

Compare – Multiple Response Variables

This procedure compares the data contained in three or more *Response* columns. It performs a one-way analysis of variance to determine whether or not there are significant differences between the column means, and also conducts a test to compare the variances within each column. Tukey's HSD multiple comparison method is used to identify which column means are significantly different than which other column means.

The data for this analysis consist of samples from m populations. Let

x_{ij} = i -th measurement in sample j

n_j = size of sample j

Access

Highlight: 3 or more *Response* columns.

Select: *Compare* from the main menu.

Output Page 1: A scatterplot of the data in each sample, together with summary statistics.

Output Page 2: Box-and-whisker plots for the data in each sample, together with the results of statistical tests that compare the samples.

Output Page 3: A plot of the sample means with Tukey's HSD intervals (or other selected uncertainty intervals).

Output Page 4: A plot of the variability within each group.

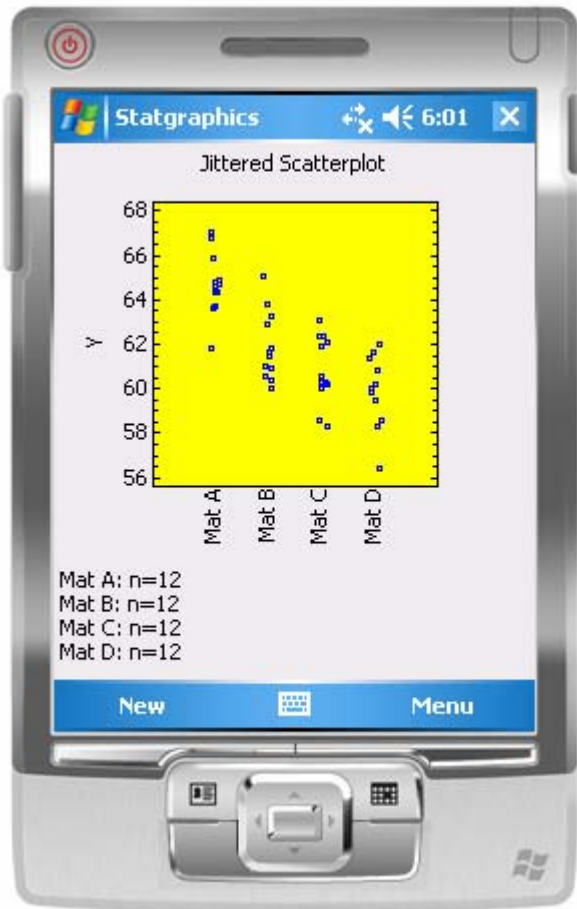
Sample Data

The file *widgets.sgm* contains data on the strength of widgets made from $m = 4$ different types of material. $n = 12$ widgets were samples from each of the four materials. A portion of the data is shown below.

Mat A	Mat B	Mat C	Mat D
64.7	60.4	58.3	60.8
64.8	61.8	62.1	60.2
66.8	63.3	62.4	59.8
67.0	61.6	60.3	58.3
64.9	61.0	60.6	56.4

Scatterplot


The *Scatterplot* displays the data in each column.



The points are jittered (randomly off in the horizontal direction) to help prevent overplotting. Also displayed is the number of observations in up to 8 columns.

Boxplots

The *Boxplots* page displays a box-and-whisker plot for the data in each column.



The plot is constructed in the following manner:

- A box is drawn extending from the *lower quartile* of each sample to the *upper quartile*. This is the interval covered by the middle 50% of the data values when sorted from smallest to largest.
- A horizontal line is drawn at the *median* (the middle value).
- A plus sign is placed at the location of the sample mean.
- Whiskers are drawn from the edges of the box to the largest and smallest data values.

Below the plot are the results of four statistical tests:

- (1) A one-way analysis of variance to compare the means of the populations from which the samples were taken. The output displays:

MSB: mean square between columns, which estimates the variance between the columns. It is calculated from the deviations of the group mean around the grand mean:

$$MSB = \frac{\sum_{j=1}^m n_j (\bar{x}_j - \bar{x})^2}{m - 1} \quad (1)$$

where the group means are given by

$$\bar{x}_j = \frac{\sum_{i=1}^{n_j} x_{ij}}{n_j} \quad (2)$$

and the grand mean is

$$\bar{x} = \frac{\sum_{j=1}^m \sum_{i=1}^{n_j} x_{ij}}{\sum_{j=1}^m n_j} \quad (3)$$

MSE: mean square error, which estimates the variance within the columns. It is calculated from the deviations of the observations around their respective column means:

$$MSE = \frac{\sum_{j=1}^m \sum_{i=1}^{n_j} (x_{ij} - \bar{x}_j)^2}{\sum_{j=1}^m n_j - m} \quad (4)$$

F-ratio: the ratio of the mean squares

$$F = \frac{MSB}{MSE} \quad (5)$$

P-value: a P-value calculated by referring the calculated F-ratio to Snedecor's F distribution. A small value of P indicates that there are significant differences amongst the group means. In the sample data, the difference is extremely significant.

- (2) **Levene's test**, which is used to determine whether or not the group variances are significantly different. Levene's test performs a one-way analysis of variance on the variables

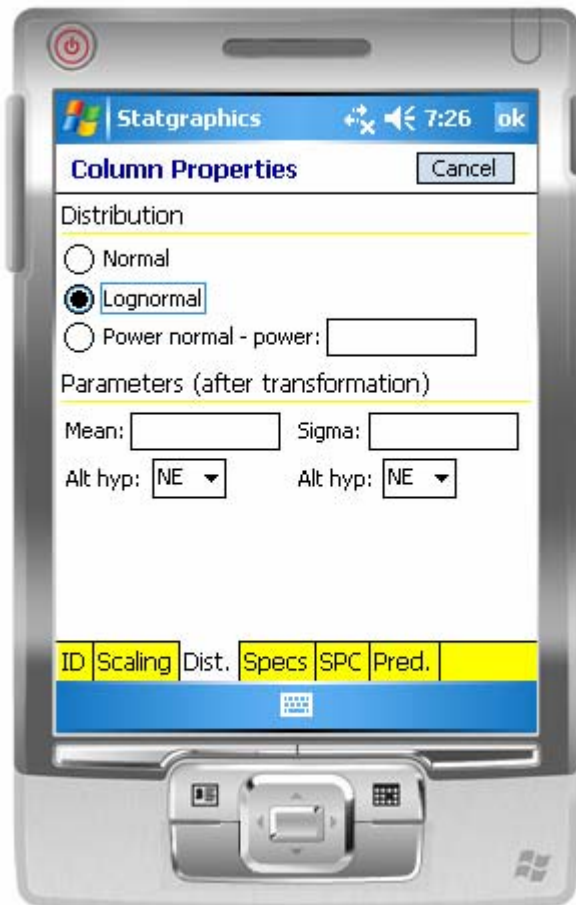
$$z_{ij} = |x_{ij} - \bar{x}_j| \quad (6)$$

The tabulated statistic is the F statistic from the resulting ANOVA table. A P-value above 0.05, such as that displayed for the sample data, means that there is not a statistically significant difference amongst the group variances at the 5% significance level.

- (3) The **Kruskal-Wallis test**, which compares the group medians by cobining all of the observations into a single sample, ranking them, and then calculating the average rank for each group. A small P-value, such as that illustrated for the sample data, indicates that the populations medians are significantly different from one another.

All of the tests here assume that the data are samples from normal distributions. To select a different distribution:

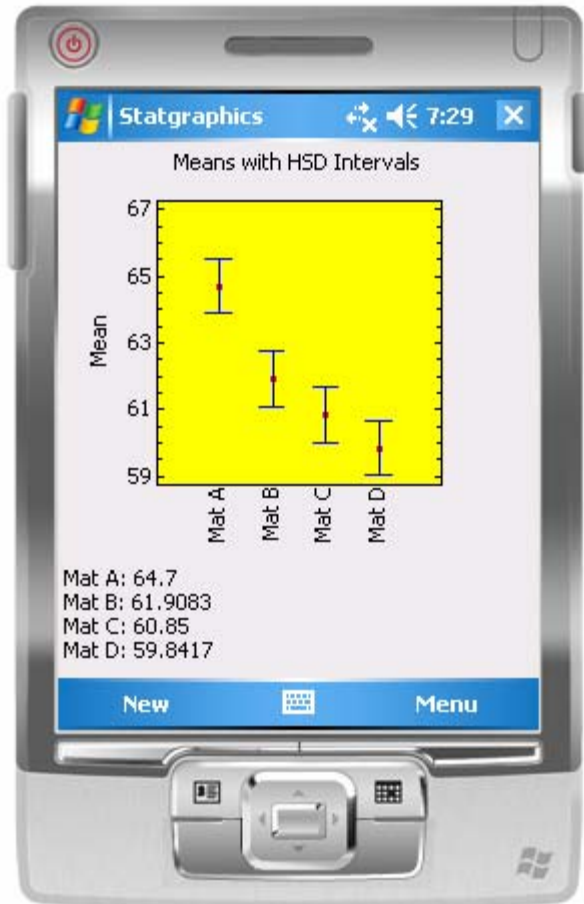
1. Access the *Properties* dialog box for the leftmost *Response* variable by double-clicking on its column header.
2. On the *Dist.* tab, select the assumed distribution. If you select *Lognormal*, the logarithms of the data will be assumed to follow a normal distribution. If you select *Power normal*, the data will be assumed to follow a normal distribution after raising them to the indicated *power*.



If a non-normal distribution is selected, the tests will be applied to the data in the transformed metric.

Means Plot

The *Means Plot* page displays the mean of each row as a point symbol, together with Tukey’s HSD (*Honestly Significant Difference*) intervals.



The point symbols are located at the sample means. The bars extend over uncertainty intervals, the type of which is determined by the settings on the *Comp.* tab of the leftmost *Response* column. Several different types of intervals may be constructed:

- **Standard errors (pooled s)** - displays the standard errors using the pooled within-sample standard deviation:

$$\bar{x}_j \pm \sqrt{\frac{MSE}{n_j}} \tag{7}$$

- **Standard errors (individual s)** - displays the standard errors using the standard deviation of each sample separately:

$$\bar{x}_j \pm \sqrt{\frac{s_j^2}{n_j}} \quad (8)$$

- **Confidence intervals (pooled s)** - displays confidence intervals for the group means using the pooled within-group standard deviation:

$$\bar{x}_j \pm t_{\alpha/2, n-m} \sqrt{\frac{MSE}{n_j}} \quad (9)$$

where n equals the total number of observations.

- **Confidence intervals (individual s)** - displays confidence intervals for the sample means using the standard deviation of each group separately:

$$\bar{x}_j \pm t_{\alpha/2, n_j-1} \sqrt{\frac{s_j^2}{n_j}} \quad (10)$$

- **LSD intervals** - designed to compare the means with the stated confidence level. The intervals are given by

$$\bar{x}_j \pm \frac{\sqrt{2}t_{\alpha/2, n-m}}{2} \sqrt{\frac{s_p^2}{n_j}} \quad (11)$$

where t represents the value of Student's t distribution with $n + q$ degrees of freedom leaving an area of $\alpha/2$ in the upper tail of the curve. If two sample sizes are equal, you can determine whether or not the population means are significantly different from each other by determining whether or not the intervals overlap.

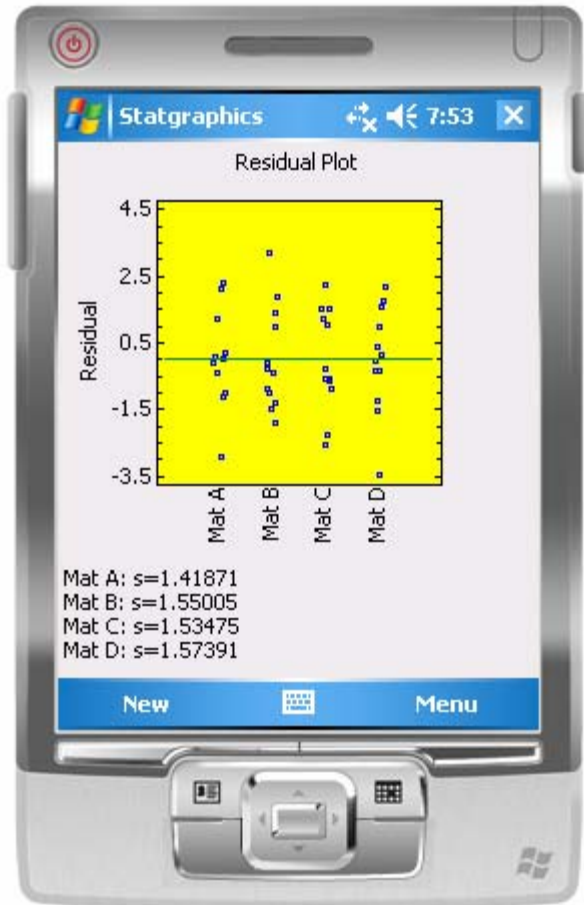
- **Tukey HSD Intervals** - designed for comparing multiple pairs of means. . The intervals are given by

$$\bar{x}_j \pm \frac{\sqrt{2}T_{\alpha/2, q, n-q}}{2} \sqrt{\frac{s_p^2}{n_j}} \quad (12)$$

which uses Tukey's T instead of Student's t . Tukey called his procedure the *Honestly Significant Difference* procedure since it allows all pairs of means to be compared while controlling the experiment-wide error rate at α . If all of the means are equal, the probability of incorrectly declaring *any* of the pairs to be significantly different in the entire experiment equals α . Tukey's procedure is more conservative than Fisher's LSD procedure, since it makes it harder to declare any particular pair of means to be significantly different.

Residuals

The *Residuals* plots the deviations of each observation from its respective group mean.



The plot may be used to help identify outliers, and to visualize any differences between the within-group variance. The sample standard deviation for up to 8 groups is displayed below the plot.