

## Stability Study



Revised: 6/28/2020



Summary .....	1
Data Input.....	3
Analysis Options .....	4
Tables and Graphs.....	5
Analysis Summary .....	6
Shelf Life Plot.....	9
Shelf Life Estimates .....	10
Observed Versus Predicted .....	11
Unusual Residuals.....	12
Residuals versus X.....	12
Residuals versus Predicted.....	14
Residuals versus Row Number .....	15
Residuals Probability Plot.....	16
Influential Points .....	18
Comparison of Alternative Models.....	19
Save Results .....	25
References.....	26

## Summary

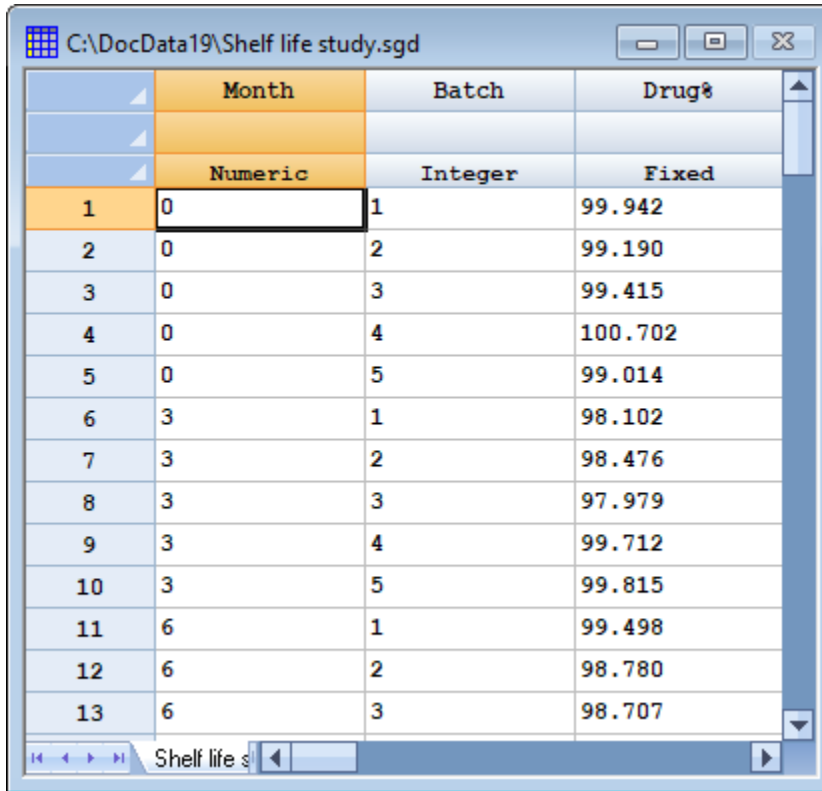
Stability studies are commonly used by pharmaceutical companies to estimate the rate of drug degradation and to establish shelf life. Measurements are typically made on samples from multiple batches at different points in time. Of primary interest is estimating the time at which the lower prediction limit from the degradation model crosses the lower specification limits for the drug.

Depending on the structure of the data, batches may be treated as either a fixed or random factor.

**Sample StatFolios:** *stability.sgp*, *stability2.sgp*

### Sample Data – Fixed Batch Effects

The data file *shelf life study.sgd* contains data from a typical stability study:



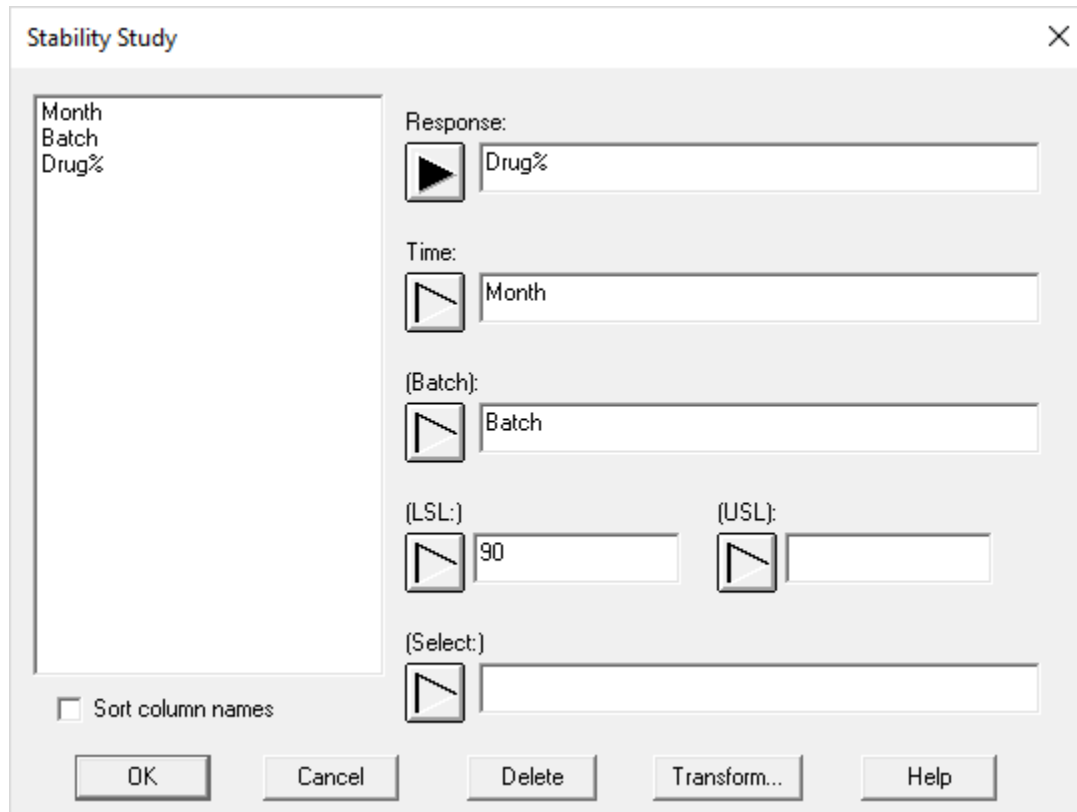
	Month	Batch	Drug%
	Numeric	Integer	Fixed
1	0	1	99.942
2	0	2	99.190
3	0	3	99.415
4	0	4	100.702
5	0	5	99.014
6	3	1	98.102
7	3	2	98.476
8	3	3	97.979
9	3	4	99.712
10	3	5	99.815
11	6	1	99.498
12	6	2	98.780
13	6	3	98.707

The file contains measurements of the *Drug%* from 5 batches of a product sampled between 0 and 48 months after which it was produced. The goal of the study is to determine the longest time after production at which one may be 95% confident that at least 95% of the product will be above the lower specification limit

$$LSL: Drug\% = 90.0.$$

## Data Input

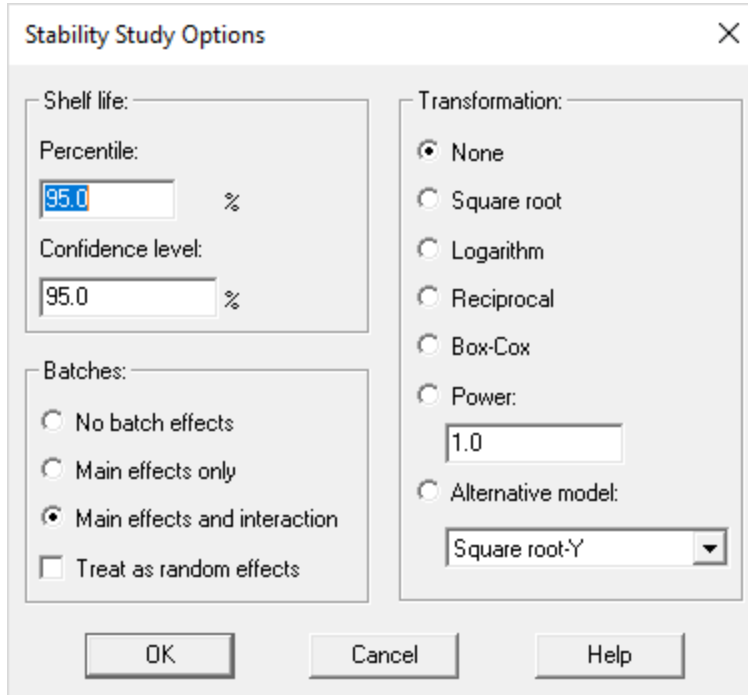
When the *Stability Study* procedure is selected from the Statgraphics menu, a data input dialog box is displayed:



- **Response:** name of a numeric column containing the data to be analyzed.
- **Time:** name of a numeric column containing the length of time after production corresponding to each data value.
- **Batch:** if sample have been taken from more than one batch, the name of a numeric or character column identifying the batch from which each data value was collected.
- **LSL:** an optional lower specification limit.
- **USL:** an optional upper specification limit.
- **Select:** optional subset selection.

## Analysis Options

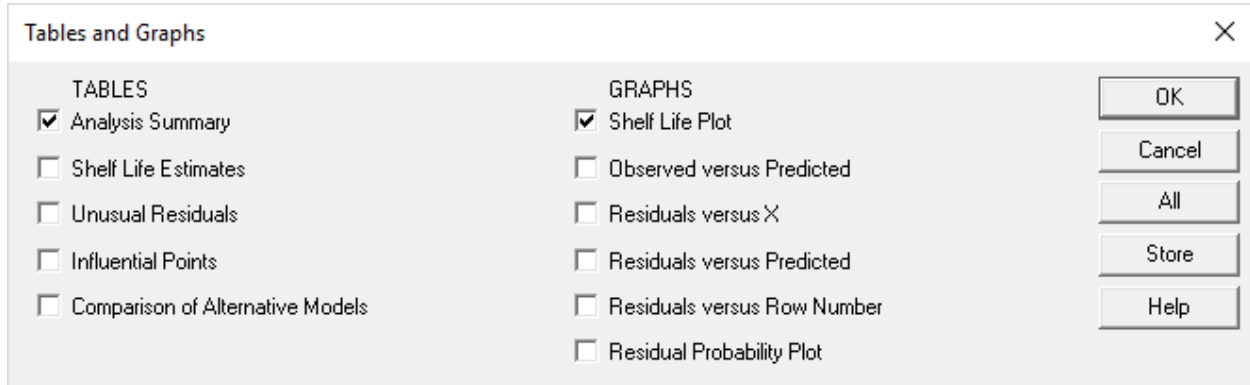
The *Analysis Options* dialog box sets various options for modeling the data:



- **Shelf life percentile:** the percentile of the data used to determine the shelf life. The fitted model will be used to determine the time at which this percentile crosses the relevant specification limit.
- **Confidence level:** the confidence level used in determining the shelf life.
- **Batches:** the manner in which batch effects are treated in the statistical model. Users may decide whether or not to *include main effects* of the batches, and whether or not to *include interactions* between batches and time. In addition, batches may be treated as either a *fixed* or *random factor*.
- **Transformation:** specifies whether or not the dependent variable in the model is a transformation of the data specified on the data input dialog box. If desired, the model may be fit to the untransformed data values, their square roots, logarithms, reciprocals, or the data values raised to the specified power. In addition, the *Box-Cox* approach may be used to automatically determine the best power transformation.
- **Alternative model:** instead of a linear model, various other curvilinear models may be fit.

## Tables and Graphs

The following tables and graphs may be created:



Tables and Graphs		×
<b>TABLES</b>	<b>GRAPHS</b>	<b>OK</b>
<input checked="" type="checkbox"/> Analysis Summary	<input checked="" type="checkbox"/> Shelf Life Plot	<b>Cancel</b>
<input type="checkbox"/> Shelf Life Estimates	<input type="checkbox"/> Observed versus Predicted	<b>All</b>
<input type="checkbox"/> Unusual Residuals	<input type="checkbox"/> Residuals versus X	<b>Store</b>
<input type="checkbox"/> Influential Points	<input type="checkbox"/> Residuals versus Predicted	<b>Help</b>
<input type="checkbox"/> Comparison of Alternative Models	<input type="checkbox"/> Residuals versus Row Number	
	<input type="checkbox"/> Residual Probability Plot	

## Analysis Summary

The *Analysis Summary* summarizes the input data and the fitted model. The top portion of the table shows the estimated model coefficients:

### Stability Study - Drug% vs. Month

Response variable: Drug%

Time variable: Month

Batch variable: Batch

Lower specification limit: 90

Number of batches: 5

Batch type: fixed

Number of complete cases: 40

Model: linear

Terms in model: Month, Batch, Month\*Batch

#### Model Coefficients

Parameter	Estimate	Standard Error	Df	T Statistic	P-Value
Constant	99.6524	0.125709	35	792.725	0.0000
Month	-0.141033	0.00546043	35	-25.8282	0.0000
Batch=1	0.200423	0.251417	35	0.797173	0.4307
Batch=2	-0.528445	0.251417	35	-2.10187	0.0428
Batch=3	-0.369701	0.251417	35	-1.47047	0.1504
Batch=4	0.806633	0.251417	35	3.20834	0.0029
Batch=5	-0.10891	0.251417	35	-0.433184	0.6675
Month*Batch=1	-0.0142641	0.0109209	35	-1.30614	0.2000
Month*Batch=2	0.0355359	0.0109209	35	3.25395	0.0025
Month*Batch=3	-0.0196544	0.0109209	35	-1.79971	0.0805
Month*Batch=4	0.00893527	0.0109209	35	0.818183	0.4188
Month*Batch=5	-0.0105526	0.0109209	35	-0.966281	0.3405

#### Regression Models

Batch	Intercept	Slope
1	99.8528	-0.155297
2	99.1239	-0.105497
3	99.2827	-0.160687
4	100.459	-0.132098
5	99.5435	-0.151586

Of particular interest are:

1. **Number of complete cases:** number of cases with no missing data for any of the variables specified. Any incomplete cases are removed before the algorithm is performed.
2. **Terms in model:** the terms included in the fitted model. The output above indicates that a *linear* model was fit which included *Month*, *fixed* effects for *Batch*, and interactions between *Month* and *Batch*.
3. **Model coefficients:** the estimated model coefficients, their standard errors, and the results of t-tests run to determine whether the coefficients are significantly different than 0. The small P-value for *Month* indicates that the slope of the linear model with respect to time is statistically significant. The t-tests for each batch indicate whether the intercept for that batch is significantly different from the overall intercept. The terms similar to *Month\*Batch=1* test whether the slope with respect for time for the indicated batch is significantly different from the overall slope.
4. **Regression models:** the estimated coefficients for each batch.

In the current example, the slope with respect to *Month* is highly significant. Many of the batch intercept terms are significant, as are some of the interaction terms.

The second portion of the *Analysis Summary* displays an analysis of variance:

Analysis of Variance					
Source	Sum of Squares	Df	Mean Square	F-Ratio	P-Value
Model	223.407	0	24.823	80.59	0.0000
Residual	10.7801	0	0.308003		
Total (Corr.)	234.187	0			

R-squared = 95.3968 percent  
R-squared (adjusted for d.f.) = 94.2131 percent  
R-squared (predicted) = 92.5022 percent  
Standard Error of Est. = 0.55498  
Mean absolute error = 0.394657  
Durbin-Watson statistic = 1.60007 (P=0.1333)  
Lag 1 residual autocorrelation = 0.19771

Further ANOVA for Variables in the Order Fitted					
Source	Sum of Squares	Df	Mean Square	F-Ratio	P-Value
Month	205.467	1	205.467	667.09	0.0000
Batch	13.7174	4	3.42934	11.13	0.0000
Month*Batch	4.22242	4	1.0556	3.43	0.0182

The F-Ratio for Model tests whether the model as a whole indicates whether or not there is a statistically significant change in the dependent variable over time. The terms in the Further ANOVA table indicate the statistical significance of *Month*, *Batch* and their interaction when

added sequentially to the model. The small P-values for *Month* and *Month\*Batch* indicate the presence of significant batch effects and interactions.

The third portion of the *Analysis Summary* displays statistics for residuals from the fitted model:

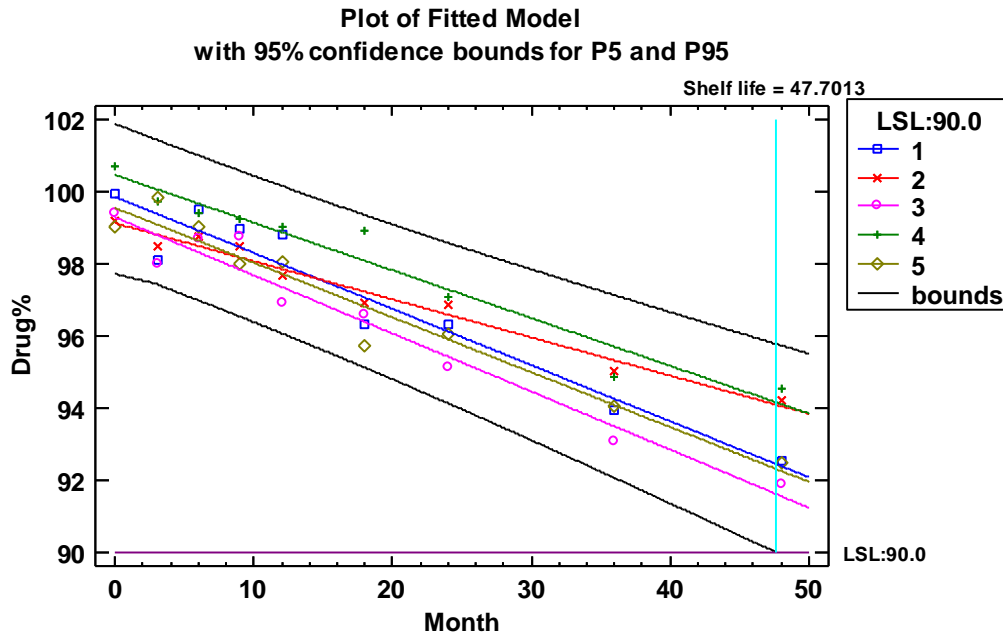
Residual Analysis		
	Estimation	Validation
n	45	
MSE	0.308003	
MAE	0.394657	
MAPE	0.405406	
ME	9.4739E-16	
MPE	-0.00249198	

It shows the number of residuals (*n*), the mean squared error (*MSE*), the mean absolute error (*MAE*), the mean absolute percentage error (*MAPE*), the mean error (*ME*), and the mean percentage error (*MPE*). If any rows of the input data were withheld from the fit using the *Select* field on the data input dialog box, separate statistics will be displayed for those observations used to fit the model (*Estimation*) and those observations that were not used to fit the model (*Validation*).



## Shelf Life Plot

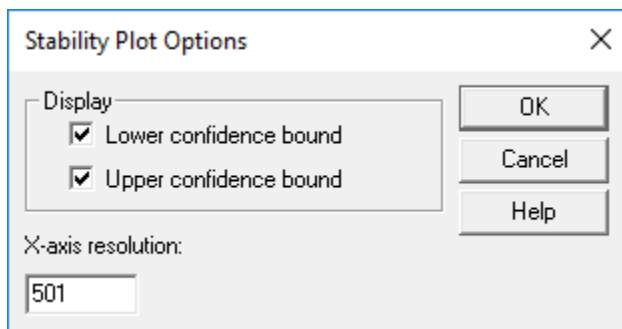
This plot shows the fitted statistical model:



A separate line is displayed for each of the batches. In addition, confidence bounds are included for the 5<sup>th</sup> and 95<sup>th</sup> percentiles of the distribution of *Drug%* as a function of *Month*. When there are different models for each batch, as in the above plot, the bounds displayed are the most extreme values of bounds calculated separately for each of the batches.

The plot also displays the estimated shelf life, which is determined by the intersection of the lower confidence bounds with the lower specification limit. In this case, the shelf life is estimated to be 47.7013 months.

### Pane Options



- **Display:** whether to display the lower and/or upper confidence bounds for the percentiles.
- **X-axis resolution:** the number of locations along the X-axis at which the confidence bounds are calculated.

## Shelf Life Estimates

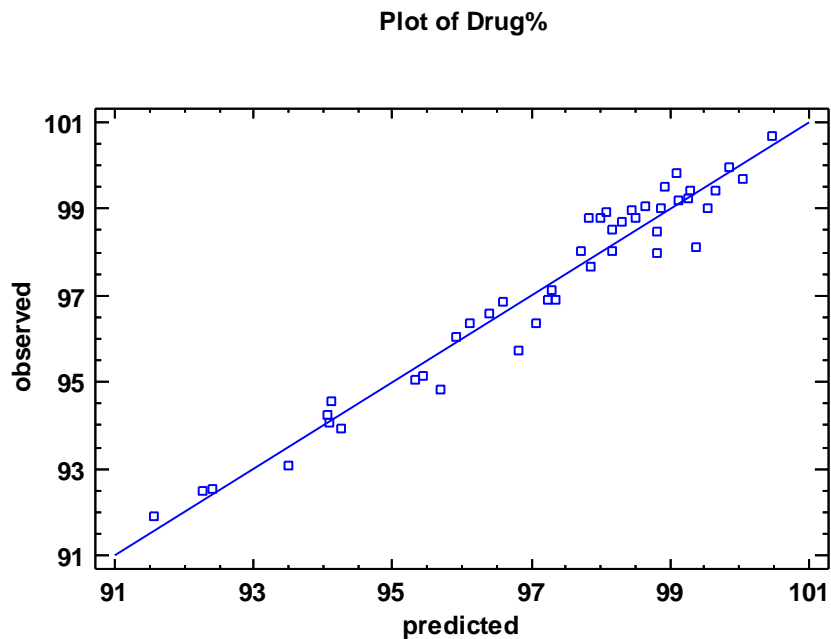
This pane shows the shelf life estimates determined for each batch.

Shelf Life Estimates		
Batch	Shelf life	Limit
1	52.4836	lower
2	67.6797	lower
3	47.7013	lower
4	64.5984	lower
5	51.8079	lower
<b>Combined</b>	47.7013	lower

The overall estimated shelf life for all batches combined is the smallest of the individual batch values.

## Observed Versus Predicted

This plot shows the observed values of the dependent variable versus the values predicted by the model:



If the model fits the data well, the points should lie close to the diagonal line. Any suggestion of curvature would cast doubt on the use of the currently selected model. For the sample data, the linear model appears to give a good fit.

## Unusual Residuals

Once the model has been fit, it is useful to study the residuals to determine whether any outliers exist that should be removed from the data. In a regression, the residuals are defined by

$$e_i = y_i - \hat{y}_i \quad (1)$$

i.e., the residuals are the differences between the observed data values and the fitted model.

The *Unusual Residuals* pane lists all observations that have Studentized residuals of 2.0 or greater in absolute value.


Unusual Residuals					
Row	X	Y	Predicted Y	Residual	Studentized Residual
6	3.0	98.102	99.3869	-1.28493	-2.86
30	18.0	95.736	96.8149	-1.07894	-2.17

Studentized residuals greater than 3 in absolute value correspond to points more than 3 standard deviations from the fitted model, which is an extremely rare event for a normal distribution. In the sample data, row #6 is more than 2.5 standard deviations out.

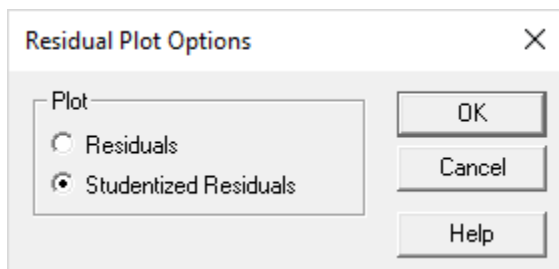
## Residuals versus X

This plot displays the residuals plotted versus the observed values of time:

Residual Plot



### Pane Options




The following residuals may be plotted on each residual plot:

1. *Residuals* – the residuals from the least squares fit.
2. *Studentized residuals* – the difference between the observed values  $y_i$  and the predicted values  $\hat{y}_i$  when the model is fit using all observations except the  $i$ -th, divided by the estimated standard error. These residuals are sometimes called *externally deleted residuals*, since they measure how far each value is from the fitted model when that model is fit using all of the data except the point being considered. This is important, since a large outlier might otherwise affect the model so much that it would not appear to be unusually far away from the line.

## Residuals versus Predicted

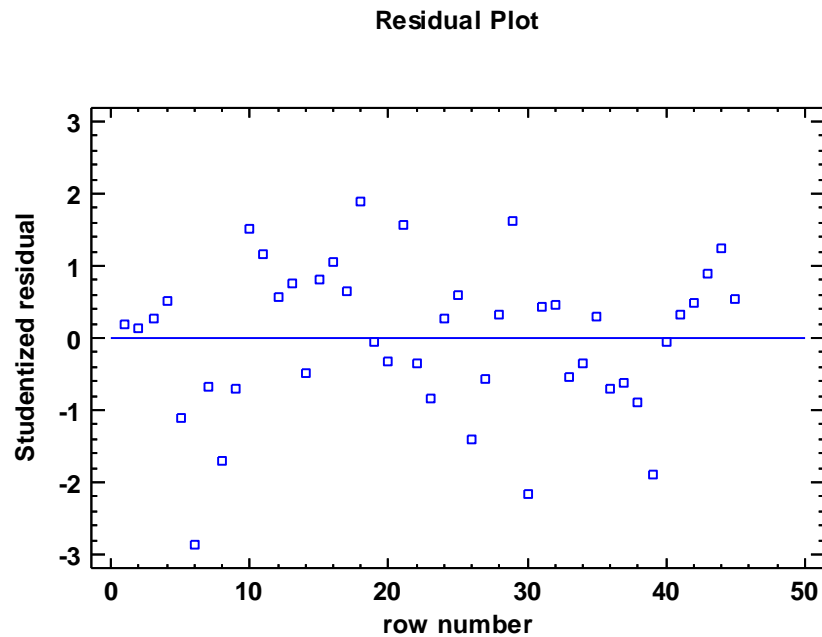
This plot displays the residuals plotted versus the values predicted by the fitted model:



This plot is helpful in detecting any heteroscedasticity in the data (changes in the error variance as the magnitude of the predicted values changes). Heteroscedasticity can sometimes be removed by fitted the logarithms of the data rather than the untransformed values.

## Residuals versus Row Number

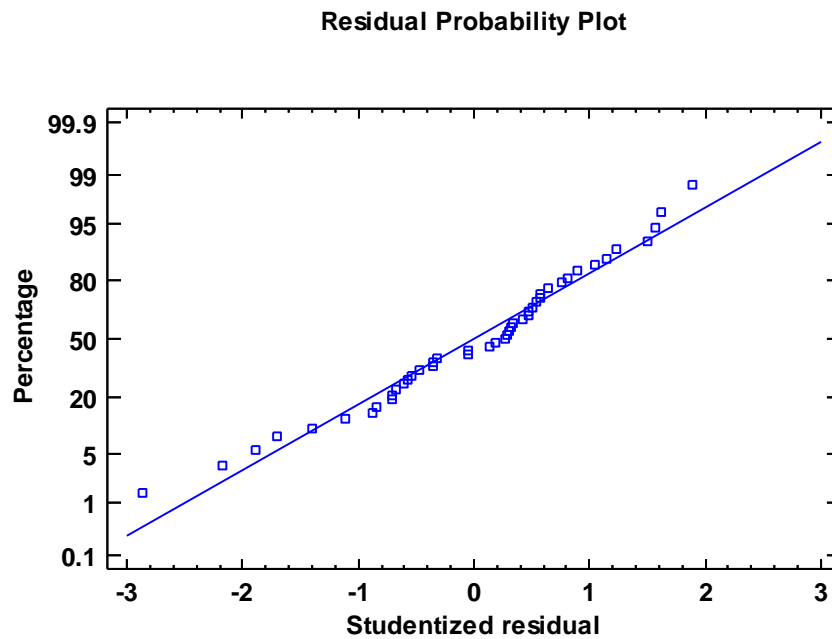
This plot shows the residuals versus row number in the datasheet:



If the data are arranged in chronological order, any pattern in the data might indicate an outside influence.

## Residuals Probability Plot

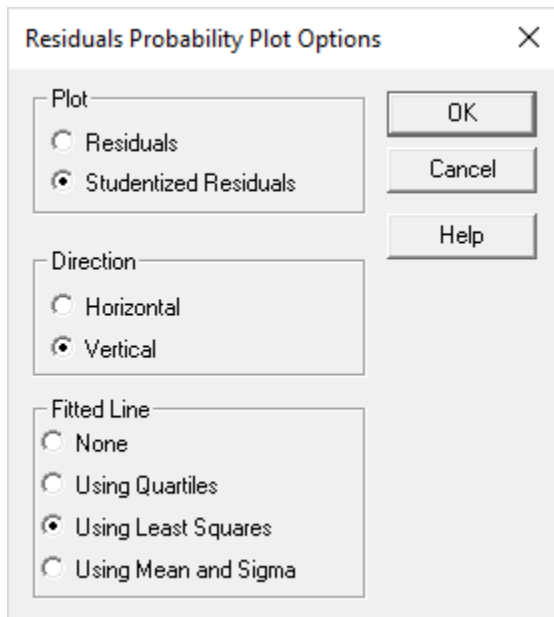
This plot can be used to help determine whether or not the deviations around the line follow a normal distribution.



If the deviations follow a normal distribution, they should fall approximately along a straight line.

### Pane Options





- **Plot:** the type of residuals to plot.
- **Direction:** the axis along which percentages are displayed.
- **Fitted line:** the method used to determine the reference line (if any) on the plot. For details, see the document titled *Normal Probability Plot*.

## Influential Points

In fitting a regression model, all observations do not have an equal influence on the parameter estimates in the fitted model. In a simple regression, points located at very low or very high values of X have greater influence than those located nearer to the mean of X. The *Influential Points* pane displays any observations that have high influence on the fitted model:

Influential Points					
Row	X	Y	Predicted Y	Studentized Residual	Leverage
Average leverage of single data point = 0.222222					

The above table shows every point with *leverage* equal to 3 or more times that of an average data point, where the leverage of an observation is a measure of its influence on the estimated model coefficients. In general, values with leverage exceeding 5 times that of an average data value should be examined closely, since they have unusually large impact on the fitted model.

## Comparison of Alternative Models

The *Comparison of Alternative Models* pane shows the R-squared values obtained when fitting each of the 27 available models (except models that cannot be fit since the data includes values of  $X=0$ ):

Comparison of Alternative Models		
Model	R-Squared	Adj. R-squared
Reciprocal-Y	95.68%	94.57%
Exponential	95.55%	94.40%
Square root-Y	95.47%	94.31%
Linear	95.40%	94.21%
Squared-Y	95.23%	94.01%
Reciprocal-Y squared-X	88.63%	85.70%
Logarithmic-Y squared-X	88.09%	85.02%
Square root-Y squared-X	87.81%	84.68%
Squared-X	87.54%	84.34%
Squared-Y square root-X	87.19%	83.89%
Double squared	86.98%	83.64%
Square root-X	86.95%	83.60%
Double square root	86.83%	83.44%
Logarithmic-Y square root-X	86.71%	83.29%
Reciprocal-Y square root-X	86.45%	82.96%
Logarithmic-X	<no fit>	
Square root-Y logarithmic-X	<no fit>	
Multiplicative	<no fit>	
Reciprocal-Y logarithmic-X	<no fit>	
Squared-Y logarithmic-X	<no fit>	
Reciprocal-X	<no fit>	
Square root-Y reciprocal-X	<no fit>	
S-curve model	<no fit>	
Double reciprocal	<no fit>	
Squared-Y reciprocal-X	<no fit>	
Logistic	<no fit>	
Log probit	<no fit>	

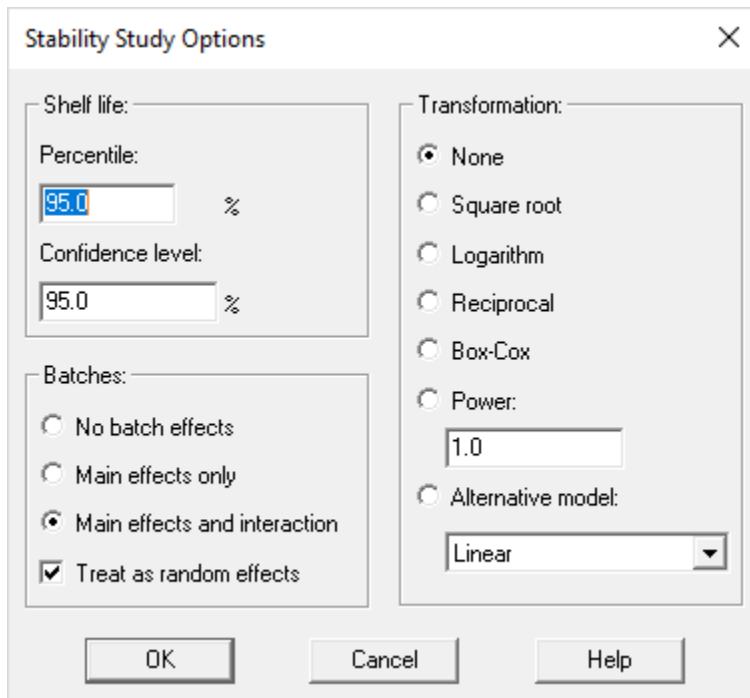
The models are listed in decreasing order of R-squared. When selecting an alternative model, consideration should be given to those models near the top of the list. However, since the R-Squared statistics are calculated after transforming X and/or Y, the model with the highest R-squared may not be the best. You should always plot the fitted model to see whether it does a good job for your data.

In the above table, the linear model has an R-squared value near the top of the list.

## Sample Data – Random Batch Effects

Rather than treating the batches as fixed effects and determining the smallest estimated shelf life amongst the batches, the batches can be treated as a random factor. Variation amongst the intercepts and slopes of the batches may then be treated as additional variance components and accounted for when the confidence bounds for the percentiles are calculated.

Using the same data used above, checking the *Treat as random effects* checkbox on the *Analysis Options* dialog box will cause batches to be treated as a random factor:



**Stability Study Options** [X]

**Shelf life:**

Percentile:  %

Confidence level:  %

**Batches:**

- No batch effects
- Main effects only
- Main effects and interaction
- Treat as random effects

**Transformation:**

- None
- Square root
- Logarithm
- Reciprocal
- Box-Cox
- Power:
- Alternative model:

OK Cancel Help

## Analysis Summary

The top portion of the *Analysis Summary* is shown below:

<b>Stability Study - Drug% vs. Month</b>					
Response variable: Drug%					
Time variable: Month					
Batch variable: Batch					
Lower specification limit: 90					
Number of batches: 5					
Batch type: random					
Number of complete cases: 45					
Model: linear					
Terms in model: Month, Batch, Month*Batch					
<b>Variance Components</b>					
Source	Variance	% of Total	S.E.	Z	P-Value
Batch	0.210933	40.7121%	0.196833	1.07164	0.1419
Month*Batch	0.00038068	0.0735%	0.000359685	1.05837	0.1449
Error	0.306795	59.2144%	0.0727517	4.21701	0.0000
Total	0.518108				
<b>Model Coefficients</b>					
Parameter	Estimate	Standard Error	Df	T Statistic	P-Value
Constant	99.6524	0.240681	4.33768	414.044	0.0000
Month	-0.141033	0.0102876	4.33768	-13.709	0.0001
<b>Model Statistics</b>					
Std. Error of Est.	MAE	R-squared	R-squared (adjusted)		
0.55389	0.664178	95.16	95.05		

The main difference from the output of a fixed effects model is the estimation of variance components. Since batch is a random factor, variation amongst batches adds additional contributions to the error variance, potentially both from differences amongst the batch means and differences between the interactions with time. The *Variance Components* table shows estimates of those variance components, together with their percentage contribution to the overall error variance. In addition, statistical tests are performed to determine whether or not those contributions are significantly different from zero.

An additional table is also displayed containing estimates of the random effects:

Random Effects	
	Estimate
Batch	
1	0.0997459
2	-0.270669
3	-0.354173
4	0.647068
5	-0.121968
Month*Batch	
1	-0.00938575
2	0.0231879
3	-0.0172481
4	0.0121081
5	-0.00866203

The model used to predict the response of the  $i^{\text{th}}$  batch at time  $x_{ij}$  when assuming random batches has the general form

$$\hat{y}_{ijk} = \hat{\alpha} + \hat{\beta}x_{ij} + \hat{\mu}_i + \hat{\eta}_ix_{ij}$$

where

$\hat{\alpha}$  = estimated mean intercept of all batches

$\hat{\beta}$  = estimated mean slope of all batches

$\hat{\mu}_i$  = estimated deviation of batch  $i$  intercept from the mean intercept

$\hat{\eta}_i$  = estimated deviation of batch  $i$  slope from the mean slope

The effects tabulated above are the estimated deviations  $\hat{\mu}_i$  and  $\hat{\eta}_i$ , which are assumed to be randomly drawn from normal distributions with variances equal to  $\sigma_{\mu}^2$  and  $\sigma_{\eta}^2$ .

## Shelf Life Estimates

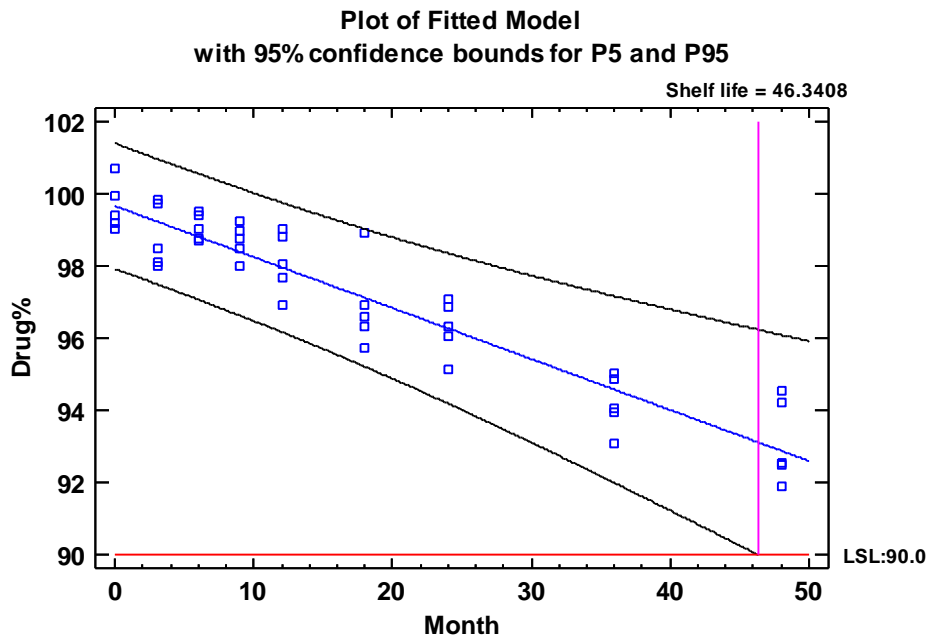
There is a single shelf life estimate that applies to all batches:

Shelf Life Estimates	
Shelf life	Limit
46.3408	lower

This is calculated determining when the confidence limit for the desired percentile crosses the specification limit, where the confidence limits accounts for the variation from all of the variance components.

## Shelf Life Plot

The *Shelf Life Plot* shows the estimated model calculated from the mean intercept and mean slope, with confidence limits for the percentile specified on the *Analysis Options* dialog box:



## Residuals and Predicted Values

In plots of predicted values and residuals, the predictions are calculated using the mean intercept and mean slope:

$$\hat{y}_{ijk} = \hat{\alpha} + \hat{\beta}x_{ij}$$

The so-called marginal residuals

$$e_{ijk} = y_{ijk} - \hat{y}_{ijk}$$

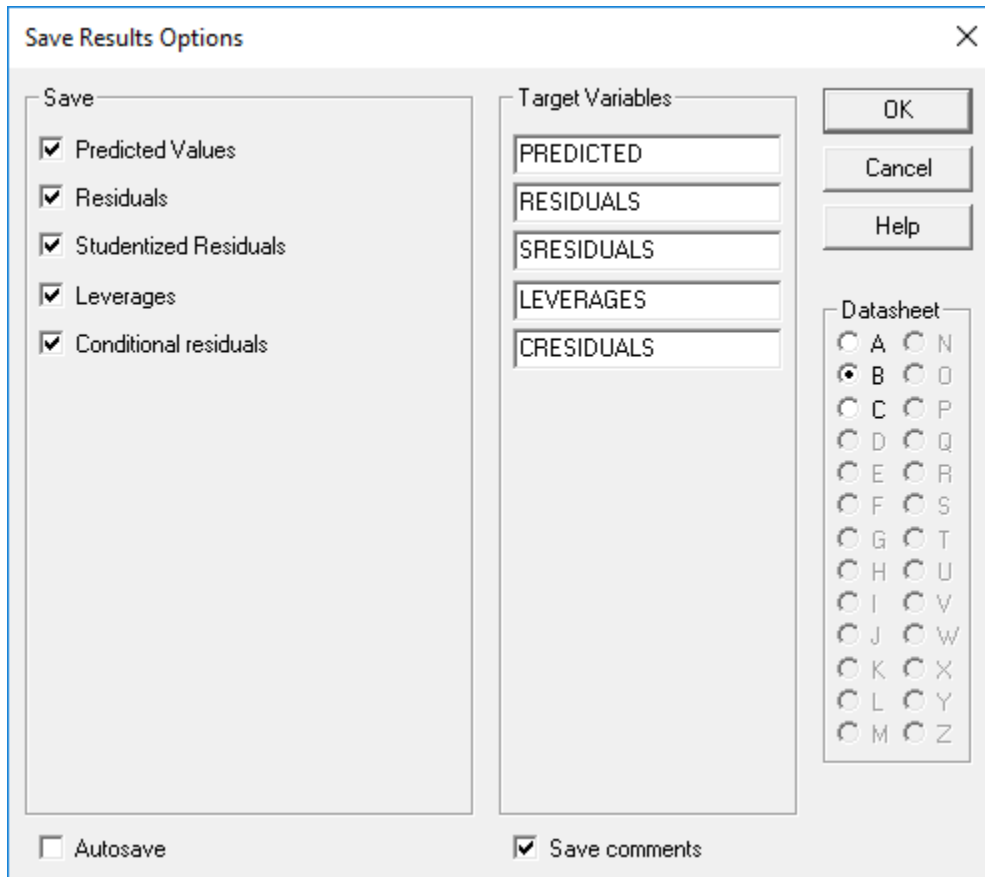
thus include the variation amongst the batch intercepts and slopes. This is accounted for when calculating the standardized residuals:

Row	X	Y	Predicted Y	Residual	Standardized Residual
29	18.0	98.907	97.1138	1.7932	2.38

The standardized residuals are calculated by dividing each residual by its standard error, using a model fitted using all of the observations. They are *not* deleted residuals as are the Studentized residuals displayed for a model with fixed batch effects.



## Save Results



The dialog box titled "Save Results Options" contains the following elements:

- Save:** A list of checkboxes for saving different types of results:
  - Predicted Values
  - Residuals
  - Studentized Residuals
  - Leverages
  - Conditional residuals
- Target Variables:** A list of text boxes containing the names of variables to be saved:
  - PREDICTED
  - RESIDUALS
  - SRESIDUALS
  - LEVERAGES
  - CRESIDUALS
- Datasheet:** A grid of radio buttons labeled with letters A through Z. Radio button 'B' is selected.
- Buttons:** OK, Cancel, and Help.
- Options:**
  - Autosave
  - Save comments

The following results may be saved to the datasheet:

1. *Predicted Values* – the fitted values corresponding to each row of the datasheet.
2. *Residuals* – the ordinary (marginal) residuals.
3. *Studentized Residuals* – the Studentized residuals (for fixed batch effects) or standardized residuals (for random batch effects).
4. *Leverages* – the leverage associated with each observation.
5. *Conditional residuals* (random batch effects only) – the residual values after accounting for the individual batch effects.

## References

Guidance for Industry: Q1E Evaluation of Stability Data (2004). U.S. Food and Drug Administration.

Statistical Design and Analysis of Stability Studies (2007). Shein-Chung Chow. Chapman and Hall.