row as independent:

$$S_r = 2(n_c - n_d)/D_r$$

Row as dependent, column as independent:

$$S_c = 2(n_c - n_d)/D_c$$

Symmetric version:

$$S_s = 4(n_c - n_d)/(D_r + D_c)$$

References for Chi-Square Tests and (Contingency) Table Associations

- [1] Agresti, A. (2010). *Analysis of Ordinal Categorical Data* (Second Ed.), New York: John Wiley and Sons.
- [2] Agresti, A. (2013). *Categorical Data Analysis* (Third Ed.), New York: John Wiley and Sons
- [3] Bishop, Y, Fienberg, S and Holland, P. (1975). *Discrete Multivariate Analysis: Theory and Practice*. Cambridge, Mass.: MIT Press.
- [4] Bonett, D.G, and Wright, T.A. (2000). "Sample size requirements for estimating Pearson, Kendall and Spearman correlations", *Psychometrika* 65 (1),pp 23-28.
- [5] Cohen, J. (1960). "A coefficient of agreement for nominal scale". *Educational and Psychological Measurement* 20 (1) pp. 37-46.
- [6] Fisher, R.A. (1915). "Frequency distribution of the values of the correlation coefficient in samples from an indefinitely large population". *Biometrika* 10 (4), pp. 507-521.
- [7] Fisher, R.A. (1922). "On the interpretation of χ^2 from contingency tables, and the calculation of P". *Journal of the Royal Statistical Society* 85 (1), pp. 87-94.
- [8] Goodman, L. A. and W. H. Kruskal, W.H. (1972). "Measures for association for cross-classification", *Journal of the American Statistical Association* 67. pp. 415-421.
- [9] IBM SPSS Statistics 24 Algorithms, Crosstabs Algorithms, pp. 161 – 182.
ftp://public.dhe.ibm.com/software/analytics/spss/documentation/statistics/24.0/en/client/Manuals/IBM_SPSS_Statistics_Algorithms.pdf
- [10] Kendall, M. (1938). "A New Measure of Rank Correlation". *Biometrika* 30 (1–2) pp. 81-89.
- [11] McNemar, Q.(1947). "Note on the sampling error of the difference between correlated proportions or percentages". *Psychometrika* 12 (2). pp. 153-157.
- [12] SAS/STAT® 14.1 Users Guide - Measures of Association,
http://support.sas.com/documentation/cdl/en/statug/68162/HTML/default/viewer.htm#statug_freq_details22.htm
- [13] Snedecor, G.W and Cochran, W.G. (1989), *Statistical Methods* 8th ed. Ames, Iowa. Blackwell Publishing.
- [14] Somers, R. H. (1962). "A new asymmetric measure of association for ordinal variables". *American Sociological Review* 27. pp. 799-811.

- [15] Cohen, J. (1988), *Statistical Power Analysis for the Behavioral Sciences* (Second Ed.), Lawrence Erlbaum Associates, pp. 215 - 227.

Multiple Comparison of Means and Variances (a.k.a. Post-Hoc Tests)

This set of tests is undertaken after an ANOVA or Equal Variance test, and are often referred to as “post hoc tests”. After the null of equal means or variances has been rejected, the following tests are computed for each of the $.5k(k - 1)$ sample pairs. Thus this permits the researcher to ascertain exactly which group(s) are responsible for the rejection of the null.

One-Way ANOVA

Fisher

Fisher’s least significant difference test is a set of individual t-tests on each of the sample pairs, but uses the pooled standard deviation from all groups, which increases the power. The test statistic is:

$$t = \frac{\bar{y}_i - \bar{y}_j}{s\sqrt{1/n_i + 1/n_j}}$$

where s is the pooled standard error for all groups, and is assumed distributed Student t.

There is no adjustment for family-wise error rate, so should only be used for $k = 3$.

Tukey

Tukey’s test is a set of individual t-tests on each of the sample pairs, but uses the pooled standard deviation from all group. The test statistic is:

$$t = \frac{\sqrt{2}(\bar{y}_i - \bar{y}_j)}{s\sqrt{1/n_i + 1/n_j}}$$

and is distributed as the Studentized range [1].

Dunnett

The Dunnett test compares each group to a given control group. For each group, a t statistic is calculated using the Fisher statistic. The distribution of this statistic is a $k - 1$ multivariate t, with all $k - 1$ arguments being the same Fisher t statistic, and with correlation matrix with off diagonal elements $\sqrt{2}$ for the balanced case, and weighted by sample size for the unbalanced case [1].

Welch's ANOVA (Unequal Variance)

Pairwise Welch

This is the pairwise two sample t-test for unequal variances. Like Fisher, there is no adjustment for family-wise error rate so should only be used for $k = 3$.

Games-Howell

The Games-Howell test does not assume equal variance for each group. The variance of each group is used. The test statistic is:

$$t = \frac{\sqrt{2} |\bar{y}_i - \bar{y}_j|}{\sqrt{s_i^2/n_i + s_j^2/n_j}}$$

And is distributed as the Studentized range [2].

Bartlett's Test for Equal Variance

F-Test

The standard F test statistic for testing the null of equal variances of two groups is:

$$F = \frac{s_i^2}{s_j^2}$$

where s_i^2 is the variance of group i , and is distributed as F_{n_i-1, n_j-1} .

There is no adjustment for family-wise error rate so should only be used for $k = 3$.

F-Test with Bonferroni Correction

The Bonferroni correction multiplies the F-Test probability by a factor of $.5k(k-1)$.

Levene's Test for Equal Variance

Pairwise Levene (a.k.a. Brown-Forsythe)

The pairwise Levene test statistic for testing the null of equal variances of two groups is a simple F statistic:

$$F = \frac{SSA/(k-1)}{SSE/(n-k)}$$

where $n = n_i + n_j$, and SSA and SSE are the sample sum of squares and the error sum of squares of the transformed variable respectively:

$$z = |x - \text{median}(x)|$$

The statistic is distributed as F with $k-1$ and $n-k$ degrees of freedom [3,4].

There is no adjustment for family-wise error rate, so should only be used for $k = 3$.

Tukey ADM (Absolute Deviations from Median)

Tukey's method is applied to the ADM (Absolute Deviations from Median). Conceptually, this is like an extension of the ANOMV-LEV chart, Analysis of Means version of the Levene Test. See [5] Nelson, Wludyka and Copeland (2005), page 65 for discussion on robustness.

References for Multiple Comparisons of Means and Variances

- [1] J.C. Hsu, J.C. (1966). *Multiple Comparisons, Theory and methods*. Chapman & Hall.
- [2] Games, P. A., and J. F. Howell. 1976. "Pairwise Multiple Comparison Procedures with Unequal N's and/or Variances: A Monte Carlo Study". *Journal of Educational Statistics*, 1, 113–125.
- [3] Levene, H. (1960). In *Contributions to Probability and Statistics: Essays in Honor of Harold Hotelling*, I. Olkin et al. eds., Stanford University Press, pp. 278-292.
- [4] Brown, M.B., and Forsythe, A.B. (1974). "Robust Tests for Equality of Variances." *Journal of the American Statistical Association* 69:364–367.
- [5] Nelson, P.R., Wludyka, P.S. and Karen A. F. Copeland, K.A. (2005) *The Analysis of Means: A Graphical Method for Comparing Means, Rates, and Proportions* ASA-SIAM Series on Statistics and Applied Probability, SIAM, Philadelphia, ASA, Alexandria, VA.

Analysis of Means (ANOM) Charts

A comparison of treatment means to determine which are significantly different from the overall average is commonly required. For example an educational professional might ask which schools in the system have a higher or lower rate of graduation compared to the average for all schools.

ANOM (Analysis of Means) addresses this problem for a fixed effects model. ANOM identifies the means that are significantly different from the overall mean. Analysis of means is often used in process control, since it has a graphical representation that is similar to a Shewhart chart. Typically, a central line represents the overall average. The treatment means, plotted as deviations from the overall average are compared with upper and lower decision limits to identify which are significantly different from the overall mean. The underlying assumptions are that the response belongs to a given distribution, and that the design is one way or two way.

The lower and upper decision levels (LDL, UDL) are used to ascertain whether the means at the factor level are different from the grand mean. A factor level mean that lies outside the UDL and LDL is statistically different from the grand mean at the specified confidence level.

Let $n = \sum_j n_j$ is the total number of observations, x_{ij} be the i^{th} observation at factor level j , and the mean for factor level j be:

$$\bar{x}_j = \frac{\sum_{i=1}^{n_j} x_{ij}}{n_j}$$

The grand mean is:

$$\bar{x} = \frac{\sum_{i,j} \bar{x}_{ij}}{n}$$

ANOM Normal One-Way

For a single factor with k factor levels, and with the response assumed normal, the LDL and UDL are given by:

$$LDL_j = \bar{x} - h_\alpha s_p w_j$$

$$UDL_j = \bar{x} + h_\alpha s_p w_j$$

where the pooled variance is:

$$s_p^2 = \sum_{j=1}^k (n_j - 1) s_j^2 / (n - k)$$

where s_j^2 is the variance of factor j .

where the weighting factor for each group j is:

$$w_j = \sqrt{(n - n_j) / (n n_j)}$$

which reduces to $w_j = \sqrt{(k - 1) / n}$ in the balanced case.

where h_α is the absolute value of the inverse CDF of the students t distribution at probability α^* with $df = n - k$ degrees of freedom, or is determined from lookup tables [4] (Table B.1 for balanced, Table B.3 for unbalanced) at the given probability level (α), and where $\alpha^* = (1 - (1 - \alpha)^{1/k}) / 2$ if $k > 2$ and $\alpha^* = \alpha / 2$ if $k = 2$.

ANOM Normal Two-Way

For a two factor model, with k_1 and k_2 factor levels respectively, and with the response assumed normal, the main effect is given by:

$$LDL_{fj} = \bar{x} - h_\alpha \sqrt{MSE w_{fj}}$$

$$UDL_{fj} = \bar{x} + h_\alpha \sqrt{MSE w_{fj}}$$

where LDL_{fj} and UDL_{fj} are the lower/upper bounds for factor f at level j respectively, where MSE is derived from an ANOVA analysis with a two factor model with cross effects, where the weighting factor for each group j is:

$$w_{fj} = (k_f - 1) / (k_f n_j)$$

which reduces to $w = (k - 1) / n$ in the balanced case.

where h_α is the absolute value of the inverse CDF of the students t distribution at probability α^* with $df = k_1 k_2 (n - 1)$ degrees of freedom, or is determined from Nelson's lookup tables at the given probability level (α),

and where $\alpha^* = (1 - (1 - \alpha)^{1/k_f}) / 2$ if $k > 2$ and $\alpha^* = \alpha / 2$ if $k = 2$.

The significance of the interaction effect is evaluated using a standard F test:

$$f = \frac{(S_r - S_u) / (f_r - f_u)}{S_u / f_u}$$

where S_u is the unrestricted sum of squares,
 where S_r is the restricted sum of squares,
 and where f_u and f_r are the degrees of freedom for the unrestricted and restricted cases respectively.

f is distributed F with $f_r - f_u$ and f_u degrees of freedom.

BY method

A straight forward method of analyzing a two factor model with k_1 and k_2 groups for factor 1 and 2 respectively is to simply evaluate the one factor model with factor 1 for each value of factor 2. Thus a total of k_2 one way models runs are undertaken.

SLICE method

This method [Wludyka 2013, 2015] is the same as the BY method, but uses the pooled MSE and degrees of freedom from the full two factor model. The SLICE factor with minimum number of levels is automatically selected.

ANOM Binomial Proportions One-Way

A binomial distribution can be used for data that consists of proportions. Thus $p_j = \sum_i^n x_{ij}/n_j$.

For a single factor with k factor levels, and $N = \sum_j n_j$ observations, and with the response assumed binomial, the LDL and UDL are given by:

$$LDL_j = \bar{p} - h_\alpha s_p w_j$$

$$UDL_j = \bar{p} + h_\alpha s_p w_j$$

where the average proportion is given by:

$$\bar{p} = \frac{\sum_j n_j p_j}{N}$$

where the pooled variance is:

$$s_{p_j}^2 = \bar{p}(1 - \bar{p})/n_j$$

where the weighting factor for each group j is:

$$w_j = \sqrt{(N - n_j)/(N)}$$

and where h_α is the absolute value of the inverse CDF of the normal distribution at probability $\alpha^* = 1 - \alpha/(2k)$ if $k > 2$, and $\alpha^* = 1 - \alpha/2$ if $k = 2$, or is determined from Nelson's lookup tables at the given probability level (α).

ANOM Binomial Proportions Two-Way

For a balanced two factor model, with k_1 and k_2 factor levels respectively, and with the response assumed binomial, the main effect is given by the one way factorial model, evaluated separately on each factor. Let k_1 be the number of levels of factor 1, with n observation per group.¹

$$LDL = \bar{p} - h_\alpha s_p w$$

$$UDL = \bar{p} + h_\alpha s_p w$$

where LDL and UDL are the lower/upper bounds respectively, where the average proportion is given by:

$$\bar{p} = \frac{\sum_j p_j}{k_1}$$

where the pooled variance is:

$$s_p^2 = \bar{p}(1 - \bar{p})/n$$

where the weighting factor is:

$$w = (k_1 - 1)/(k_1)$$

and where h_α is the absolute value of the inverse CDF of the normal distribution at probability $\alpha^* = 1 - \alpha/(2k)$ if $k > 2$, and $\alpha^* = 1 - \alpha/2$ if $k_1 = 2$, or is determined from Nelson's lookup tables at the given probability level (α).

The significance of the interaction effect is evaluated using a likelihood ratio test:

$$q = -2(\text{llf}_u - \text{llf}_r)$$

where llf_u is the unrestricted log likelihood,

where llf_r is the restricted log likelihood

and where the likelihood is specified as:

$$\text{llf} = \sum n (p \ln(\hat{p}) + (1 - p) \ln(1 - \hat{p}))$$

¹ A similar analysis follows for factor 2.

where n is the vector of sample size, p is the vector of probabilities, and \hat{p} is the predicted probability given by $\exp(x\beta)/(1 + \exp(x\beta))$. x is the factor matrix, and β is the maximum likelihood coefficient vector.

q is distributed χ^2 with degrees of freedom equal to the number of restrictions.

ANOM Poisson Rates One-Way

A Poisson distribution can be used for data that consists of counts. Thus the average count for group j is $c_j = \sum_i^n c_{ij}/n_j$.

For a single factor with k factor levels, and $N = \sum_j n_j$ observations, and with the response assumed distributed Poisson, the LDL and UDL are given by:

$$LDL_j = \bar{c} - h_\alpha s_p w_j$$

$$UDL_j = \bar{c} + h_\alpha s_p w_j$$

where the average count is given by:

$$\bar{c} = \frac{\sum_j n_j c_j}{N}$$

where the pooled variance is:

$$s_p^2 = \bar{c}$$

where the weighting factor for each group j is:

$$w_j = \sqrt{(N - n_j)/(N n_j)}$$

and where h_α is the absolute value of the inverse CDF of the normal distribution at probability $\alpha^* = 1 - \alpha/(2k)$ if $k > 2$, and $\alpha^* = 1 - \alpha/2$ if $k = 2$, or is determined from Nelson's lookup tables at the given probability level (α).

ANOM Poisson Rates Two-Way

For a balanced two factor model, with k_1 and k_2 factor levels respectively, and with the response assumed Poisson, the main effect is given by the one way factorial model, evaluated separately on each factor.

The significance of the interaction effect is evaluated using a likelihood ratio test:

$$q = -2(\text{llf}_u - \text{llf}_r)$$

where llf_u is the unrestricted log likelihood,

where llf_r is the restricted log likelihood

and where the likelihood is specified as:

$$\text{llf} = \sum y \ln(\lambda) - \lambda - \ln(y!)$$

where y is the vector of counts, $\lambda = \exp(x\beta + \ln(n))$, n is the vector of sample size, x is the factor matrix, and β is the maximum likelihood coefficient vector.

q is distributed χ^2 with degrees of freedom equal to the number of restrictions.

ANOM Nonparametric Transformed Ranks

When the distribution of a data set is not known, or cannot be assumed to be normal, Poisson or binomial, then a non-parametric methodology is used to compare means. Only a single factor (one way) is permitted.

The rank (r) of each element in the data set is derived. Then the rank is transformed using inverse normal scores:

$$x = \Phi^{-1}\left(\frac{0.5+r}{2N-1}\right)$$

where N is the number of observations, and Φ^{-1} is the inverse of the standard normal distribution function. x is then distributed normally, and the usual ANOM procedure can be applied.

ANOM Variances

ANOMV can be used to test whether the variance for any sample is significantly different from the pooled variance - it is thus a test of homogeneity. Typically, consistent output is required from a set of production lines, and ANOMV can be used to identify which line has variability significantly different from the average.

Let $N = \sum_j n_j$ be the total number of observations, x_{ij} be the i^{th} observation at factor level j , and the sample variance for factor level j be s_j^2 .

One Way balanced

For a single factor with k factor levels, n observations for each group, and with the response assumed normal, the LDL, CL, and UDL are given by:

$$LDL = L_{\alpha,k,n-1} k MS_e$$

$$CL = MS_e$$

$$UDL = U_{\alpha,k,n-1} k MS_e$$

where the pooled MSE is:

$$MS_e = \sum_{j=1}^k s_j^2$$

and where $L_{\alpha,k,n-1}$ and $U_{\alpha,k,n-1}$ are determined from Nelson's lookup tables at the given probability level (α) and k and $n-1$ degrees of freedom.

One Way unbalanced

For a single factor with k factor levels, n_j observations for group j , and with the response assumed normal, the LDL, CL, and UDL are given by:

$$LDL_j = L_{\alpha,k,n_j}^* \frac{N-k}{n_j-1} MS_e$$

$$CL_j = MS_e$$

$$UDL_j = U_{\alpha,k,n_j}^* \frac{N-k}{n_j-1} MS_e$$

where the pooled MSE is:

$$MS_e = \frac{\sum_{j=1}^k (n_j - 1) s_j^2}{N - k}$$

where L_{α,k,n_j}^* is given by:

$$L_{\alpha,k,n_j}^* = \beta^{-1}(\alpha^*, \theta_j, \phi_j)$$

where U_{α,k,n_j}^* is given by:

$$U_{\alpha,k,n_j}^* = \beta^{-1}(1 - \alpha^*, \theta_j, \phi_j)$$

and where $\alpha^* = (1 - \alpha)/(2k)$, $\theta_j = (n_j - 1)/2$, $\phi_j = \sum_{i \neq j} \theta_i$, and β^{-1} is the inverse of the cdf of the beta function.

Large Sample - Balanced

For data sets with large samples - all groups have at least 35 observations, for a single factor with k factor levels, n observations for each group, and with the response assumed normal, the LDL, CL,

and UDL are given by:

$$LDL = MS_e - h_{\alpha,k,\infty} MS_e w$$

$$CL = MS_e$$

$$UDL = MS_e + h_{\alpha,k,\infty} MS_e w$$

where the pooled MSE is:

$$MS_e = \sum_{j=1}^k s_j^2$$

where w is given by:

$$\sqrt{\frac{k-1}{k} \left(\frac{2}{n-1} + \frac{\gamma}{n} \right)}$$

where $h_{\alpha,k,\infty}$ is determined from Nelson's lookup tables at the given probability level α and k and ∞ degrees of freedom.

and where γ is an estimate of the common kurtosis - for samples distributed normal, this will be zero..

A similar formula is used for the large sample unbalanced case - see [4] for details.

ANOM Levene Robust Variances

When the distribution of a data set is not known, or cannot be assumed to be normal, then a non-parametric methodology is used to compare variances. The methodology is to first transform the data to the absolute deviation from the median for each group ADM_{ij} :

$$ADM_{ij} = |x_{ij} - md_j|$$

where md_j is the median for group j . Then, the ANOM Normal process is used on the ADM.

References for Analysis of Means (ANOM) Charts

- [1] Bakir, S.T. (2010) "A Nonparametric Test For Homogeneity Of Variances: Application To GPAs Of Students Across Academic Majors" *American Journal of Business Education*, 3(3) pp. 47-54.
- [2] Garrett, L and Nash, J.C. (2001). "Issues in Teaching the Comparison of Variability to Non-Statistics Students", *Journal of Statistics Education* ((2), pp. 1-16.
- [3] Nelson, L.S. (1983). "Exact Critical Values for Use with the Analysis of Means", *Journal of Quality Technology*, 15, pp. 40-44.

- [4] Nelson, P.R., Wludyka, P.S. and Karen A. F. Copeland, K.A. (2005) *The Analysis of Means: A Graphical Method for Comparing Means, Rates, and Proportions* ASA-SIAM Series on Statistics and Applied Probability, SIAM, Philadelphia, ASA, Alexandria, VA.
- [5] Wludyka, P.S., (2013). "Analyzing Multiway Models with ANOM Slicing" *SESUG Working Paper SD-06* 21st Southeast SAS users Group Conference.
- [6] Wludyka, P.S., (2015). "Using ANOM Slicing for Multi-Way Models with Significant Interaction", *Journal of Quality Technology* 47(2), pp. 193-203.
- [7] Wludyka, P.S. and Noguera, J.G. (2016, November). "Using Anom Slicing For Multiway Models With Binomial Or Poisson Data," Poster session presented at INFORMS Annual Meeting, Nashville, TN.
- [8] Wludyka, P.S., Bergen, A. & Liu, X. (2002) "A Monte Carlo Study of Two Normal Based Homogeneity of Variance Tests for Unbalanced Designs", *Joint Statistical Meetings - Statistical Computing Section*, pp. 3758-3759.
- [9] Wludyka, P.S. & Nelson, P.R. (1997). "An Analysis-of-Means-Type Test for Variances from Normal Populations", *Technometrics*, 39(3), pp. 274-285
- [10] Wludyka, P.S. & Nelson, P.R. (1999) "Two non-parametric, analysis-of-means-type tests for homogeneity of variances", *Journal of Applied Statistics*, 26:2, pp. 243-256,

Time Series Forecasting and Control Charts for Autocorrelated Data

Autocorrelation (ACF), Partial Autocorrelation (PACF) and Cross Correlation (CCF)

Just as correlation measures the extent of a linear relationship between two variables, autocorrelation measures the linear relationship between lagged values of data [9].

The autocorrelation coefficient (r_k) at lag k for a data set y is given by:

$$r_k = \frac{\sum_{t=k+1}^T (y_t - \bar{y})(y_{t-k} - \bar{y})}{\sum_{t=1}^T (y_t - \bar{y})^2}$$

where T is length of the time series (i.e., number of observations). This is known as the Autocorrelation Function (ACF) [9].

The maximum lag (k_{max}) may be user provided; the default value is the maximum of $10\log(T)$ and $3 * \text{seasonal frequency}$, but not exceeding $T/3$.

Partial autocorrelation is similar to autocorrelation but adjusts for correlation inherent in lags, e.g., y_t and y_{t-2} might be correlated simply because they are both connected to y_{t-1} , rather than because of any new information contained in y_{t-2} [9].

Partial autocorrelation is the autocorrelation between y_t and y_{t+k} after removing any linear dependence on $y_{t+1}, y_{t+2}, \dots, y_{t+k-1}$; thus it measures the correlation between y_t and y_{t+k} not accounted for by an Autoregressive AR($k-1$) process. Each partial autocorrelation can be estimated as the last coefficient in an autoregressive model. Specifically, α_k , the k th partial autocorrelation coefficient, is equal to the estimate of ϕ_k in an AR(k) model [9].

The partial autocorrelation coefficient (ϕ_k) at lag $k, k > 1$ is given by:

$$\begin{aligned}\hat{y}_{t+k} &= \beta_1 y_{t+k-1} + \beta_2 y_{t+k-2} + \dots + \beta_{k-1} y_{t+1} \\ \hat{y}_t &= \beta_1 y_{t+1} + \beta_2 y_{t+2} + \dots + \beta_{k-1} y_{t+k-1} \\ \phi_k &= \Phi(y_{t+k} - \hat{y}_{t+k}, y_t - \hat{y}_t)\end{aligned}$$

where Φ is the correlation function.

If there are missing values, the autocorrelation and partial autocorrelation coefficients are evaluated by skipping correlations that include missing values and adjusting the sample count.

Under the null, the sample autocorrelation and partial autocorrelation coefficients are distributed approximately normal, with a mean of zero, and a variance of $1/n$; thus, confidence intervals can be used to test the null hypothesis of no correlation. The bargraph of the sequence of autocorrelations with confidence intervals is called the Autocorrelation Function (ACF) Plot (sometimes also called a Correlogram). The bargraph of the sequence of partial autocorrelations is called the Partial Autocorrelation Function (PACF) Plot.

Cross Correlation is the standard method of estimating the degree to which two series are correlated. The bargraph of the sequence of cross correlations is called the Cross Correlation Function (CCF) Plot. The cross correlation coefficient ($r_{x,y}$) for data sets x and y is given by:

$$r_{x,y} = \frac{1}{n\sigma_x\sigma_y} \sum_{t=1}^n (x_t - \bar{x})(y_t - \bar{y})$$

where n is the number of observations, \bar{x} and \bar{y} are the means, and σ_x and σ_y are the standard errors of x and y respectively. The cross-correlation coefficient can also be estimated between lagged values of x and y ; for example, between y_{t+k} and x_t . The cross-correlation coefficient ranges in value between +1 (perfect positive correlation), and -1 (perfect negative correlation).

If x is autocorrelated, the CCF is affected by its time series structure and any “in common” trends the x and y series may have over time. Pre-whitening solves this problem by removing the autocorrelation and trends. The Pre-Whiten Data option does the following:

1. Automatic ARIMA is applied to the x input; residuals are stored.
2. The y data is filtered using the same model and parameter values; residuals are stored.
3. A CCF is then created using the x and y residuals.

For further details, see “Pre-whitening as an Aid to Interpreting the CCF”

<https://newonlinecourses.science.psu.edu/stat510/lesson/9/9.1>.

Ljung-Box Test

In addition to looking at the ACF plot, we can also do a more formal test for autocorrelation by considering a whole set of r_k values as a group, rather than treating each one separately. This is the Ljung-Box Q Test:

$$Q = T(T + 2) \sum_{k=1}^h (T - k)^{-1} r_k^2,$$

where h is the maximum lag being considered and T is the number of observations.

If the autocorrelations did come from a white noise series, then Q would have a χ^2 distribution with $(h - k)$ degrees of freedom, where k is the number of parameters in the model [9]. This is also known as a Portmanteau test.

Box-Cox Transformation

The Box-Cox procedure transforms the vector x as follows:

$$y = \frac{(x - \tau)^\lambda - 1}{\lambda}$$

where $\tau < x_{min}$ is a threshold parameter, such that $x - \tau$ is always positive. The object of this procedure is to transform non-normal data into data that is closer to normal using Sequential Quadratic Programming (SQP) to maximize a likelihood function. The likelihood includes trend and seasonality dummy variables.

If λ is zero, then $y = \ln(x)$. λ is constrained to be ≥ 0 and ≤ 1 . If Rounded Lambda is selected the choices for λ are $\ln(x)$ or \sqrt{x} . This conservative approach was found to give the best results with competition data. The user may also specify any value for λ and threshold.

If an invalid threshold is specified, a threshold value is specified as $x_{min} - \delta$, where $\delta = \min(0.001, 0.001(x_{max} - x_{min}))$.

If a threshold parameter is not specified, and x has negative values, then τ is estimated along with λ ; these parameters are selected to maximize the constrained likelihood using SQP. The constraint is used to ensure that $x - \tau$ is always positive. Trend and seasonal dummy variables are not included in the likelihood for this case.

In both ARIMA and Exponential Smoothing, the user has the option of invoking a Box-Cox transform. Forecast values are specified using the inverse transform.

Seasonal Frequency

Seasonal Frequency is the number of observations per “cycle” unit of time. Some examples include:

- Quarterly data, seasonal frequency = 4 observations per year
- Daily data, seasonal frequency = 7 observations per week
- Monthly data, seasonal frequency = 12 observations per year
- Hourly data, seasonal frequency = 24 observations per day
- Half-hourly data, seasonal frequency = 48 observations per day
- Weekly data, seasonal frequency = 52 observations per year
- Hourly data, seasonal frequency = 168 observations per week
- Half-hourly data, seasonal frequency = 336 observations per week
- Daily data, seasonal frequency = 365 observations per year (excluding leap years).

Seasonal decomposition should be used for specified frequencies higher than 24.

SigmaXL follows the convention used in R time series objects (ts) and the R forecast package [12], but note that it is the opposite of the definition of frequency in physics, or in engineering Fourier analysis, where “period” is the length of the cycle, and “frequency” is the inverse of period (see

<https://robjhyndman.com/hyndsight/seasonal-periods/>).

Seasonal Trend Decomposition

A time series can be decomposed into three components:

$$y_t = T_t + S_t + I_t$$

where T_t is the trend and cyclical component, S_t is the seasonal component, and I_t is the remainder, the irregular or noise component. STL is a versatile and robust method for decomposing time series into these three components. STL is an acronym for “Seasonal and Trend decomposition using Loess”. Loess is locally weighted regression smoothing. The STL method was developed by Cleveland, et al. [3]. Note that it is strictly additive, but one can simply specify a Box-Cox Transformation with Lambda = 0 (Ln) to obtain an equivalent to multiplicative decomposition. Unlike the use of Box-Cox in forecasting, the results remain transformed, i.e., the inverse transformation is not applied to the decomposed data – the additive model is retained.

If seasonal frequency is not specified, the trend is computed using Freidman’s “Supersmoother” [4], which uses local linear regression with adaptive bandwidths. The basic idea is to first estimate a number of fixed bandwidth smooths by local linear regression. The leave-one-out cross-validated residuals from each of those initial estimates are then smoothed using a constant bandwidth. A set of smoothed bandwidths are selected and the smoothed outcomes are linearly interpolated and smoothed with a fixed bandwidth.

Multiple Seasonal Decomposition (MSD) uses STL iteratively and was developed by Hyndman, et al. [12].

Spectral Density and Automatic Detection of Seasonal Frequency

Spectral Density determines the seasonal frequency of a time series using spectral analysis with Fast Fourier Transforms, using a specified uniform window (see https://en.wikipedia.org/wiki/Window_function).

The Fisher Kappa white noise test is automatically applied (see Fuller [5], Chapter 7) and, if significant (p-value < .05), returns the rounded frequency with the largest spectral power value. If not significant, a frequency of unity is returned (nonseasonal).

A further decomposition based seasonal strength test is applied to the rounded frequency (<https://otexts.com/fpp2/seasonal-strength.html>) and if the statistic is $\geq .5$, the above frequency value is used, otherwise it is set to 1. This helps to distinguish between cyclical and seasonal time series.

For the Spectral Density Plot or MSD option, if the seasonal frequency is > 1 , the above process is

repeated two more times with the dominant frequency removed by seasonal decomposition, spectral analysis applied to the remainder and seasonal strength test applied to the new dominant frequency.

Seasonally Adjusted Linear Interpolation of Missing Values

Linear interpolation is carried out by first estimating missing values using a regression model with trend and dummy coding for seasonal. The time series is then decomposed into three components - trend, seasonal and remainder - using STL decomposition for seasonal data. A data point is estimated by linear interpolation of the nonseasonal components of the previous and subsequent values, and then adding the seasonal component for the observation in question [12, 17]. For nonseasonal data, linear interpolation is utilized.

Once a model has been specified, missing values are estimated using one step ahead forecasting.

Information Criteria for Model Comparison

Parameters are estimated by maximizing the Log-Likelihood function (which is similar to minimizing the residual sum-of-squares). Given a set of candidate models for the data, the preferred model is the one with the minimum Information Criteria value. The Information Criteria rewards goodness of fit (as assessed by the likelihood function), but it also includes a penalty that is an increasing function of the number of estimated parameters. The penalty discourages overfitting, because increasing the number of parameters in the model almost always improves the goodness of the fit (see https://en.wikipedia.org/wiki/Akaike_information_criterion).

Akaike's Information Criterion (AIC) is:

$$AIC = -2\ln(L) + 2(k + 1)$$

where L is the likelihood of the model and k is the total number of parameters (and initial states for exponential smoothing) that have been estimated.

The AIC corrected for small sample bias (AICc) is defined as:

$$AIC_c = AIC + \frac{2(k + 1)(k + 2)}{T - k - 2}$$

The Bayesian Information Criterion (BIC) is:

$$BIC = AIC + (k + 1)[\ln(T) - 2]$$

AICc is recommended as the default Information Criterion, based on forecast error performance with M3 competition data [12].

Notes on AIC from Rob Hyndman (see <https://robjhyndman.com/hyndsight/aic/>):

- You can compare very different models with AIC, however, make sure the likelihoods are computed on the same data. For example, you cannot compare an ARIMA model with differencing to an ARIMA model without differencing, because you lose one or more observations via differencing.
- For a similar reason, you cannot compare the AIC from an Exponential Smoothing model with the AIC from an ARIMA model. The likelihood function calculations are different and the two models treat initial values differently. For example, after differencing, an ARIMA model is computed on fewer observations, whereas an Exponential Smoothing model is always computed on the full set of data. Even when the models are equivalent (e.g., an ARIMA(0,1,1) and an ETS(A,N,N)), the AIC values will be different.

Forecast Accuracy

As noted above, Information Criteria AIC/AICc cannot be used to compare ARIMA and ETS models or ARIMA models with different d , D values, but forecast accuracy metrics can be compared. A forecast error e_t is the difference between an observed value and its forecast. It is not the same as a model residual because residuals may be in different units due to Box-Cox transformation or a Multiplicative model.

Types of forecast error include:

- In-Sample One-Step-Ahead Forecast. This is less useful because the model may be over-fitted.
- Out-of-Sample (Withhold) One-Step-Ahead. Model parameter estimates do not use any withhold data, but the forecast updates with every new withhold observation.
- Out-of-Sample (Withhold) Multi-Step-Ahead. This is important if one is assessing forecast accuracy over a horizon. This is used in forecast competitions.

Forecast accuracy measures used in SigmaXL include:

$$\text{Root mean squared error: RMSE} = \sqrt{\text{mean}(e_t^2)}$$

$$\text{Mean absolute error: MAE} = \text{mean}(|e_t|)$$

$$\text{Mean absolute percentage error: MAPE} = \text{mean}\left(\left|\frac{100e_t}{y_t}\right|\right)$$

$$\text{Mean absolute scaled error: MASE} = \text{mean}(|e_t|)/\text{scale}$$

where *scale* is the MAE of the in-sample naïve or seasonal naïve forecast (set all forecasts to be the value of the last observation/period). A scaled error is less than one if it arises from a better forecast than the average naïve/seasonal naïve forecast. Conversely, it is greater than one if the forecast is worse than the average naïve forecast [9].

Exponential Smoothing - ETS

Simple Exponential Smoothing:

Simple Exponential Smoothing forecasts are calculated using weighted averages, where the weights decrease exponentially as observations come from further in the past with the smallest weights associated with the oldest observations:

$$\hat{y}_{t+1} = \alpha y_t + \alpha(1 - \alpha) y_{t-1} + \alpha(1 - \alpha)^2 y_{t-2} + \dots$$

where $0 \leq \alpha \leq 1$ is the smoothing parameter.

An equivalent formulation for simple exponential smoothing is:

$$\hat{y}_{t+1} = \alpha y_t + (1 - \alpha)\hat{y}_t$$

with the starting forecast value (initial level) \hat{y}_1 to be estimated and denoted as level l_0 .

Hyndman gives an alternative representation known as the component form [9]. For simple exponential smoothing, the only component included is the level l_t . Other methods may also include a trend b_t and/or a seasonal component s_t .

Component form representations of exponential smoothing methods comprise a forecast equation and a smoothing equation for each of the components included in the method. The component form of simple exponential smoothing is given by:

$$\begin{array}{ll} \text{Forecast equation} & \hat{y}_{t+h|t} = l_t \\ \text{Smoothing equation} & l_t = \alpha y_t + (1 - \alpha)l_{t-1}, \end{array}$$

where l_t is the level (or the smoothed value) of the series at time t . $\hat{y}_{t+h|t}$ denotes an h multi-step forecast, taking account of all observations up to time t . Setting $h = 1$, a one step ahead forecast gives the in-sample fitted values.

The “error correction” form is:

$$\begin{aligned} l_t &= l_{t-1} + \alpha(y_t - l_{t-1}) \\ &= l_{t-1} + \alpha e_t \end{aligned}$$

where $e_t = y_t - l_{t-1} = y_t - \hat{y}_{t|t-1}$ is the residual at time t . The equations of the additive error model can then be written as:

$$\begin{aligned} y_t &= l_{t-1} + \varepsilon_t \\ l_t &= l_{t-1} + \alpha \varepsilon_t. \end{aligned}$$

Multiplicative errors are one-step-ahead relative errors:

$$\varepsilon_t = \frac{y_t - \hat{y}_{t|t-1}}{\hat{y}_{t|t-1}}$$

The equations of the multiplicative error model can be written as:

$$\begin{aligned} y_t &= \ell_{t-1}(1 + \varepsilon_t) \\ \ell_t &= \ell_{t-1}(1 + \alpha\varepsilon_t). \end{aligned}$$

Error, Trend, Seasonal (ETS) Models:

Error, Trend, Seasonal (ETS) models expand on simple exponential smoothing to accommodate trend and seasonal components as well as additive or multiplicative errors. Simple Exponential Smoothing is an Error Model. Error, Trend model is Holt's Linear, also known as double exponential smoothing. Error, Trend, Seasonal model is Holt-Winters, also known as triple exponential smoothing.

Rob Hyndman has developed a complete taxonomy that describes all of the combinations of exponential smooth models in a consistent manner [9]:

- Error:
 - Additive or Multiplicative
 - The point forecasts produced by the models are identical if they use the same smoothing parameter values. Multiplicative will, however, generate different prediction intervals to accommodate change in variance.
 - An alternative to multiplicative is to use the Ln transformation (Box-Cox transformation with $\lambda = 0$).
 - Error models include the smoothing parameter α and initial level value.

- Trend:
 - None, Additive, Additive Damped
 - Multiplicative Trend is not recommended as it tends to produce poor forecasts [9], so is not included in SigmaXL
 - Trend models add a smoothing parameter β and initial trend value.
 - Damped trend models add a smoothing parameter ϕ that "dampens" the trend to a flat line some time in the future.

- Seasonal:
 - None, Additive, Multiplicative
 - The additive method is preferred when the seasonal variations are roughly constant through the series, while the multiplicative method is preferred when the seasonal variations are changing proportional to the level of the series.
 - Seasonal models add a smoothing parameter γ and initial seasonal values.
 - # of initial values estimated = seasonal frequency – 1
 - constrained to sum to 0 for additive or seasonal frequency m for multiplicative

Summary of ETS Models in SigmaXL:

Short hand (Error, Trend, Seasonal)	Method
(A, N, N)	Simple Exponential Smoothing with Additive Errors - Exponentially Weighted Moving Average (EWMA)
(M, N, N)	Simple Exponential Smoothing with Multiplicative Errors
(A, A, N)	Additive Trend Method with Additive Errors (Holt's Linear)
(M, A, N)	Additive Trend Method with Multiplicative Errors (Holt's Linear)
(A, A, A)	Additive Trend, Additive Seasonal Method with Additive Errors (Holt-Winters)
(M, A, A)	Additive Trend, Additive Seasonal Method with Multiplicative Errors (Holt-Winters)
(A, N, A)	Additive Seasonal Method with Additive Errors
(M, N, A)	Additive Seasonal Method with Multiplicative Errors
(A, Ad, N)	Additive Damped Trend Method with Additive Errors
(M, Ad, N)	Additive Damped Trend Method with Multiplicative Errors
(A, Ad, A)	Additive Damped Trend, Additive Seasonal Method with Additive Errors
(M, Ad, A)	Additive Damped Trend, Additive Seasonal Method with Multiplicative Errors
(M, A, M)	Additive Trend, Multiplicative Seasonal Method with Multiplicative Errors (Holt-Winters)
(M, N, M)	Multiplicative Seasonal Method with Multiplicative Errors
(M, Ad, M)	Additive Damped Trend, Multiplicative Seasonal Method with Multiplicative Errors (Holt-Winters)

ETS Model Equations:

The equations for each of the models (also known as the “innovations state space model”) are given as [9]:

ADDITIVE ERROR MODELS

Trend	Seasonal		
	N	A	M
N	$y_t = \ell_{t-1} + \varepsilon_t$ $\ell_t = \ell_{t-1} + \alpha\varepsilon_t$	$y_t = \ell_{t-1} + s_{t-m} + \varepsilon_t$ $\ell_t = \ell_{t-1} + \alpha\varepsilon_t$ $s_t = s_{t-m} + \gamma\varepsilon_t$	$y_t = \ell_{t-1}s_{t-m} + \varepsilon_t$ $\ell_t = \ell_{t-1} + \alpha\varepsilon_t/s_{t-m}$ $s_t = s_{t-m} + \gamma\varepsilon_t/\ell_{t-1}$
A	$y_t = \ell_{t-1} + b_{t-1} + \varepsilon_t$ $\ell_t = \ell_{t-1} + b_{t-1} + \alpha\varepsilon_t$ $b_t = b_{t-1} + \beta\varepsilon_t$	$y_t = \ell_{t-1} + b_{t-1} + s_{t-m} + \varepsilon_t$ $\ell_t = \ell_{t-1} + b_{t-1} + \alpha\varepsilon_t$ $b_t = b_{t-1} + \beta\varepsilon_t$ $s_t = s_{t-m} + \gamma\varepsilon_t$	$y_t = (\ell_{t-1} + b_{t-1})s_{t-m} + \varepsilon_t$ $\ell_t = \ell_{t-1} + b_{t-1} + \alpha\varepsilon_t/s_{t-m}$ $b_t = b_{t-1} + \beta\varepsilon_t/s_{t-m}$ $s_t = s_{t-m} + \gamma\varepsilon_t/(\ell_{t-1} + b_{t-1})$
Ad	$y_t = \ell_{t-1} + \phi b_{t-1} + \varepsilon_t$ $\ell_t = \ell_{t-1} + \phi b_{t-1} + \alpha\varepsilon_t$ $b_t = \phi b_{t-1} + \beta\varepsilon_t$	$y_t = \ell_{t-1} + \phi b_{t-1} + s_{t-m} + \varepsilon_t$ $\ell_t = \ell_{t-1} + \phi b_{t-1} + \alpha\varepsilon_t$ $b_t = \phi b_{t-1} + \beta\varepsilon_t$ $s_t = s_{t-m} + \gamma\varepsilon_t$	$y_t = (\ell_{t-1} + \phi b_{t-1})s_{t-m} + \varepsilon_t$ $\ell_t = \ell_{t-1} + \phi b_{t-1} + \alpha\varepsilon_t/s_{t-m}$ $b_t = \phi b_{t-1} + \beta\varepsilon_t/s_{t-m}$ $s_t = s_{t-m} + \gamma\varepsilon_t/(\ell_{t-1} + \phi b_{t-1})$

MULTIPLICATIVE ERROR MODELS

Trend	Seasonal		
	N	A	M
N	$y_t = \ell_{t-1}(1 + \varepsilon_t)$ $\ell_t = \ell_{t-1}(1 + \alpha\varepsilon_t)$	$y_t = (\ell_{t-1} + s_{t-m})(1 + \varepsilon_t)$ $\ell_t = \ell_{t-1} + \alpha(\ell_{t-1} + s_{t-m})\varepsilon_t$ $s_t = s_{t-m} + \gamma(\ell_{t-1} + s_{t-m})\varepsilon_t$	$y_t = \ell_{t-1}s_{t-m}(1 + \varepsilon_t)$ $\ell_t = \ell_{t-1}(1 + \alpha\varepsilon_t)$ $s_t = s_{t-m}(1 + \gamma\varepsilon_t)$
A	$y_t = (\ell_{t-1} + b_{t-1})(1 + \varepsilon_t)$ $\ell_t = (\ell_{t-1} + b_{t-1})(1 + \alpha\varepsilon_t)$ $b_t = b_{t-1} + \beta(\ell_{t-1} + b_{t-1})\varepsilon_t$	$y_t = (\ell_{t-1} + b_{t-1} + s_{t-m})(1 + \varepsilon_t)$ $\ell_t = \ell_{t-1} + b_{t-1} + \alpha(\ell_{t-1} + b_{t-1} + s_{t-m})\varepsilon_t$ $b_t = b_{t-1} + \beta(\ell_{t-1} + b_{t-1} + s_{t-m})\varepsilon_t$ $s_t = s_{t-m} + \gamma(\ell_{t-1} + b_{t-1} + s_{t-m})\varepsilon_t$	$y_t = (\ell_{t-1} + b_{t-1})s_{t-m}(1 + \varepsilon_t)$ $\ell_t = (\ell_{t-1} + b_{t-1})(1 + \alpha\varepsilon_t)$ $b_t = b_{t-1} + \beta(\ell_{t-1} + b_{t-1})\varepsilon_t$ $s_t = s_{t-m}(1 + \gamma\varepsilon_t)$
Ad	$y_t = (\ell_{t-1} + \phi b_{t-1})(1 + \varepsilon_t)$ $\ell_t = (\ell_{t-1} + \phi b_{t-1})(1 + \alpha\varepsilon_t)$ $b_t = \phi b_{t-1} + \beta(\ell_{t-1} + \phi b_{t-1})\varepsilon_t$	$y_t = (\ell_{t-1} + \phi b_{t-1} + s_{t-m})(1 + \varepsilon_t)$ $\ell_t = \ell_{t-1} + \phi b_{t-1} + \alpha(\ell_{t-1} + \phi b_{t-1} + s_{t-m})\varepsilon_t$ $b_t = \phi b_{t-1} + \beta(\ell_{t-1} + \phi b_{t-1} + s_{t-m})\varepsilon_t$ $s_t = s_{t-m} + \gamma(\ell_{t-1} + \phi b_{t-1} + s_{t-m})\varepsilon_t$	$y_t = (\ell_{t-1} + \phi b_{t-1})s_{t-m}(1 + \varepsilon_t)$ $\ell_t = (\ell_{t-1} + \phi b_{t-1})(1 + \alpha\varepsilon_t)$ $b_t = \phi b_{t-1} + \beta(\ell_{t-1} + \phi b_{t-1})\varepsilon_t$ $s_t = s_{t-m}(1 + \gamma\varepsilon_t)$

Automatic Model Selection (Identification):

The identification process is used to ascertain which ETS model to use. The best model is determined by estimation of all possible combinations of error, trend, seasonality and damping. AICc is recommended as the default Information Criteria, based on forecast error performance with M3 competition data.

Some of the model combinations lead to numerical instability and are not considered in the selection process: (A,N,M) (A,A,M) (A,Ad,M).

If a Box-Cox transformation is used, multiplicative models are not considered.

Estimation:

Given a specific ETS model, the values of the weights (α for error, β for trend, γ for season, and ϕ for damping) are estimated using exact constrained maximum likelihood in state space using recursion, employing the Nelder Mead simplex algorithm. Effectively, the objective is to minimize the residual sum of squares:

$$SSE = \sum_{t=1}^n (y_t - \hat{y}_t)^2$$

The optimization returns the weights (α for error, β for trend, γ for seasonality, and ϕ for damping), as well as initial state values. There are a number of constraints on the weights (in addition to the weights each lying between 0 and 1) that are required for stability [8, 10]. If missing values are encountered, seasonally adjusted linear interpolation is used. A Box-Cox transform is permitted prior to estimation, although such a transformation should only be used for an additive ETS model.

Covariates cannot be estimated as part of an integrated ETS estimation (see <https://robjhyndman.com/hyndsight/ets-regressors/> for a discussion on this issue).

A starting value of 0.5 is used for α , 0.05 for β , 0.025 for γ , and 0.978 for ϕ . For covariates, the initial values are taken from the OLS estimates of a regression of y against a constant and the scaled covariates.

Forecasting:

Once a model is specified, future values of the series can be evaluated easily. Given data y_1, \dots, y_n , then the conditional mean of y_{n+1} is derived using the point forecast. The process is repeated dynamically.

See <https://otexts.com/fpp2/ets-forecasting.html> for a brief discussion of ETS model prediction intervals. Full formula details are quite complex and given in Hyndman et al. [11].

Autoregressive Integrated Moving Average - ARIMA

The Autoregressive Integrated Moving Average (ARIMA) model was developed by Box and Jenkins [2]. The automatic determination of the best model order in SigmaXL uses the method of Rob Hyndman [9,10,12].

Stationarity and Differencing:

ARIMA assumes that the time series is stationary, i.e., it has the property that the mean, variance and autocorrelation structure do not change over time. If a time series mean is not stationary (e.g. trending), this can be corrected by differencing, computing the differences between consecutive observations for nonseasonal and between consecutive periods (e.g. months) for seasonal data (Jan 2019 – Jan 2018, etc.). For nonseasonal, this will typically involve 1 or 2 orders of differencing. This order is the Integrated term d . For seasonal, this will typically involve 1 order of differencing. This order is the Seasonal Integrated term D . A constant term c is optional:

- If $d+D = 0$, a constant term in the model is the mean.
- If $d+D = 1$, a constant term in the model is a trend (drift).
- If $d+D > 1$, a constant term would be a quadratic trend, so constant should not be included.
- It is recommended that $d+D$ should not be > 3 .
- If the variance is not stationary, use a Box-Cox transformation.

Autoregressive (AR) Model:

In an autoregressive model, we forecast the variable of interest using a linear combination of past values of the variable. The term autoregressive indicates that it is a regression of the variable against itself:

$$y_t = c + \phi_1 y_{t-1} + \phi_2 y_{t-2} + \dots + \phi_p y_{t-p} + \varepsilon_t$$

where ε_t is white noise. This is like a multiple regression but with lagged values of y_t as predictors. We refer to this as an AR(p) model, an autoregressive model of order p [9].

Moving Average (MA) Model:

Rather than using past values of the forecast variable in a regression, a moving average model uses past forecast errors in a regression-like model [9].

$$y_t = c + \varepsilon_t + \theta_1 \varepsilon_{t-1} + \theta_2 \varepsilon_{t-2} + \dots + \theta_q \varepsilon_{t-q}$$

We refer to this as an MA(q) model, a moving average model of order q [9].

Autoregressive Integrated Moving Average (ARIMA) Model:

If we combine differencing with autoregression and a moving average model, we obtain a

nonseasonal ARIMA model.

$$y'_t = c + \phi_1 y'_{t-1} + \dots + \phi_p y'_{t-p} + \theta_1 \varepsilon_{t-1} + \dots + \theta_q \varepsilon_{t-q} + \varepsilon_t$$

where y'_t is the differenced series. This is the ARIMA (p, d, q) model, where p is the number of autoregressive terms, d is the degree of differencing and q is the number of moving average terms.

To form more complicated models, it is much easier to work with the backshift notation [9].

$$\begin{array}{ccc} (1 - \phi_1 B - \dots - \phi_p B^p) & (1 - B)^d y_t & = & c + (1 + \theta_1 B + \dots + \theta_q B^q) \varepsilon_t \\ \uparrow & \uparrow & & \uparrow \\ \text{AR}(p) & d \text{ differences} & & \text{MA}(q) \end{array}$$

Seasonal ARIMA:

For seasonal, the model consists of terms that are similar to the nonseasonal components of the model, but involve backshifts of the seasonal period. The seasonal model is ARIMA (P, D, Q) and combined we have ARIMA $(p, d, q) (P, D, Q)$.

For example, an ARIMA(1,1,1)(1,1,1) model for quarterly data (seasonal frequency = 4), can be written as [9]:

$$(1 - \phi_1 B) (1 - \Phi_1 B^4)(1 - B)(1 - B^4)y_t = (1 + \theta_1 B) (1 + \Theta_1 B^4)\varepsilon_t.$$

Automatic Model Selection (Identification):

The identification process is required to determine the degree of differencing (d, D) necessary to generate a series that is stationary, and to then ascertain the order of the model (p, q, P, Q) . This can be done manually using the correlogram (ACF) and partial autocorrelogram (PACF) both before and after differencing. This process is automated in SigmaXL.

First, the degree of seasonal differencing (D) is determined with the Seasonal Strength (SEAS) test [12, 19] which uses the ratio of the seasonal component to the detrended component and compares this with the SEAS critical value. If the ratio exceeds the critical value, then this implies the rejection of the no unit root null hypothesis, implying $D = 1$, and the data are seasonally differenced.

The degree of differencing for the nonseasonal component (d) is then determined using the KPSS unit root test [12, 13]. The data is tested for a unit root, and if found the data is differenced, and the process continues; it stops when no significant unit root is found or $d = 2$. SigmaXL uses the Schwert truncation lag parameter = $\text{floor}(12*(n/100)^{.25})$ for $n \leq 70$ and Quadratic Spectral Kernel for $n > 70$ [21, 22]. If $d = 0$ and $n \leq 500$, a further test is applied comparing the standard deviation of the (0,0,0) with constant model against the standard deviation of (0,1,0) with constant.

If $d = 1$ produces a lower standard deviation value, then it is used as the order for nonseasonal differencing (see <https://people.duke.edu/~rnau/411arim2.htm>). While there is a potential risk of overdifferencing, this additional test results in dramatic improvement in forecast accuracy using M3 and M4 competition data. If the process is thought to be stationary and SigmaXL returns $d = 1$, the user should manually specify a $d = 0$ model and compare the two models.

Other tests for seasonal and nonseasonal stationarity (unit-roots) are available, but SEAS and KPSS are recommended by Hyndman based on forecast accuracy evaluation with M3 competition data (see <https://robjhyndman.com/hyndsight/forecast83/> and <https://robjhyndman.com/hyndsight/unit-root-tests/>). Currently, SigmaXL does not include alternative unit root/stationarity tests. The test statistic values for SEAS and KPSS are used internally but not reported.

Having ascertained the degree of seasonal and nonseasonal differencing, model selection implies the choice of the AR and MA components - p, q, P, Q - as well as whether a constant should or should not be part of the model. One method of choosing the best model is to exhaustively estimate all possible combinations, subject to upper bounds on the AR/MA terms. A more efficient approach is to estimate a series of ARIMA models using a stepwise ordering suggested by Hyndman and Khandakar [10]. The preferred model is the one with the lowest fit metric, the default being AICc. Only models that are stationary and invertible are considered.

Estimation:

Given a specific (p, d, q, P, D, Q) ARIMA model, the values of the AR and MA terms are estimated. Covariates (which are scaled automatically) are permitted, as well as a constant if there is no differencing, and a trend if d and D equals $\{ 0 \ 1 \}$ or $\{ 1 \ 0 \}$. Two general models are available - the conditional sum of squares (CSS) and the state space Kalman maximum likelihood. The likelihood for the CSS model is:

$$\log L = \sum_i [\log(2\pi)/2 - 0.5\log(\sigma^2) - e_i^2/2\sigma^2]$$

The CSS is always used for initial estimates and is used if $n > 500$ or frequency > 12 . The residuals used are derived from the differenced time series – for CSS, AR/MA filtering is used to generate the residuals, while for Kalman, which is a recursive estimator, the residuals are derived from a maximum likelihood AR/MA process. In each case, the maximum likelihood process uses the constrained Nelder Mead simplex algorithm. The constraints impose a maximum value of unity on AR/MA terms by using a tanh transformation, as well as using a unit root test to ensure that the AR terms are stationary, and the MA terms are invertible.

For CSS, if missing values are encountered, the largest contiguous range is used. For Kalman, missing values are evaluated by skipping the update phase in the state based model.

Confidence bounds on the coefficients are derived from the diagonal of the inverse of either the Hessian or the outer product of the gradient (OPG).

Covariates (predictors) are treated in the traditional ARIMAX model by adding them to the right-

hand side before differencing. In order to estimate an ARIMA model with covariates, the covariate terms are subtracted from the data before evaluating the residuals in each iteration of the evaluation of the likelihood. The parameters of the covariate terms are estimated concurrently with all the other parameters. For prediction and forecasting, the reverse process occurs, and the covariate terms are added back to the predicted/forecast values. Thus, the effect of the exogenous variables is always additive, irrespective as to the ARIMA model.

A starting value of 0.1 is used for each parameter, with the following exceptions: For the constant term, the mean is used if there is no differencing, and the mean of the differenced data, adjusted by the number of AR terms, otherwise. For covariates, the initial values are taken from the OLS estimates of an ARX model with m AR terms, scaled covariates and trend, where m is the seasonal frequency.

ARIMA Parameter Estimates include significance tests; P-Values < .05 are significant and highlighted in red. This is useful for model refinement with multiple predictors. Note that for AR/MA model order selection, minimum AICc should be used rather than significance tests [25].

In-Sample Fitted (One-Step Prediction)

The predicted values for an arima process are calculated using the Kalman process on the de-trended, differenced data, y_d . For the initial data point, the Kalman process uses a predicted value y_p of zero; subsequent predicted values are derived from the Kalman process. The calculated residuals are derived as the difference between y_d and y_p , and the predicted data, which are one step ahead predictions, are derived from actual data plus the calculated residual.

The predicted for an arima process are calculated using either the Kalman process or the CSS process on the de-trended, differenced data, y_d . For the initial data points, the Kalman process uses a predicted value y_p of zero; subsequent predicted values are derived from the Kalman process; these initial values are excluded under CSS. The calculated residuals are derived as the difference between y_d and y_p , and the predicted data, after removing trend. These are one-step-ahead predictions, and are derived from actual data plus the calculated residual. If differencing occurs, the first $d + D$ predicted values are set to missing.

Forecasting:

The raw coefficients are transformed into the ϕ^* and θ^* vectors, which can be used on the original time series. A forecast is undertaken from the end of the series for the number of forecast periods specified. Under the default, the forecasts are evaluated dynamically, such that the predicted values at any point provide the basis for the next period's prediction. A confidence band is produced for each forecast value, based on the standard deviation of the forecast error for observation $n + m$:

$$\sigma_{n+m} = \sqrt{\sigma_w^2 \sum_{j=0}^{m-1} \Psi_j^2}$$

where σ_w^2 is the residual variance derived from the n observations, and Ψ_j is the weight on the $t - j$ observation when the residual for the t observation is written as an infinite order MA model.

Assessing Forecast Accuracy with Forecast Competitions

The “M” competitions organized by Spyros Makridakis have had an enormous influence on the field of forecasting. They focused attention on what models produced good forecasts, rather than on the mathematical properties of those models.

The M3 competition involved 3003 time series with various seasonal frequencies and areas of application. The M3 data have continued to be used since 2000 for testing new time series forecasting methods. The recent M4 competition involved 100,000 time series! (See <https://robjhyndman.com/hyndsight/m4comp/> and <https://www.mcompetitions.unic.ac.cy/>).

Using a hybrid average of automatic Exponential Smoothing and ARIMA, SigmaXL (unofficially) ranked 10th out of 60 in the M4 Overall Weighted Average score, ahead of three well known commercial forecast software packages.

Control Charts for Autocorrelated Data

Statistical process control for autocorrelated processes typically use the EWMA (Exponentially Weighted Moving Average) one-step-ahead forecast model [1, 2, 7, 14, 15]. The formula used for EWMA is the same as simple exponential smoothing, but the smoothing parameter λ is typically used instead of α and X_t instead of y_t :

$$EWMA_{t+1} = \lambda X_t + (1 - \lambda)EWMA_t$$

with the starting forecast value $EWMA_1$ estimated as the data mean or target value. In the case of an EWMA control chart with independent data, the smoothing parameter λ is determined by desired average run length characteristics and is typically 0.2. For SPC for autocorrelated data, the smoothing parameter and initial level are determined by minimizing the sum-of-square forecast errors (residuals):

$$SSE = \sum_{t=1}^T (y_t - \hat{y}_t)^2 = \sum_{t=1}^T e_t^2.$$

using non-linear minimization methods like Newton-Raphson or Nelder-Mead Simplex.

The time series model forecasts the motion in the mean and an Individuals control chart is plotted of the residuals to detect assignable causes. Failure to account for the autocorrelation will produce limits that are too narrow resulting in excessive false alarms, or limits that are too wide resulting in misses.

Woodall and Faltin give some helpful guidelines on dealing with autocorrelation [20]:

- If possible, one should first attempt to remove the source of the autocorrelation.
- If the source of autocorrelation cannot be removed directly, then it may be possible to model the autocorrelation and use a feedback control scheme to reduce variability about a

specified target value.

- If the source of the autocorrelation cannot be removed directly, and feedback control is not a viable option, then it is important to monitor the process with control charts which do not repeatedly give signals due to presence of the autocorrelation.

SigmaXL provides an option to create an Individuals and Moving Limits (with One-Step-Ahead Forecast) control chart for the residuals of any of the forecast models, including:

- Exponential Smoothing
- Multiple Seasonal Decomposition with Exponential Smoothing
- ARIMA
- ARIMA with Predictors
- Multiple Seasonal Decomposition with ARIMA

The Moving Limits Residuals chart uses the one step prediction as the center line, so the control limits will move with the center line. If a Box-Cox transformation is used then an inverse transformation is applied to calculate the control limits. If the residuals are exponential smoothing multiplicative, the control limits are approximate and out-of-control signals may not exactly match the Individuals Chart. If that occurs, the Individuals Chart should be used to determine what points are out-of-control.

Tests for Special Causes are not recommended to be used with these control charts. While the residuals are independent, any sustained shift in the process mean is quickly adjusted for in the time series model, so the additional runs rules would not work as designed in order to detect small shifts.

Distinguishing an assignable cause as an outlier or as a sustained shift can be done by adding coded predictors in an ARIMA model and comparing the regression coefficient p-values. A potential outlier would be designated with a "1" and all other observations are "0". To test for a shift, all values up to the observation of interest would be a "0", and the remaining are designated as a "1" [23].

A Moving Range Residuals Chart is not included as that has too many false alarms [24, pp. 114-115].

If Tests for Special Causes or a Moving Range Chart are desired, simply create an Individuals and Moving Range Chart on the Forecast Residuals. Note that the "Add Data" feature with automatic update of residuals is not supported in this case.

References for Time Series Forecasting and Control Charts for Autocorrelated Data

- [1] Alwan, L.C., and Roberts, H.V. (1988), "Time Series Modeling for Statistical Process Control", *Journal of Business and Economic Statistics*, 6, 87-95.
- [2] Box, G.E.P., Jenkins, G.M., Reinsel, G.C. and Ljung, G.M. (2016). *Time Series Analysis, Forecasting and Control*, 5th edition, Wiley.
- [3] Cleveland, R.B., Cleveland, W.S., McRae, J.E., and Terpenning, I. (1990) "STL: A Seasonal-Trend Decomposition Procedure Based on Loess", *Journal of Official Statistics*, Vol. 6(1), pp. 3-73.
- [4] Friedman, J. H. (1984). "A variable span smoother. Laboratory for Computational Statistics", *Technical Report(5)*, Department of Statistics, Stanford University.
- [5] Fuller, W.A. (1996). *Introduction to Statistical Time Series*, 2nd ed. John Wiley and Sons, New York.
- [6] Gardner, G, Harvey, A.C. and Phillips, G.D.A. (1980). "Algorithm AS154. An algorithm for exact maximum likelihood estimation of autoregressive - moving average models by means of Kalman filtering". *Applied Statistics*, Vol. 29, pp. 311 - 322.
- [7] Hunter, J.S. (1986), "The Exponentially Weighted Moving Average," *Journal of Quality Technology*, 18, 203-210.
- [8] Hyndman, R.J., M. Akram, and B.C. Archibald, (2006). "The admissible parameter space for exponential smoothing models." *Annals of the Institute of Statistical Mathematics* Vol. 60. pp. 407-426.
- [9] Hyndman, R.J., & Athanasopoulos,G. (2018). *Forecasting: principles and practice*, 2nd edition, OTexts: Melbourne, Australia. OTexts.com/fpp2.
- [10] Hyndman, R.J. and Y. Khandakar, (2008). "Automatic Time Series Forecasting: The forecast Package for R." *Journal of Statistical Software*, Vol. 27(3), pp. 1-22.
- [11] Hyndman, R.J., A.B. Koehler, J. Keith Ord and R.D. Snyder (2008). *Forecasting with Exponential Smoothing. The State Space Approach* Springer-Verlag, Berlin Heidelberg.
- [12] Hyndman R, Athanasopoulos G, Bergmeir C, Caceres G, Chhay L, O'Hara-Wild M, Petropoulos F, Razbash S, Wang E, Yasmeeen F (2019). `_forecast`: Forecasting functions for time series and linear models_. R package version 8.9.
- [13] Kwiatkowski, D., Phillips, P.C.B, Schmid, P. and Shin, T. (1992). "Testing the null hypothesis of stationarity against the alternative of a unit root: How sure are we that economic series have a unit root", *Journal of Econometrics*, Vol. 54, pp. 159-178.
- [14] Montgomery, D. C., and Mastrengelo, C.M. (1991), "Some Statistical Process Control Methods for Autocorrelated Data", *Journal of Quality Technology*, 23, 179-204.
- [15] Montgomery, D. C. (2013). *Introduction to Statistical Quality Control*, 7th edition, Wiley.
- [16] Montgomery, D. C., Jennings, C.L., and Kulahci, M. (2015). *Introduction to Time Series Analysis and Forecasting*, 2nd edition, Wiley.
- [17] Moritz, S., Sarda, A. et al. (2015). "Comparison of different Methods for Univariate Time Series Imputation in R", <https://arxiv.org/ftp/arxiv/papers/1510/1510.03924.pdf>.
- [18] *NIST/SEMATECH e-Handbook of Statistical Methods*,

-
- <https://www.itl.nist.gov/div898/handbook>.
- [19] Wang, X., Smith, K. and Hyndman, R. (2006) "Characteristic-Based Clustering for Time Series Data", *Data Mining and Knowledge Discovery*, Vol. 13, pp. 335-364.
- [20] Woodall, W.H. and Faltin, F.W. "Autocorrelated Data and SPC", *ASQC Statistics Division Newsletter*, 13(4).
- [21] Hobijn, B., P. H. Franses, and M. Ooms. 1998. "Generalizations of the KPSS-test for stationarity", *Econometric Institute Report 9802/A*, Econometric Institute, Erasmus University Rotterdam.
- [22] Schwert, "G.W. Tests for Unit Roots: A Monte Carlo Investigation", *Journal of Business and Economic Statistics*, 7, 1989, 147-160.
- [23] The ARIMA Procedure – Detecting Outliers, SAS/ETS(R) 9.3 User's Guide, http://support.sas.com/documentation/cdl/en/etsug/63939/HTML/default/viewer.htm#etsug_arima_sect045.htm.
- [24] Fuentes, Jesús Cuéllar (2008), "Statistical Monitoring of a Process with Autocorrelated Output and Observable Autocorrelated Measurement Error" (Doctoral Dissertation), Baylor University. <https://baylor-ir.tdl.org/handle/2104/5185>.
- [25] Kostenko, A.V. and Hyndman, R.J. (2008), "Forecasting without significance tests?", <https://robjhyndman.com/papers/sst2.pdf>



About SigmaXL, Inc.

Established in 1998, SigmaXL Inc. is a leading provider of user-friendly Excel Add-ins for Lean Six Sigma graphical and statistical tools and Monte Carlo simulation.

Our flagship product, SigmaXL[®], was designed from the ground up to be a cost-effective, powerful, but easy to use tool that enables users to measure, analyze, improve and control their service, transactional, and manufacturing processes. As an add-in to the already familiar Microsoft Excel, SigmaXL[®] is ideal for Lean Six Sigma training and application, or use in a college statistics course.

DiscoverSim™ enables you to quantify your risk through Monte Carlo simulation and minimize your risk with global optimization. Business decisions are often based on assumptions with a single point value estimate or an average, resulting in unexpected outcomes. DiscoverSim™ allows you to model the uncertainty in your inputs so that you know what to expect in your outputs.

Contact Information:

Technical Support: 1-866-475-2124 (Toll Free in North America) or

1-519-579-5877

Sales: 1-888-SigmaXL (888-744-6295)

E-mail: support@SigmaXL.com Web: www.SigmaXL.com