Chapter 240

Nondetects-Data Group Comparison

Introduction

This procedure computes summary statistics, generates EDF plots, and computes hypothesis tests appropriate for two or more groups for data with nondetects (left-censored) values. Following the recommendation of Helsel (2005), pp. 77-78, the methods for this procedure are valid only if fewer than 50% of the values are nondetects (left-censored).

Nondetects analysis is the analysis of data in which one or more of the values cannot be measured exactly because they fall below one or more detection limits. Detection limits often arise in environmental studies because of the inability of instruments to measure small concentrations. Some examples of sampling scenarios that lead to datasets with nondetects values are finding pesticide concentrations in water, determining chemical composition of soils, or establishing the number of particulates of a compound in the air.

A common practice for dealing with values which fall below the detection threshold is substitution. Often, each value which is below the detection limit is substituted with one half the detection limit. Summary statistics and comparisons are then carried out using standard techniques (means, confidence intervals, t-tests, ANOVA, etc.) with the substituted data. Helsel (2005) warns of the potential data analysis biases that result if nondetects values are substituted. He particularly warns about the arbitrariness of substituting one half the detection limit (or zero, or the detection limit). Alternatively, techniques based on survival analysis methods have been developed for appropriate use of the information contained in the nondetected observations. The general approach is to convert the nondetects data (left-censored) to survival data (right-censored), use the survival analysis techniques on the newly created survival data, and then convert the survival summaries back to original scale (In **NCSS**, these conversions are performed automatically). The resulting summary statistics and hypothesis tests are analogs to the common techniques, but which appropriately account for nondetected observations. For example, medians are used rather than means, EDF plots replace box plots and histograms, and logrank tests are used instead of two-sample t-tests and ANOVA.

The technical details of survival analysis are found in the Kaplan-Meier Survival Curves chapter. For a complete account of nondetects analysis, we suggest the book by Helsel (2005).

Technical Details

Flipping Constant

To convert nondetects data to the format of survival data, each response, including nondetected values, must be subtracted from a suitable flipping constant. The flipping constant can be any number which is larger than the maximum of the nondetects data. The resulting right-censored data are

$$Flip_i = M - x_i,$$

where M is the flipping constant and the x_i are the original observations.

For example, consider the first 10 of 25 dioxin concentrations (fg/cubic meter) with lower detection limit 50 fg/cubic meter (these data can be found in the DIOXIN dataset):

DIOXIN Dataset (Subset)

Dioxin
391
724
603
50
482
656
50
797
190
444

A suitable flipping constant is any value larger than the maximum value. Suppose M = 1000 is arbitrarily chosen as the flipping constant. The flipped data would then become

Dioxin	***M* – Dioxin**	**Flip**
391	1000 – 391	609
724	1000 – 724	276
603	1000 – 603	397
<50	1000 – <50	>950
482	1000 – 482	518
656	1000 – 656	344
<50	1000 – <50	>950
797	1000 – 797	203
190	1000 – 190	810
444	1000 – 444	556
.	.	.
.	.	.
.	.	.

The flipped data is now in the survival data format.

Once the data are converted to the survival data format, the nonparametric Kaplan-Meier methods can be used for estimating summary statistics (i.e., median, quantiles, standard errors, confidence limits), and for group comparisons. The summary statistics of location (i.e., median, quantiles, and confidence limits) are converted back to the original scale using the same flipping constant *M*. For example, to convert the median of the survival data to the median of the original units, the formula

$$Median = M - SurvivalMedian$$

is used. For the Dioxin data, the survival median (of the flipped data) is 556 fg/cubic meter. The median on the original scale would then be Median = 1000 – 556 = 444 fg/cubic meter. The standard error statistics for the flipped survival data are the same as those of the original scale and need not be converted. All of the calculations involving conversion and re-conversion based on the flipping constant are done automatically in **NCSS**.

The Empirical Distribution Function (EDF)

The empirical distribution function (EDF) provides an approximation of the true cumulative distribution function of the measured response. It is useful for viewing or obtaining sample percentiles (quantiles) for each of the observed responses. The EDF is produced using the Kaplan-Meier product-limit estimator (estimated survival distribution) of the flipped data. The resulting survival distribution is then converted to the EDF by re-subtracting all values from the flipping constant. We now examine the technical details of the estimation of the survival distribution.

Hypothesis Tests

This section presents methods for testing that the distribution functions of two or more populations are equal. The null hypothesis is that the distribution functions of all populations are equal at all values greater than the minimum observed value. The alternative hypothesis is that at least two of the distribution functions are different at some value greater than the observed minimum value.

Five different choices of tests are available in **NCSS** to test the above hypotheses. The tests differ in the manner in which different responses are weighted. The most commonly used test is the logrank test, which has equal weighting. The other four tests shift the heaviest weighting to the larger or smaller responses. Although five tests are displayed, only one should be used. Because of the different weighting patterns, they will often give quite different results. The test that will be used should be justified and designated before viewing the data or test results.

The following table describes the weighting scheme for each of these tests.

Test	Comments
Logrank	This is the most commonly used test and the one we recommend. Equal weights across all times are used.
Gehan	Places very heavy weight on large responses.
Tarone-Ware	Places heavy weight on small responses.
Peto-Peto	Places a little more weight on large responses.
Modified Peto-Peto	Places a little more weight on large responses.

Data Structure

Nondetects datasets are specified using up to four components: the response value (e.g., concentration or amount), an optional indicator of whether or not each observation was detected, an optional group specification, and an optional frequency (count) specification. If no detection indicator is included, all response values represent detected responses. If there is no group specification, a single group is assumed. If the frequency (count) variable is omitted, all counts are assumed to be one.

Sample Dataset

The table below shows a dataset (fictitious) reporting sediment arsenic concentrations for three different regions of a lake. A single sample was taken from each of twenty randomly selected locations of each region. In this dataset, the response is the concentration of arsenic in mg/Kg (dry weight). The instruments used in the study to determine arsenic concentration are unable to detect concentrations below 10 mg/Kg. A value of zero in the ANondet column indicates arsenic was detected. A value of one in the ANondet column indicates arsenic was not detected. These data are contained in the ARSENIC dataset.

Arsenic Dataset (Subset)

Arsenic	ANondet	Region
14	0	1
10	1	1
31	0	1
26	0	1
10	1	1
.	.	.
.	.	.
.	.	.
15	0	2
10	1	2
25	0	2
21	0	2
27	0	2
.	.	.
.	.	.
.	.	.
29	0	2
26	0	2
18	0	3
26	0	3
.	.	.
.	.	.
.	.	.

EDF Plot Format Window Options

This section describes the specific options available on the EDF Plot Format window, which is displayed when the EDF Plot Format button is clicked. Common options, such as axes, labels, legends, and titles are documented in the Graphics Components chapter.

EDF Plot Tab

EDF Line Section

You can modify the attributes of the EDF line using the options in this section.

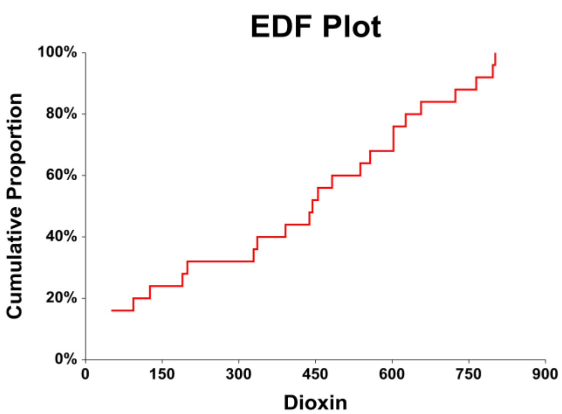

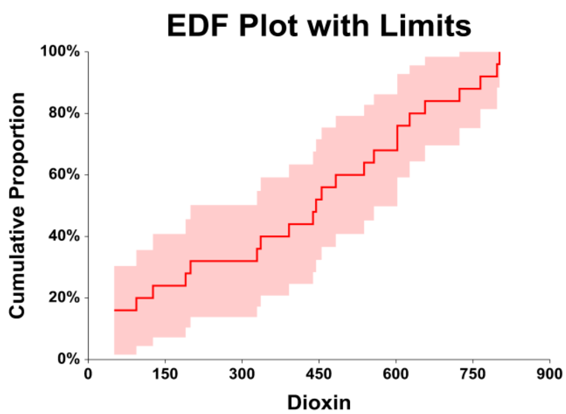

Symbols Section

You can modify the attributes of the EDF line using the options in this section.

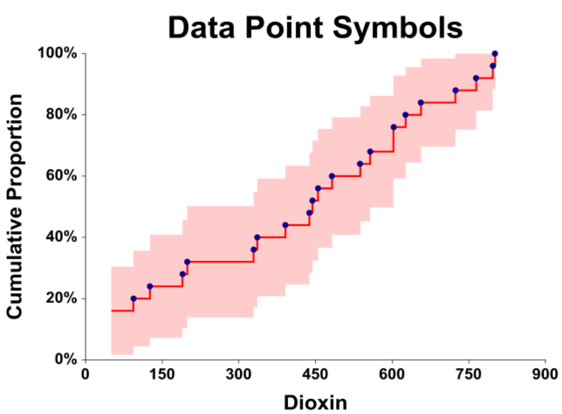

Reference Lines Tab

Lines from [Vertical] / [Horizontal] Axis Sections

You can modify the attributes of the EDF line using the options in this section.

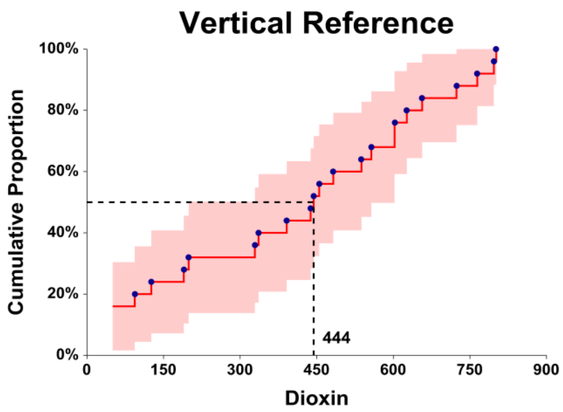

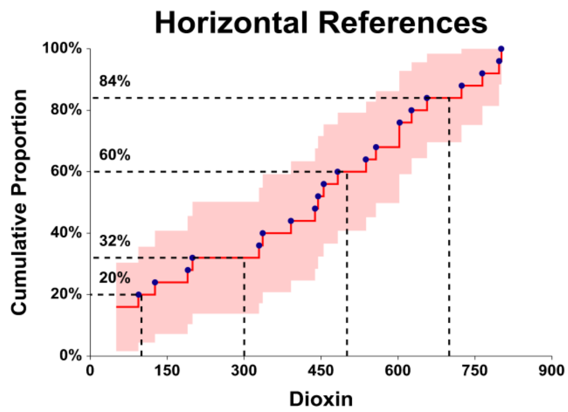

Titles, Legend, Numeric Axis, Group Axis, Grid Lines, and Background Tabs

Details on setting the options in these tabs are given in the Graphics Components chapter.

Example 1 – Analysis of Data with Nondetects

This section presents an example of how to analyze a typical set of nondetects data. Twenty-five air quality locations were randomly chosen to determine dioxin concentration (fg/cubic meter). The lower detection limit of the measurement instrument is 50 fg/cubic meter. Four of the 25 concentrations were not detected, and thus, are known only to be less than 50.

The data used are recorded in the Dioxin dataset.

Setup

To run this example, complete the following steps:

1 Open the Dioxin example dataset

- From the File menu of the NCSS Data window, select **Open Example Data**.
- Select **Dioxin** and click **OK**.

2 Specify the Nondetects-Data Group Comparison procedure options

- Find and open the **Nondetects-Data Group Comparison** procedure using the menus or the Procedure Navigator.
- The settings for this example are listed below and are stored in the **Example 1** settings file. To load these settings to the procedure window, click **Open Example Settings File** in the Help Center or File menu.

Variables Tab

Response Variable ....**Dioxin**
Nondetection (Censor) Variable......................**DNondet**
Detected **0**
Not Detected..**1**

Reports Tab

Specific Responses ....**100:500(100)**

3 Run the procedure

- Click the **Run** button to perform the calculations and generate the output.

Data Summary Section

Data Summary Section

Type	Rows	Count	Minimum	Maximum
Detected	21	21	94	801
Not Detected	4	4	50	50
Total	25	25	50	801

Data Summary Section: Response Quartiles

Quartile	Estimate	Lower 95.0% C.L.	Upper 95.0% C.L.
First (Q1)	190.000	50.000	438.000
Median (Q2)	444.000	199.000	603.000
Third (Q3)	603.000	455.000	724.000

This report displays a summary of the amount of data that were analyzed and the three quartiles. Scan this report to determine if there were any obvious data errors by double checking the counts and the minimum and maximum responses.

Specific Response Detail: Estimated Cumulative Proportion

Specific Response Detail: Estimated Cumulative Proportion

Response (R)	Cumulative Proportion P(R)	Standard Error of P(R)	Lower 95.0% C.L. for P(R)	Upper 95.0% C.L. for P(R)	Cum. Count
100.000	0.2000	0.0800	0.0432	0.3568	5
200.000	0.3200	0.0933	0.1371	0.5029	8
300.000	0.3200	0.0933	0.1371	0.5029	8
400.000	0.4400	0.0993	0.2454	0.6346	11
500.000	0.6000	0.0980	0.4080	0.7920	15

This report displays the Kaplan-Meier cumulative proportions at the specified responses. The standard error and confidence limits are also shown.

Response (R)

This is the specific response being reported on this line. The response values were specified in the Specific Responses box under the Reports tab.

Cumulative Proportion P(R)

This is the estimated proportion of responses less than the specified response (R).

Standard Error of P(R)

This is the estimated standard error, the square root of the variance estimate given by Greenwood's formula.

Lower and Upper Confidence Limits for S(T)

The lower and upper confidence limits provide a pointwise confidence interval for the cumulative proportion at each response. These limits are constructed so that the probability that the true proportion lies between them is $1 - \alpha$.

Three difference confidence intervals are available. All three confidence intervals perform similarly for large samples. The linear (Greenwood) interval is the most commonly used. However, the log-transformed and the arcsine-square intervals behave better in small to moderate samples, so they are recommended. The formulas for these limits are given in the Kaplan-Meier Survival Curves chapter and are not repeated here.

Cumulative Count

This value is the number of less than or equal to the specified response (R).

Quantiles of Responses

Quantiles of Responses

Proportion of Response	Estimated Quantile	Lower 95.0% C.L. Quantile	Upper 95.0% C.L. Quantile
0.0500		50.000	126.000
0.1000		50.000	190.000
0.1500		50.000	329.000
0.2000	126.000	50.000	336.000
0.2500	190.000	50.000	438.000
0.3000	199.000	50.000	444.000
0.3500	329.000	94.000	455.000
0.4000	391.000	126.000	482.000
0.4500	438.000	190.000	537.000
0.5000	444.000	199.000	603.000
0.5500	455.000	336.000	603.000
0.6000	537.000	391.000	626.000
0.6500	557.000	438.000	656.000
0.7000	603.000	444.000	724.000
0.7500	603.000	455.000	724.000
0.8000	656.000	537.000	764.000
0.8500	724.000	557.000	797.000
0.9000	764.000	603.000	801.000
0.9500	797.000	626.000	801.000

This report displays the estimated quantiles for various response proportions. For example, it gives the median response if it can be estimated.

Proportion of Response

This is the response proportion that is reported on this line. The proportion values were specified in the Quantiles box under the Reports tab.

Estimated Quantile

This is the response value corresponding to the response proportion. For example, this table estimates that 65% of the concentrations are less than or equal to 557 fg/m^3.

Lower and Upper Confidence Limits on Quantiles

These values provide a pointwise $100(1-\alpha)\%$ confidence interval for the estimated quantiles. For example, if the proportion of response 0.50, this provides a confidence interval for the median survival time.

Three methods are available for calculating these confidence limits. The method is designated under the Variables tab in the Confidence Limits box. The formulas for these confidence limits are given in the Kaplan-Meier Survival Curves chapter and are not repeated here.

Because of censoring, estimates and confidence limits are not available for all response proportions.

Response Detail

Response Detail

Response (R)	Cumulative Proportion P(R)	Standard Error of P(R)	Lower 95.0% C.L. for P(R)	Upper 95.0% C.L. for P(R)	Cum. Count	Count
<50.000					4	4
94.000	0.1600	0.0733	0.0163	0.3037	5	1
126.000	0.2000	0.0800	0.0432	0.3568	6	1
190.000	0.2400	0.0854	0.0726	0.4074	7	1
199.000	0.2800	0.0898	0.1040	0.4560	8	1
329.000	0.3200	0.0933	0.1371	0.5029	9	1
336.000	0.3600	0.0960	0.1718	0.5482	10	1
391.000	0.4000	0.0980	0.2080	0.5920	11	1
438.000	0.4400	0.0993	0.2454	0.6346	12	1
444.000	0.4800	0.0999	0.2842	0.6758	13	1
455.000	0.5200	0.0999	0.3242	0.7158	14	1
482.000	0.5600	0.0993	0.3654	0.7546	15	1
537.000	0.6000	0.0980	0.4080	0.7920	16	1
557.000	0.6400	0.0960	0.4518	0.8282	17	1
603.000	0.6800	0.0933	0.4971	0.8629	19	2
626.000	0.7600	0.0854	0.5926	0.9274	20	1
656.000	0.8000	0.0800	0.6432	0.9568	21	1
724.000	0.8400	0.0733	0.6963	0.9837	22	1
764.000	0.8800	0.0650	0.7526	1.0000	23	1
797.000	0.9200	0.0543	0.8137	1.0000	24	1
801.000	0.9600	0.0392	0.8832	1.0000	25	1

This report displays the Kaplan-Meier product-limit distribution values along with confidence limits. The formulas used are given in the Kaplan-Meier Survival Curves chapter.

Response (R)

This is the response being reported on this line. The response are the unique responses that occurred in the data.

Note that observations which are nondetects are marked with a less than sign (<). Estimated proportions are not calculated for nondetects observations.

Cumulative Proportion P(R)

This is the estimated proportion of responses less than the response (R).

Standard Error of S(T)

This is the estimated standard error, the square root of the variance estimate given by Greenwood's formula.

Lower and Upper Confidence Limits for S(T)

The lower and upper confidence limits provide a pointwise confidence interval for the cumulative proportion at each response. These limits are constructed so that the probability that the true proportion lies between them is $1 - \alpha$.

Three difference confidence intervals are available. All three confidence intervals perform similarly for large samples. The linear (Greenwood) interval is the most commonly used. However, the log-transformed and the arcsine-square intervals behave better in small to moderate samples, so they are recommended. The formulas for these limits are given in the Kaplan-Meier Survival Curves chapter and are not repeated here.

Cumulative Count

This value is the number of less than or equal to the specified response (R).

Count

This is the number of observations with this specific response value.

EDF Plot

EDF Plot(s)

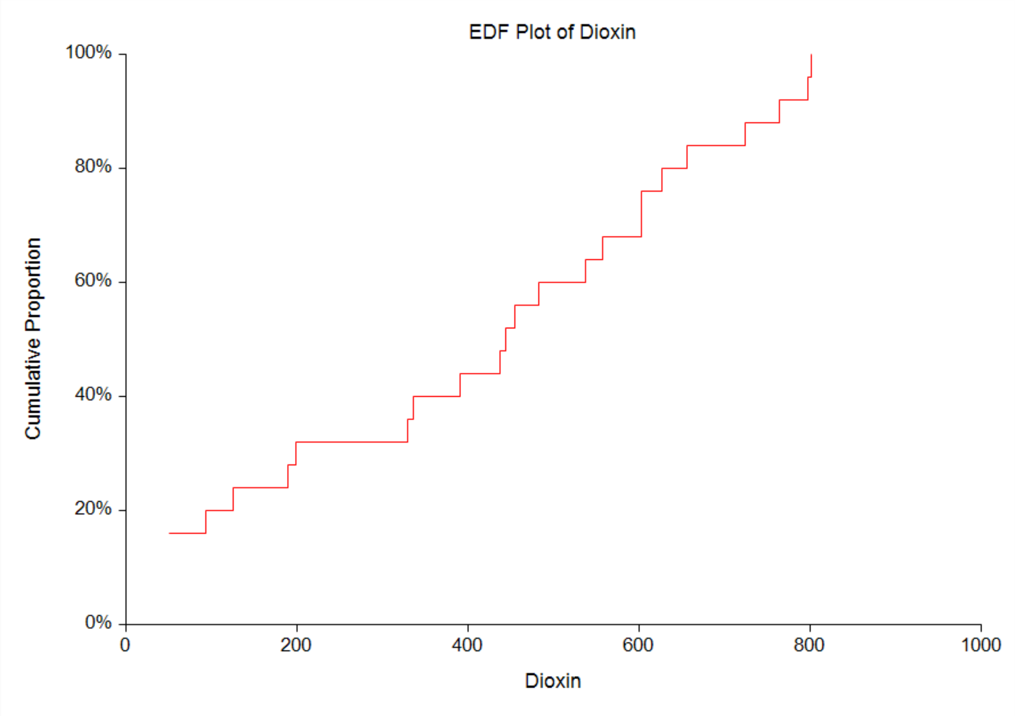

This plot shows the empirical distribution function (EDF). If there are several groups, a separate line is drawn for each group.

Example 2 – Group Comparisons with Nondetects

The research purpose of this example is comparing sediment arsenic concentrations for three different regions of a lake. A single sample was taken from each of twenty randomly selected locations of each region. The response is the concentration of arsenic in mg/Kg (dry weight). The instruments used in the study to determine arsenic concentration are unable to detect concentrations below 10 mg/Kg.

The data used are recorded in the variables Arsenic, ANondet, and Region of the Arsenic dataset.

Setup

To run this example, complete the following steps:

1 Open the Arsenic example dataset

- From the File menu of the NCSS Data window, select **Open Example Data**.
- Select **Arsenic** and click **OK**.

2 Specify the Nondetects-Data Group Comparison procedure options

- Find and open the **Nondetects-Data Group Comparison** procedure using the menus or the Procedure Navigator.
- The settings for this example are listed below and are stored in the **Example 2** settings file. To load these settings to the procedure window, click **Open Example Settings File** in the Help Center or File menu.

Variables Tab

Response Variable ....**Arsenic**
Nondetection (Censor) Variable......................**ANondet**
Group Variable...**Region**

Reports Tab

Logrank Test Summary ....**Checked**
Logrank Test Detail..**Checked**

3 Run the procedure

- Click the **Run** button to perform the calculations and generate the output.

Logrank Tests Section

Logrank Tests Section

Hypotheses
H0: Distribution Functions are Equal Among Groups
HA: At Least One Group Distribution Functions Differs

Test Name	Chi-Square	DF	Prob Level	Reject H0 (Alpha = 0.05)
Logrank	26.680	2	0.0000	Yes
Gehan-Wilcoxon	35.265	2	0.0000	Yes
Tarone-Ware	32.241	2	0.0000	Yes
Peto-Peto	35.479	2	0.0000	Yes
Mod. Peto-Peto	35.589	2	0.0000	Yes

Multiple Pairwise Tests Section

Hypotheses
H0: Distribution Functions are Equal
HA: Distribution Functions Differ

Test Name	Chi-Square	DF	Prob Level	Reject H0 (Alpha =0.05)	Bonferroni Adjusted Prob Level	Reject H0 (Alpha =0.05)
Group Pair Tested: 1 vs. 2						
Logrank	0.374	1	0.5409	No	1.0000	No
Gehan-Wilcoxon	0.326	1	0.5683	No	1.0000	No
Tarone-Ware	0.389	1	0.5327	No	1.0000	No
Peto-Peto	0.267	1	0.6055	No	1.0000	No
Mod. Peto-Peto	0.265	1	0.6069	No	1.0000	No
Group Pair Tested: 1 vs. 3						
Logrank	16.239	1	0.0001	Yes	0.0002	Yes
Gehan-Wilcoxon	19.657	1	0.0000	Yes	0.0000	Yes
Tarone-Ware	18.787	1	0.0000	Yes	0.0000	Yes
Peto-Peto	19.418	1	0.0000	Yes	0.0000	Yes
Mod. Peto-Peto	19.457	1	0.0000	Yes	0.0000	Yes
Group Pair Tested: 2 vs. 3						
Logrank	15.978	1	0.0001	Yes	0.0002	Yes
Gehan-Wilcoxon	20.474	1	0.0000	Yes	0.0000	Yes
Tarone-Ware	19.109	1	0.0000	Yes	0.0000	Yes
Peto-Peto	20.391	1	0.0000	Yes	0.0000	Yes
Mod. Peto-Peto	20.453	1	0.0000	Yes	0.0000	Yes

Notes:
The most commonly used test is the Logrank test.

This report gives the results of the five logrank type tests that are provided by this procedure. We strongly suggest that you select the test that will be used before viewing this report. We recommend the Logrank test.

The tests are divided into two groups: overall tests and pairwise tests. The overall tests test for significant differences between groups, but do not indicate which groups are different from each other. The pairwise tests indicate which groups have significantly different distribution functions. Adjusted probability levels should be used to account for multiplicity of tests.

Chi-Square

This is the chi-square value of the test. Each of these tests is approximately distributed as a chi-square in large samples.

DF

This is the degrees of freedom of the chi-square distribution associated with each test. It is one less than the number of groups being compared in a particular test.

Prob Level

This is the significance level of the test. If this value is less than than chosen significance level (often 0.05), the test is significant, indicating evidence of a difference in distribution functions. For pairwise tests the Bonferroni adjusted probability level should be used to account for multiple testing.

Reject H0

This is an indicator based on the comparison of the probability level to the specified alpha. 'Yes' indicates rejection of the null hypothesis (evidence that the true distribution functions are different). 'No' indicates the null hypothesis should not be rejected (not sufficient evidence that the true distribution functions are different).

Bonferroni Adjusted Prob Level

When more than two groups are compared, the number of pairwise comparisons is greater than one. Bonferroni adjusted probability levels account for the multiplicity of hypothesis tests. The Bonferroni adjustment to the probability level is made by multiplying the given probability level by the number of tests that are performed (with a ceiling of 1.0). In this example, three pairwise comparisons are made. Thus, each probability level is multiplied by three. Any adjusted probability level greater than one is set to one. The Bonferroni adjusted probability level for the last two longrank tests in this example appears to be only two times the base probability level. This is due to rounding. If more decimal places are specified, it is seen that the adjusted probability levels are three times the base probability levels.

Logrank Test Detail Section

Logrank Test Detail Section

Group	Z-Value	Standard Error	Standardized Z-Value
1	-7.561	3.398	-2.225
2	-4.484	3.380	-1.327
3	12.044	2.340	5.146

Result

Probability Level = 0.0000

Gehan-Wilcoxon Test Detail Section

Group	Z-Value	Standard Error	Standardized Z-Value
1	-349.000	132.199	-2.640
2	-270.000	132.219	-2.042
3	619.000	104.394	5.929

Result

Probability Level = 0.0000

Tarone-Ware Test Detail Section

Group	Z-Value	Standard Error	Standardized Z-Value
1	-51.277	20.460	-2.506
2	-35.076	20.428	-1.717
3	86.353	15.249	5.663

Result

Probability Level = 0.0000

Peto-Peto Test Detail Section

Group	Z-Value	Standard Error	Standardized Z-Value
1	-5.568	2.114	-2.634
2	-4.453	2.114	-2.106
3	10.021	1.684	5.949

Result

Probability Level = 0.0000

Mod. Peto-Peto Test Detail Section

Group	Z-Value	Standard Error	Standardized Z-Value
1	-5.452	2.065	-2.640
2	-4.377	2.066	-2.119
3	9.830	1.650	5.959

Result
Probability Level = 0.0000

This report gives the details of each of the five logrank tests that are provided by this procedure. We strongly suggest that you select the test that will be used before viewing this report. We recommend that you use the Logrank test.

Group

This is the group reported on this line.

Z-Value

The details of the z-value are given in the Kaplan-Meier Survival Curves chapter and are not repeated here.

Standard Error

This is the standard error of the above z-value. It is used to standardize the z-values.

Standardized Z-Value

The standardized z-value is created by dividing the z-value by its standard error. This provides an index number that will usually very between -3 and 3. Extreme values represent groups that are quite different from the typical group, at least at some response values.

Example 3 – Validation of Summary Statistics using Helsel (2005)

This section presents validation of nondetects analysis summary statistics. Helsel (2005) presents an example on pages 103-113 involving lead concentrations. These data are contained in the Lead dataset.

On page 108, Helsel (2005) finds the median to be 1 – 0.984483 = 0.015517. The first and third quartiles are 1 – 0.985714 = 0.014286 and 1 – 0.975472 = 0.024528, respectively. The cumulative proportion for a lead concentration of 0.034 is 0.777778. The (B-C Sign) 95% confidence interval for the median lead concentration is presented on page 112 as (0.014, 0.019).

Setup

To run this example, complete the following steps:

1 Open the Lead example dataset

- From the File menu of the NCSS Data window, select **Open Example Data**.
- Select **Lead** and click **OK**.

2 Specify the Nondetects-Data Group Comparison procedure options

- Find and open the **Nondetects-Data Group Comparison** procedure using the menus or the Procedure Navigator.
- The settings for this example are listed below and are stored in the **Example 3** settings file. To load these settings to the procedure window, click **Open Example Settings File** in the Help Center or File menu.

Variables Tab

Response Variable ....**Lead**
Nondetection (Censor) Variable......................**LNondet**

Reports Tab

Data Summary Section...................................**Checked**
Specific Response Detail...............................**Unchecked**
Quantiles of Responses.................................**Unchecked**
Response Detail ....**Checked**
Logrank Test Summary ....**Unchecked**
Logrank Test Detail..**Unchecked**
Response **6**

3 Run the procedure

- Click the **Run** button to perform the calculations and generate the output.

Output

Data Summary Section

Type	Rows	Count	Minimum	Maximum
Detected	12	12	0.01372549	0.2689655
Not Detected	15	15	0.02	0.02
Total	27	27	0.01372549	0.2689655

Data Summary Section: Response Quartiles

Quartile	Estimate	Lower 95.0% C.L.	Upper 95.0% C.L.
First (Q1)	**0.014286**	0.013725	0.018644
Median (Q2)	**0.015517**	**0.014286**	**0.018644**
Third (Q3)	**0.024528**	0.015517	0.106061

Response Detail

Response (R)	Cumulative Proportion P(R)	Standard Error of P(R)	Lower 95.0% C.L. for P(R)	Upper 95.0% C.L. for P(R)	Cum. Count	Count
0.013725	0.0000				1	1
0.014286	0.1759	0.1539	0.0000	0.4776	2	1
0.015517	0.3519	0.1813	0.0000	0.7073	3	1
0.018644	0.5278	0.1660	0.2024	0.8531	4	1
<0.020000					19	15
0.023529	0.7037	0.0879	0.5315	0.8759	20	1
0.024528	0.7407	0.0843	0.5754	0.9060	21	1
0.033962	**0.7778**	0.0800	0.6210	0.9346	22	1
0.049153	0.8148	0.0748	0.6683	0.9613	23	1
0.106061	0.8519	0.0684	0.7179	0.9858	24	1
0.174074	0.8889	0.0605	0.7703	1.0000	25	1
0.177049	0.9259	0.0504	0.8271	1.0000	26	1
0.268966	0.9630	0.0363	0.8917	1.0000	27	1

You can check this table to see that the results are the same as those of Helsel (2005).

Example 4 – Validation of Group Comparison Statistics using Helsel (2005)

This section presents validation of the group comparison statistics. Helsel (2005) presents an example of results for comparing concentrations among three groups. These data are contained in the Concentration dataset.

The results for the overall test for determining difference in concentration patterns across groups is found on page 180. The log rank test results in a chi-square statistic of 16.2794 with probability level 0.000. The Gehan (Wilcoxon) test gives a chi-square statistic of 16.0761 with probability level 0.000. The results of the individual group comparison Gehan (Wilcoxon) tests are given on page 181. For comparing the low group to the medium group, the chi-square value is 0.68890 with probability level 0.407. For comparing the low group to the high group, the chi-square value is 7.09906 with probability level 0.008. For comparing the medium group to the high group, the chi-square value is 11.5275 with probability level 0.001.

Setup

To run this example, complete the following steps:

1 Open the Concentration example dataset

- From the File menu of the NCSS Data window, select **Open Example Data**.
- Select **Concentration** and click **OK**.

2 Specify the Nondetects-Data Group Comparison procedure options

- Find and open the **Nondetects-Data Group Comparison** procedure using the menus or the Procedure Navigator.
- The settings for this example are listed below and are stored in the **Example 4** settings file. To load these settings to the procedure window, click **Open Example Settings File** in the Help Center or File menu.

Variables Tab

Response Variable ....**Conc**
Nondetection (Censor) Variable......................**CNondet**
Group Variable...**Group**

Reports Tab

Data Summary Section...................................**Unchecked**
Specific Response Detail...............................**Unchecked**
Quantiles of Responses.................................**Unchecked**
Response Detail ....**Unchecked**
Logrank Test Summary ....**Checked**
Logrank Test Detail..**Unchecked**

3 Run the procedure

- Click the **Run** button to perform the calculations and generate the output.

Output

Logrank Tests Section

Hypotheses
H0: Distribution Functions are Equal Among Groups
HA: At Least One Group Distribution Functions Differs

Test Name	Chi-Square	DF	Prob Level	Reject H0 (Alpha = 0.05)
Logrank	**16.280**	**2**	**0.0003**	**Yes**
Gehan-Wilcoxon	**16.076**	**2**	**0.0003**	**Yes**
Tarone-Ware	16.669	2	0.0002	Yes
Peto-Peto	16.359	2	0.0003	Yes
Mod. Peto-Peto	16.369	2	0.0003	Yes

Multiple Pairwise Tests Section

Hypotheses
H0: Distribution Functions are Equal
HA: Distribution Functions Differ

Test Name	Chi-Square	DF	Prob Level	Reject H0 (Alpha =0.05)	Bonferroni Adjusted Prob Level	Reject H0 (Alpha =0.05)
Group Pair Tested: High vs. Low						
Logrank	7.360	1	0.0067	Yes	0.0200	Yes
Gehan-Wilcoxon	**7.099**	**1**	**0.0077**	**Yes**	**0.0231**	**Yes**
Tarone-Ware	7.282	1	0.0070	Yes	0.0209	Yes
Peto-Peto	7.385	1	0.0066	Yes	0.0197	Yes
Mod. Peto-Peto	7.378	1	0.0066	Yes	0.0198	Yes
Group Pair Tested: High vs. Medium						
Logrank	11.398	1	0.0007	Yes	0.0022	Yes
Gehan-Wilcoxon	**11.528**	**1**	**0.0007**	**Yes**	**0.0021**	**Yes**
Tarone-Ware	11.931	1	0.0006	Yes	0.0017	Yes
Peto-Peto	11.454	1	0.0007	Yes	0.0021	Yes
Mod. Peto-Peto	11.470	1	0.0007	Yes	0.0021	Yes
Group Pair Tested: Low vs. Medium						
Logrank	1.125	1	0.2888	No	0.8663	No
Gehan-Wilcoxon	**0.689**	**1**	**0.4065**	**No**	**1.0000**	**No**
Tarone-Ware	0.796	1	0.3723	No	1.0000	No
Peto-Peto	1.109	1	0.2923	No	0.8769	No
Mod. Peto-Peto	1.092	1	0.2961	No	0.8884	No

Notes:
The most commonly used test is the Logrank test.

You can check this table to see that the results are the same as those of Helsel (2005).