Chapter 8
Hypothesis Tests

Chapter Table of Contents

Chapter 8
Hypothesis Tests

Introduction

Hypothesis tests are frequently performed for one and two samples. For one sample, you are often interested in whether a population characteristic such as the mean is equivalent to a certain value. For two samples, you may be interested in whether the true means are different. When you have paired data, you may be interested in whether the mean difference is zero.

Statistical hypothesis tests depend on a statistic designed to measure the degree of evidence for various alternative hypotheses. You compute the value of the statistic for your sample. If the value is improbable under the hypothesis you want to test, then you reject the hypothesis.

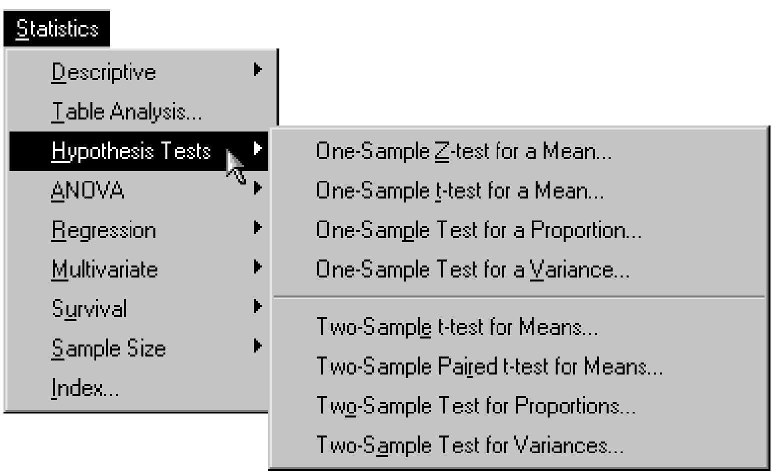

Figure 8.1. Hypothesis Tests Menu

The Analyst Application enables you to perform hypothesis tests for means, proportions, and variances for either one or two samples.

The examples in this chapter demonstrate how you can use the Analyst Application to perform a one-sample *t*-test, a paired *t*-test, a two-sample test for proportions, and a two-sample test for variances. Additionally, the section “Discussion of Other Tests” on page 177 provides information on other hypothesis tests you can perform with the Analyst Application.

One-Sample t-Test

The One-Sample *t*-Test task enables you to test whether the mean of a variable is less than, greater than, or equal to a specific value. The observed mean of the variable is compared to this value.

The data set analyzed in the following example, Bthdth92, is taken from the 1995 Statistical Abstract of the United States, and it contains measures of the birth rate and infant mortality rate for 1992 in the United States. Information is provided for the 50 states and the District of Columbia, grouped by region.

Suppose you want to determine whether the average infant mortality rate in the United States is equal to a specific value. Note that the one-sample *t*-test is appropriate in this situation because the standard deviation of the population from which the data arise is unknown. When you know the standard deviation of the population, use the One-Sample Z-Test for a Mean task (see the section “Discussion of Other Tests” on page 177 for more information).

Open the Bthdth92 Data Set

The data are provided in the Analyst Sample Library. To access this Analyst sample data set, follow these steps:

1. Select **Tools** → **Sample Data** . . .
2. Select Bthdth92.
3. Click **OK** to create the sample data set in your Sasuser directory.
4. Select **File** → **Open By SAS Name** . . .

5. Select Sasuser from the list of **Libraries**.
6. Select Bthdth92 from the list of members.
7. Click **OK** to bring the Bthdth92 data set into the data table.

Request a One-Sample t-Test

To test whether the average infant mortality rate is equal to 8, follow these steps:

1. Select **Statistics** → **Hypothesis Tests** → **One-Sample t-Test for a Mean** . . .
2. Select death as the variable to be analyzed.
3. Enter **8** in the box labeled **Null: Mean =** and press **Enter**.

Your alternative hypothesis can be that the mean is less than, greater than, or not equal to a specified value. In this example, the alternative hypothesis is that the mean of the variable death is not equal to 8.

In Figure 8.2, the one-sample *t*-test dialog defines the null and alternative hypotheses and specifies death as the variable to be tested.

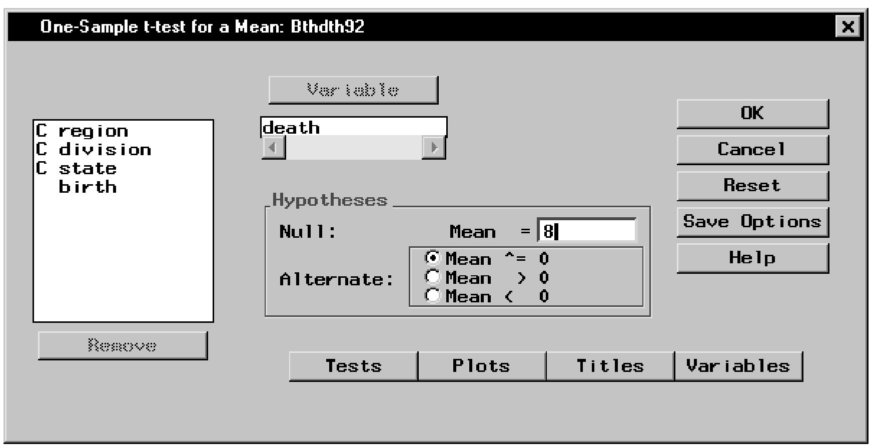

Figure 8.2. One-Sample t-Test Dialog

The default one-sample *t*-test task includes sample statistics for the variable death and the hypothesis test results.

Compute a Confidence Interval for the Mean

To produce a confidence interval for the mean in addition to the hypothesis test, follow these steps:

1. Click on the **Tests** button in the main dialog.
2. Select **Interval** to request a two-sided confidence interval for the mean.

You can choose either a one-sided or a two-sided confidence interval for the mean. The selections **Lower bound** and **Upper bound** specify one-sided confidence bounds.

The default confidence level is 95%. You can click on the down arrow to select another confidence level, or you can enter a confidence level in the box.

3. Click **OK** to return to the main dialog.

Figure 8.3 displays the selection of a 95% two-sided confidence interval for the mean. Note that you can also request a retrospective power analysis of the test in the **Power Analysis** tab.

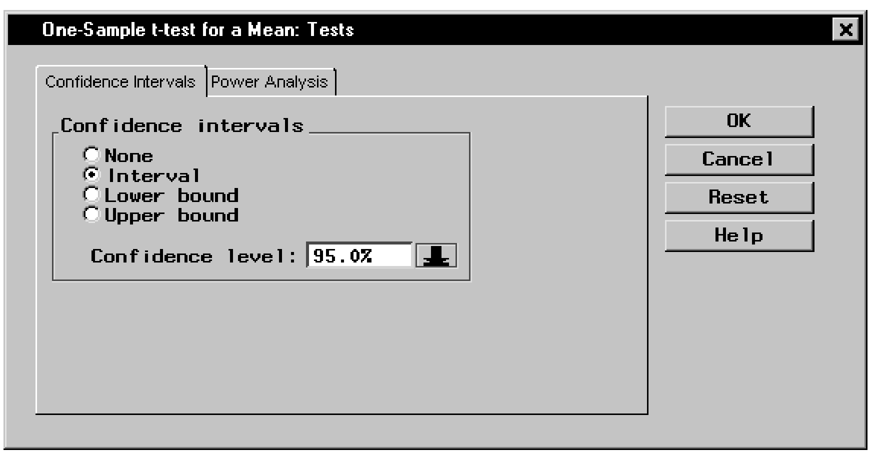

Figure 8.3. One-Sample t-Test: Tests Dialog

Request a t Distribution Plot

To request a *t* distribution plot in addition to the hypothesis test, follow these steps:

1. Click on the **Plots** button in the main dialog.
2. Select **t distribution plot**.
3. Click **OK** to return to the main dialog.

Figure 8.4 displays the Plots dialog with **t distribution plot** selected.

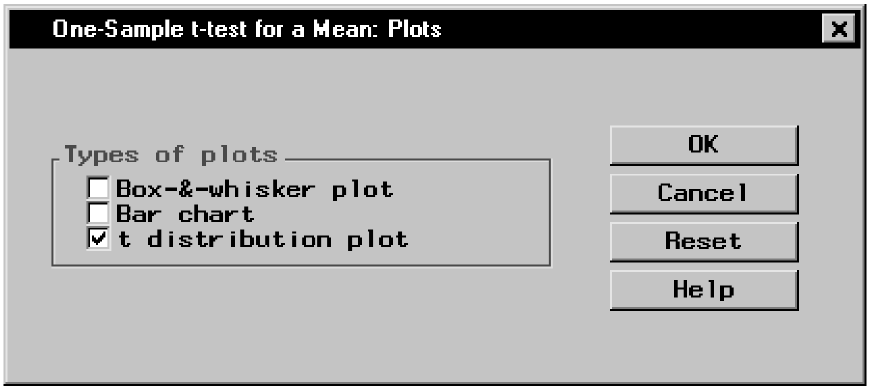

Figure 8.4. One-Sample t-Test: Plots Dialog

Click **OK** in the main dialog to perform the analysis.

Review the Results

The results of the hypothesis test are displayed in Figure 8.5. The output includes the "Sample Statistics" table for the variable death, the hypothesis test results, and the 95% confidence interval for the mean.

The mean of the variable death is 8.61, which is greater than the specified test value of 8.

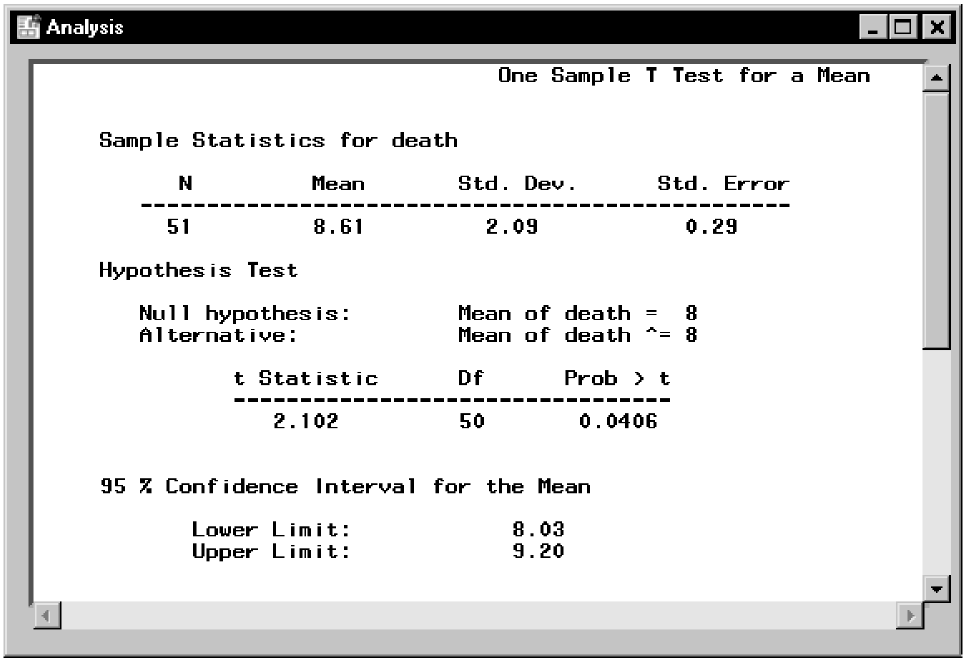

```
Analysis

                                   One Sample T Test for a Mean


   Sample Statistics for death

        N         Mean       Std. Dev.      Std. Error
      -----------------------------------------------
       51         8.61         2.09           0.29

   Hypothesis Test

      Null hypothesis:        Mean of death =  8
      Alternative:            Mean of death ^= 8

             t Statistic      Df      Prob > t
             ---------------------------------
                2.102         50       0.0406


   95 % Confidence Interval for the Mean

          Lower Limit:            8.03
          Upper Limit:            9.20
```

Figure 8.5. One-Sample t-Test: Output

The *t* statistic of 2.102 and the associated *p*-value (0.0406) provide evidence at the $\alpha = 0.05$ level that the average infant mortality rate is not equal to 8. The confidence interval indicates that you can be 95% confident that the true mean of the variable lies within the interval $[8.03, 9.20]$.

The requested *t* distribution plot is displayed in Figure 8.6. The plot depicts the calculated *t* statistic superimposed on a *t* distribution density function with 50 degrees of freedom.

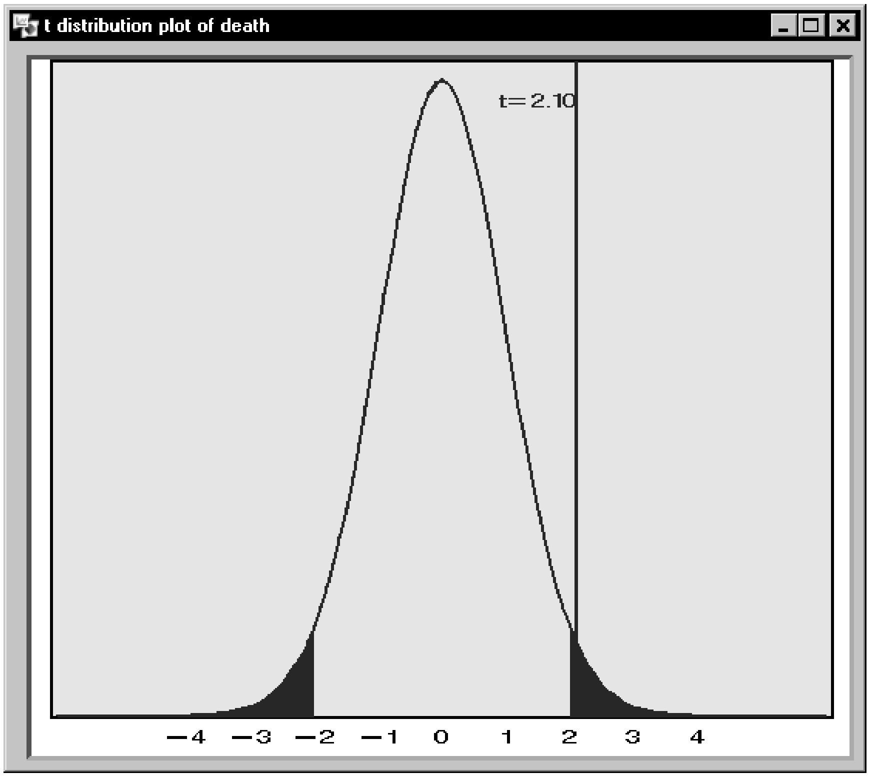

Figure 8.6. One-Sample t-Test: t Distribution Plot

Because this analysis requests a two-tailed test, two critical regions are shaded, one in each of the left and right tails. The alpha level for the test is 0.05; thus, each region represents 2.5% of the area under the curve. In a one-tailed test at the $\alpha = 0.05$ level, the critical region appears in one tail only, and it represents 5% of the area under the curve.

Here, the *t* statistic falls in the shaded region. Thus, the null hypothesis is rejected.

Paired t-test

The Paired *t*-test enables you to determine whether the means of paired samples are equal. The term *paired* means that there is a correspondence between observations from each population. For example, the birth and death data analyzed in the preceding section are considered to be paired data because, in each observation, the variables birth and death correspond to the same state.

Suppose that you want to determine whether the means for the birth rate and the infant mortality rate are equal. Analyst provides the Two-Sample Paired *t*-test for Means task, which tests the equality of means of two paired samples. The two samples in this example are the birth rate (birth) and the infant mortality rate (death) for each state.

Open the Bthdth92 Data Set

The data are provided in the Analyst Sample Library. To access this Analyst sample data set, follow these steps:

1. Select **Tools** → **Sample Data** . . .
2. Select Bthdth92.
3. Click **OK** to create the sample data set in your Sasuser directory.
4. Select **File** → **Open By SAS Name** . . .
5. Select Sasuser from the list of **Libraries**.
6. Select Bthdth92 from the list of members.
7. Click **OK** to bring the Bthdth92 data set into the data table.

Request a Paired t-Test

To perform this analysis, follow these steps:

1. Select **Statistics** → **Hypothesis Tests** →
   Two-Sample Paired t-test for Means . . .
2. Select the variable birth as the Group 1 variable.
3. Select the variable death as the Group 2 variable.

The test of interest is whether the difference of the means is zero. This is the default value in Analyst, although you can specify other values as well.

You can choose one of three alternative hypotheses. The default is that the difference between the means is not equal to the specified difference, which is the two-sided alternative. The one-sided alternatives are that the difference is greater than, or less than, the difference specified in the null hypothesis.

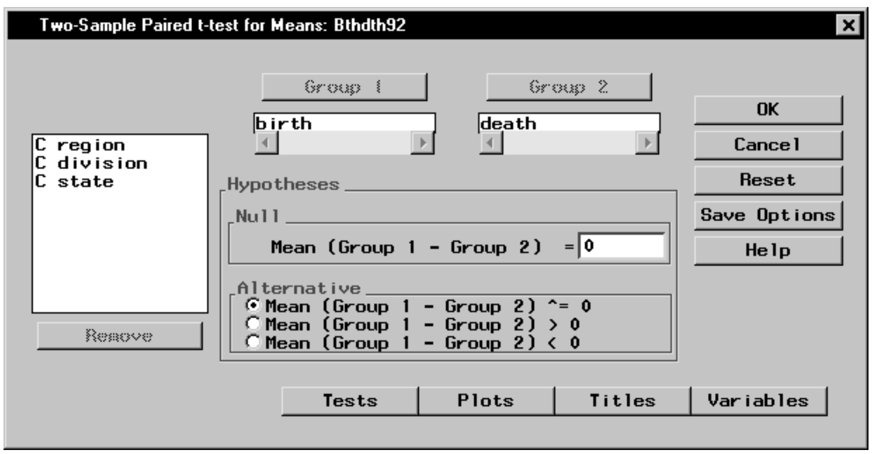

Figure 8.7. Paired t-test Dialog

In Figure 8.7, the null hypothesis specifies that the means of the variables birth and death are equal (or, equivalently, that the difference between the means is 0). The alternative hypothesis is that the two means are not equal.

Request Plots

To specify a box-and-whisker plot and a means plot in addition to the hypothesis test, follow these steps:

1. Click on the **Plots** button in the main dialog.
2. Select **Box-&-whisker plot**.
3. Select **Means plot**.
4. Click **OK**.

Figure 8.8 displays the Plots dialog with **Box-&-whisker plot** and **Means plot** selected.

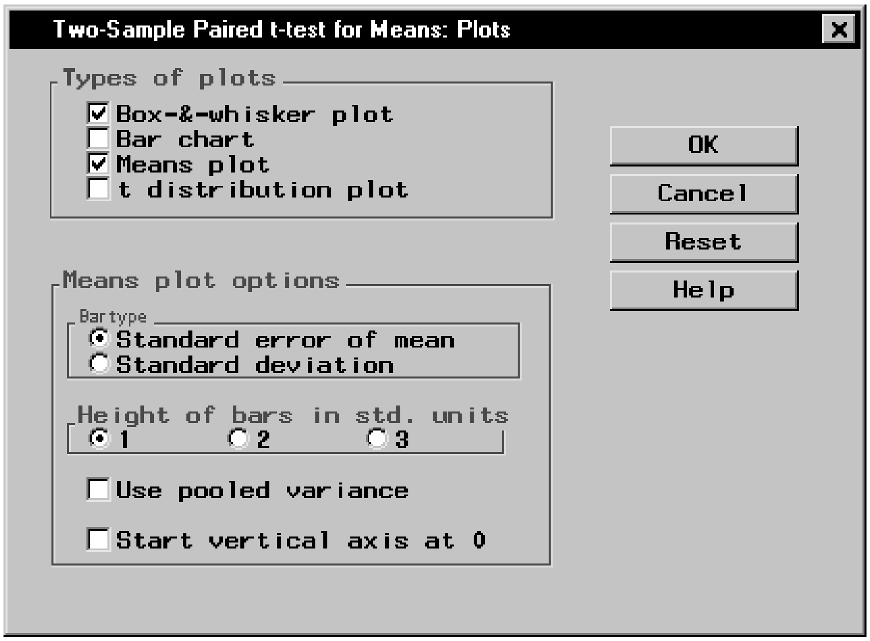

Figure 8.8. Paired t-test: Plots Dialog

Click **OK** in the main dialog to perform the analysis.

Review the Results

The results of the analysis, displayed in Figure 8.9, contain the mean, standard deviation, and standard error of the mean for both variables. The “Hypothesis Test” table provides the observed *t* statistic, the degrees of freedom, and the associated *p*-value of the test.

Analysis

Two Sample Paired t-test for the Means of birth and death

Sample Statistics

Group	N	Mean	Std. Dev.	Std. Error
birth	51	15.48431	1.7202	0.2409
death	51	8.613725	2.0851	0.292

Hypothesis Test

Null hypothesis: Mean of (birth - death) = 0
Alternative: Mean of (birth - death) ^= 0

t Statistic	Df	Prob > t
19.926	50	<.0001

Figure 8.9. Paired t-test: Results

In Figure 8.9, the “Sample Statistics” table shows that the mean of the variable birth is larger than that of the variable death. In the “Hypothesis Test” table, the *t* statistic (19.926) and associated *p*-value (< 0.0001) indicate that the difference between the two means is statistically very significant.

Figure 8.10 displays the side-by-side box plots of birth and death. Observations that fall beyond the whiskers are individually identified with a square symbol.

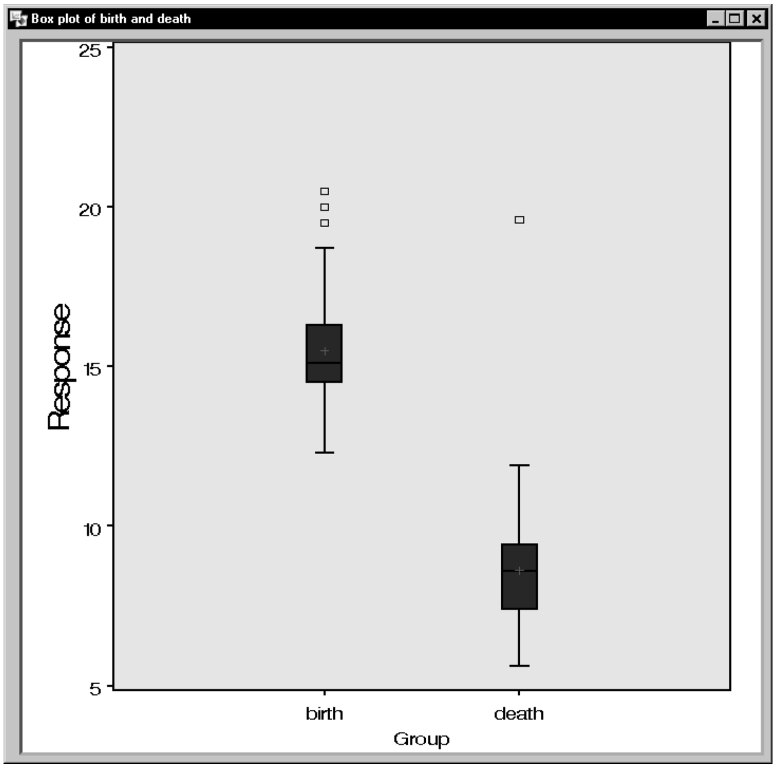

Figure 8.10. Paired t-test: Box-and-Whisker Plot

The means and standard error plot displayed in Figure 8.11 provides another view of the two variables. The means plot depicts an interval centered on the sample mean for each variable. The vertical line interval extends two standard deviations on either side of the mean.

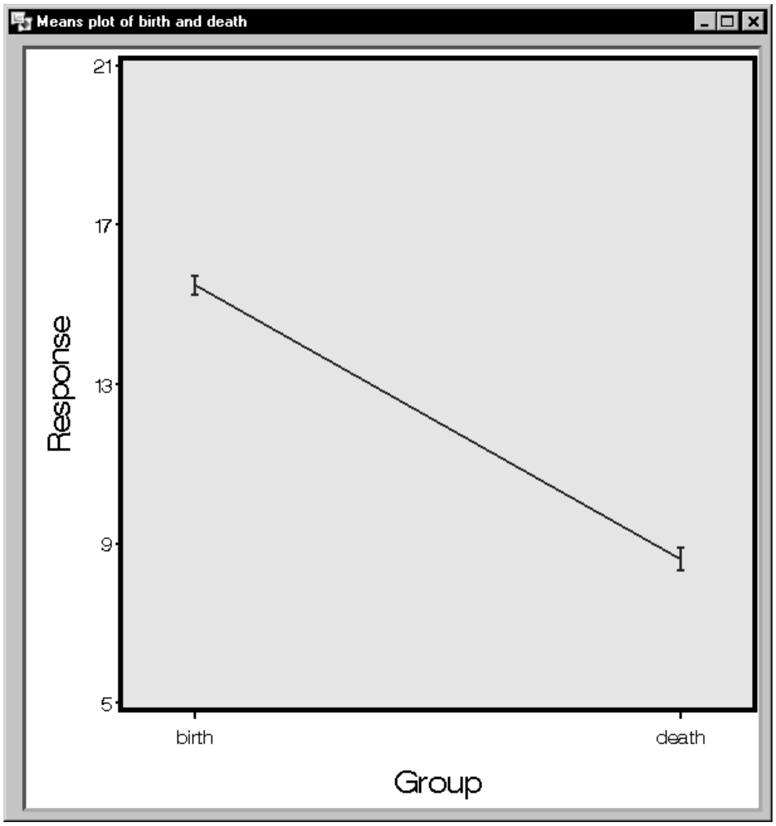

Figure 8.11. Paired t-test: Means Plot

Two-Sample Test for Proportions

In the Two-Sample Test for Proportions task, you can determine whether two probabilities are the same.

The data analyzed in this example are taken from a study measuring the accuracy of two computer programs. Each program searches the World Wide Web and returns a list of web pages that meet a particular set of specified criteria. The data set Search contains two samples in which each observation is either 'yes' or 'no'. A response of 'yes' indicates that the program returns the desired page at the top of the list of potential pages; a value of 'no' indicates that this is not the case. The data set contains the results of 535 searches using an older search program and 409 searches using a new program. The variables containing the results for the old and new programs are named oldfind and newfind, respectively.

Suppose that you want to determine whether the probability of a correct search by the new algorithm is higher than that for the old algorithm. That is, you want to determine whether you can reject the null hypothesis that the two probabilities are equal in favor of the alternative that the new probability is larger. The values for analysis are contained in the two variables oldfind and newfind.

Open the Search Data Set

The data are provided in the Analyst Sample Library. To access this Analyst sample data set, follow these steps:

1. Select **Tools** → **Sample Data** . . .
2. Select Search.
3. Click **OK** to create the sample data set in your Sasuser directory.
4. Select **File** → **Open By SAS Name** . . .
5. Select Sasuser from the list of **Libraries**.
6. Select Search from the list of members.
7. Click **OK** to bring the Search data set into the data table.

Request a Two-Sample Test for Proportions

To perform the analysis, follow these steps:

1. Select **Statistics**→ **Hypothesis Tests** → **Two-Sample Test for Proportions** . . .
2. Select **Two variables** in the box labeled **Groups are in**.
3. Select the variable newfind as the Group 1 variable.
4. Select the variable oldfind as the Group 2 variable.
5. Select the **Level of Interest** by clicking on the down arrow and selecting **yes** to test whether the two groups have the same proportions of success.
6. Specify the **Alternative** hypothesis by selecting **Prop 1 - Prop 2 > 0**.

Note that, if your data are arranged so that the values for the two groups are contained in a single variable, you can define the dependent and group variables by selecting **One variable** in the box labeled **Groups are in**.

Figure 8.12 displays the Two-Sample Test for Proportions dialog.

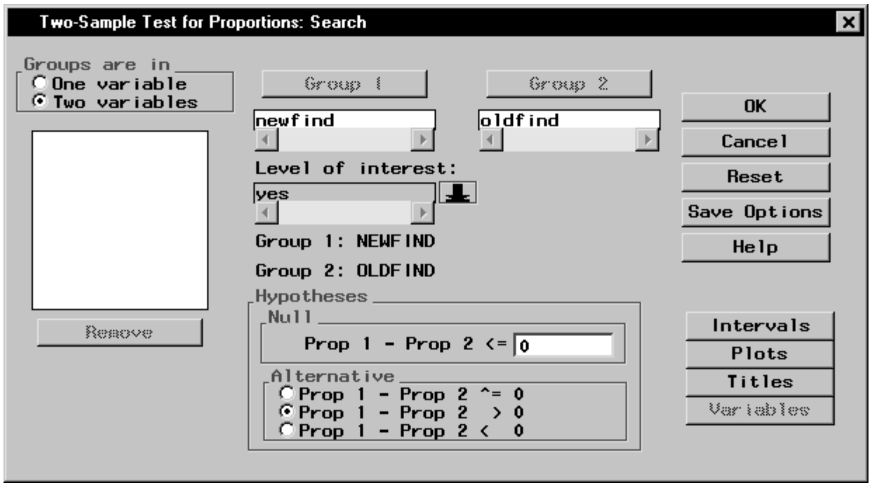

Figure 8.12. Two-Sample Test for Proportions Dialog

In Figure 8.12, the null hypothesis specifies that the proportions of success for the algorithms are equal (or, equivalently, that the difference between the proportions is 0). The alternative hypothesis is that the probability of a correct search by the new algorithm is higher than that for the old algorithm.

Click **OK** in the main dialog to perform the analysis.

Review the Results

The results of the hypothesis test are displayed in Figure 8.13.

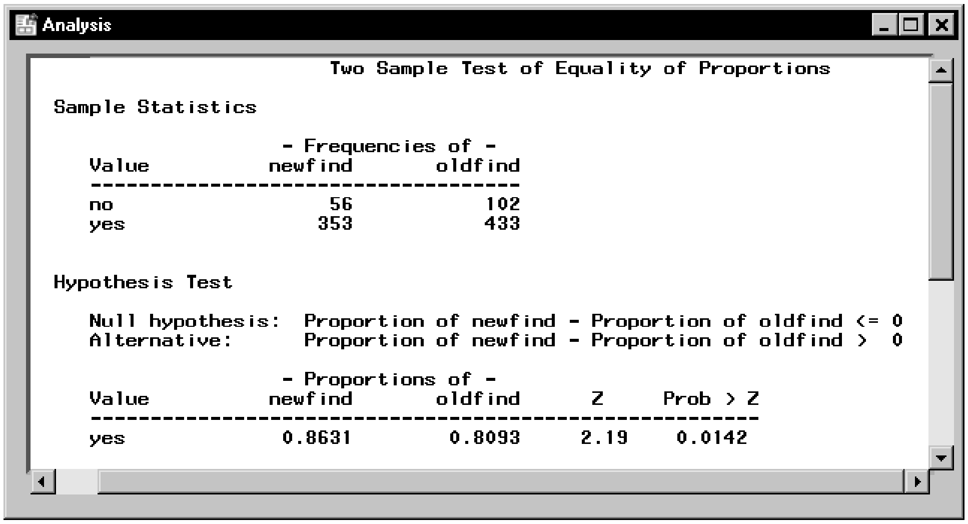

```
Analysis

                         Two Sample Test of Equality of Proportions

 Sample Statistics

                   - Frequencies of -
    Value          newfind       oldfind
    ------------------------------------
    no                  56           102
    yes                353           433


 Hypothesis Test

    Null hypothesis:  Proportion of newfind - Proportion of oldfind <= 0
    Alternative:      Proportion of newfind - Proportion of oldfind >  0

                   - Proportions of -
    Value          newfind       oldfind      Z     Prob > Z
    --------------------------------------------------------
    yes             0.8631        0.8093     2.19    0.0142
```

Figure 8.13. Two-Sample Test for Proportions: Results

The "Sample Statistics" table lists the frequency of 'yes' and 'no' responses for each variable. The "Hypothesis Test" table displays the null and alternative hypotheses and the results of the test.

The observed proportion of 'yes' responses is 0.8631 for the newfind variable, and 0.8093 for the oldfind variable. The Z statistic of 2.19 and associated p-value of 0.0142 indicate that the proportion of successful searches is significantly larger for the new search algorithm.

Two-Sample Test for Variances

In the Two-Sample Test for Variances task, you can test whether two variables have different variances, or, if you have a single variable that contains values for two groups, you can determine whether the variance differs between the groups.

The data set analyzed in this example, Gpa, contains test scores for 224 students. The data include the students' grade point averages (the variable gpa), high school scores in mathematics, science, and English (the variables hsm, hss, and hse, respectively), and SAT math and verbal scores (the variables satm and satv, respectively).

Suppose that you want to examine the difference in grade point averages between males and females. You can use the two-sample test for variances to test whether the variance of the grade point average differs between males and females.

Open the Gpa Data Set

The data are provided in the Analyst Sample Library. To access this Analyst sample data set, follow these steps:

1. Select **Tools** → **Sample Data** . . .
2. Select GPA.
3. Click **OK** to create the sample data set in your Sasuser directory.
4. Select **File** → **Open By SAS Name** . . .
5. Select Sasuser from the list of **Libraries**.
6. Select Gpa from the list of members.
7. Click **OK** to bring the Gpa data set into the data table.

Request a Two-Sample Test for Variances

To perform the hypothesis test, follow these steps:

1. Select **Statistics** → **Hypothesis Tests** → **Two-Sample Test for Variances** . . .
2. Ensure that **One variable** is selected in the box labeled **Groups are in**.
3. Select the variable gpa as the Dependent variable.
4. Select the variable sex as the Group variable.

If your data are arranged so that the values for both groups are contained in two variables, you can define the two groups by checking the **Two variables** selection in the box labeled **Groups are in**.

The null hypothesis for the test is that the two variances are equal (or, equivalently, that their ratio is equal to 1). You can specify the type of alternative hypothesis. The three choices are that Variance 1 is not equal to, is greater than, or is less than Variance 2. In Figure 8.14, the alternative hypothesis states that the two variances are not equal, which is the two-sided alternative hypothesis.

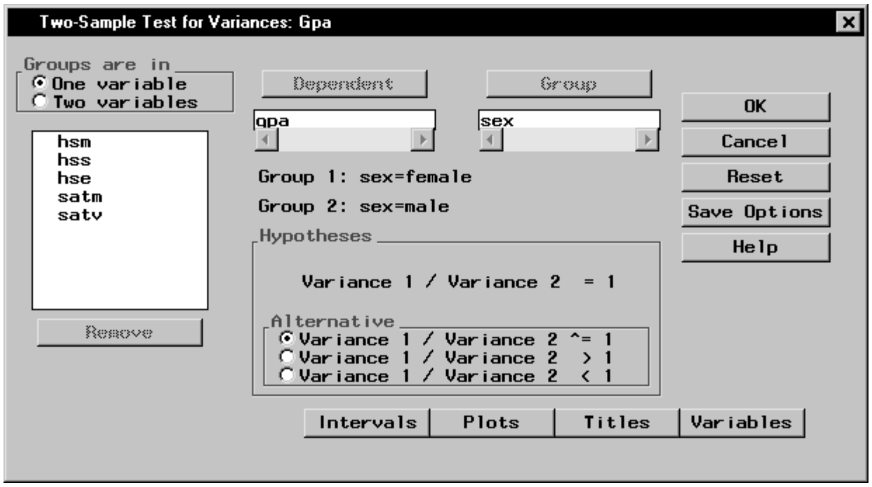

Figure 8.14. Two-Sample Test for Variances Dialog

Request a Box-&-Whisker Plot

To request a box-and-whisker plot in addition to the hypothesis test, follow these steps:

1. Click on the **Plots** button.
2. Select **Box-&-whisker plot**.
3. Click **OK**.

Figure 8.15 displays the Plots dialog with **Box-&-whisker plot** selected. Note that the plot is constructed to have a mean of zero.

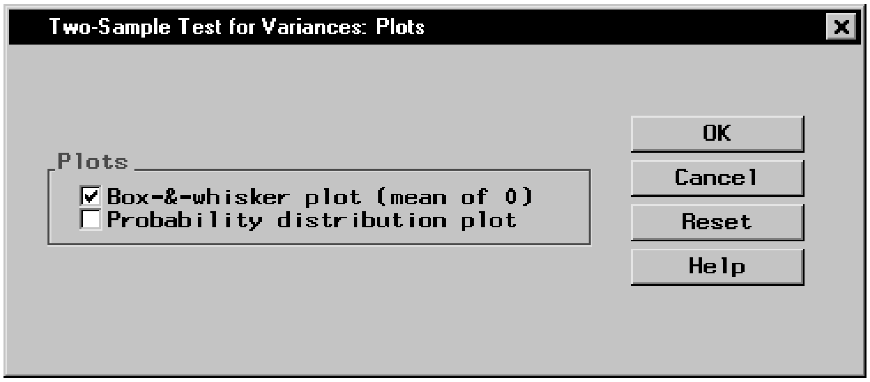

Figure 8.15. Two-Sample Test for Variances: Plots Dialog

Click **OK** in the Two-Sample Test for Variances dialog to perform the hypothesis test.

Review the Results

Figure 8.16 displays the results of the hypothesis test. The output contains the results of the hypothesis test, including summary statistics, the *F* statistic, and the associated *p*-value.

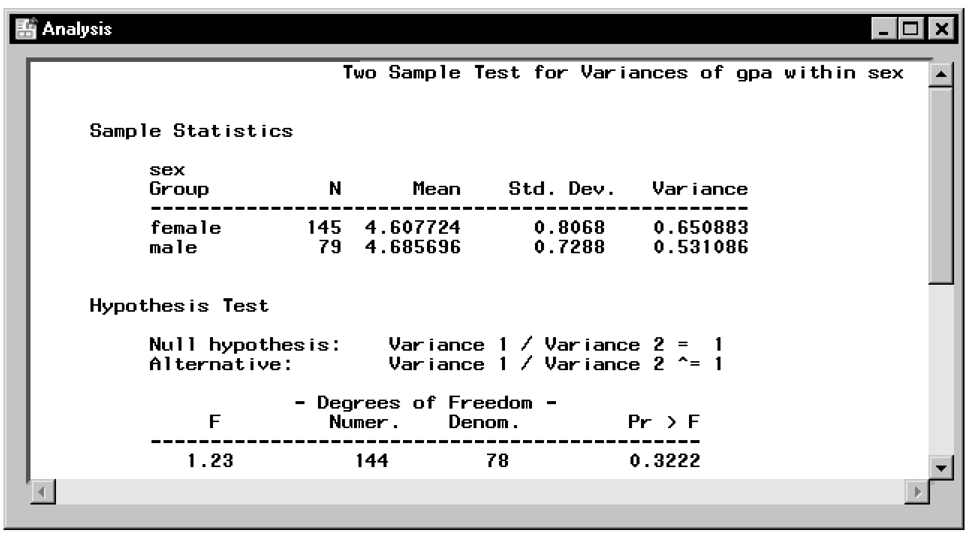

```
Analysis
                         Two Sample Test for Variances of gpa within sex

     Sample Statistics

          sex
          Group          N      Mean    Std. Dev.   Variance
          --------------------------------------------------
          female       145  4.607724      0.8068    0.650883
          male          79  4.685696      0.7288    0.531086

     Hypothesis Test

          Null hypothesis:    Variance 1 / Variance 2 =  1
          Alternative:        Variance 1 / Variance 2 ^= 1

                      - Degrees of Freedom -
               F         Numer.    Denom.        Pr > F
          ----------------------------------------------
             1.23          144        78          0.3222
```

Figure 8.16. Two-Sample Test for Variances: Output

The "Sample Statistics" table displays the variance of the variable gpa for both females (0.6509) and males (0.5311). The "Hypothesis Test" table displays the test statistics: the F value is 1.23 and the resulting p-value is 0.3222. Thus, the data give no evidence for rejecting the hypothesis of equal variances.

Figure 8.17 displays the box-and-whisker plot. Observations that fall beyond the whiskers are identified with a square symbol.

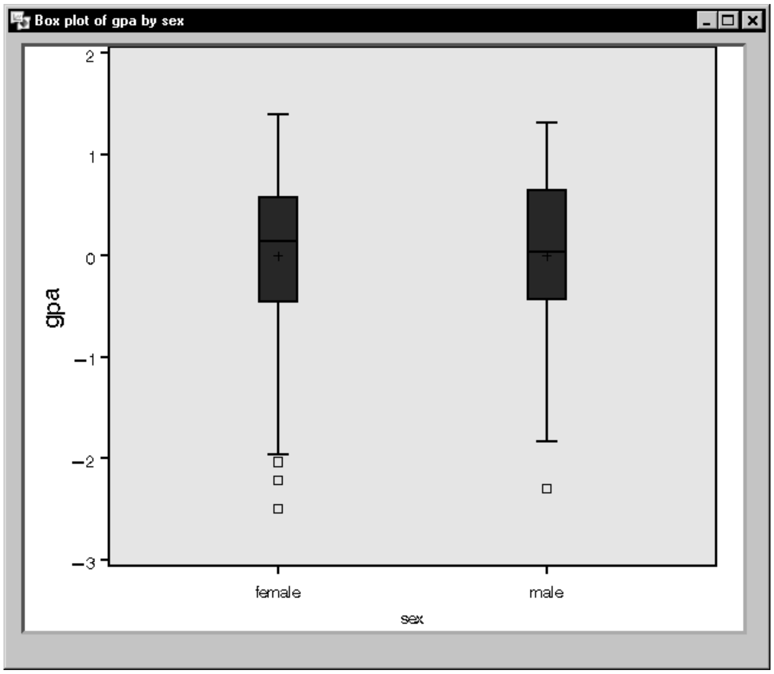

Figure 8.17. Two-Sample Test for Variances: Box-and-whisker Plot

The box-and-whisker plot displays the amount of spread and the range for the two variables. The two groups do not appear to be appreciably different.

Discussion of Other Tests

The following descriptions provide an overview of other hypothesis tests available in the Analyst Application.

One-Sample Z-Test for a Mean

In the One-Sample Z-Test for a Mean task, you can test whether the mean of a population is equal to the value you specify in the null hypothesis. This test is appropriate when the population standard deviation or variance is known, and your data are either normally distributed or you have a large number of observations. Generally, a sample size of at least 30 is considered to be sufficient.

The default output from the test includes summary statistics for the selected variable, the Z statistic, and the associated p-value.

One-Sample Test for a Proportion

In the One-Sample Test for a Proportion task, you can test whether the proportion of a population giving a certain response is equal to the proportion you specify in the null hypothesis.

The default output from this test provides a frequency table of responses versus the analysis variable, the observed proportion, the Z statistic, and the associated p-value.

One-Sample Test for a Variance

In the One-Sample Test for a Variance task, you can test whether the variance of a population is equal to the value you specify in the null hypothesis.

The default output from this test includes summary statistics for the selected variable, the chi-square statistic, and the associated p-value.

Two-Sample t-Test for Means

In the Two-Sample t-Test for Means task, you can test whether the means of two populations are equal or, optionally, whether they differ by a specified amount. Two-sample data arise when two independent samples are observed, possibly with different sample sizes. Note that, if the two samples are not independent, the two-sample t-test is inappropriate and you should use instead the Two-Sample

Paired t-Test for Means task (see the section "Paired t-test" beginning on page 164 for more information).

The default output from the test includes summary statistics for the two samples, two t statistics, and the associated p-values. The first t statistic assumes the population variances of the two groups are equal; the second statistic is an approximate t statistic and should be used when the population variances of the two groups are potentially unequal.

References

SAS Institute Inc. (1999), *SAS/STAT User's Guide, Version 7-1*, Cary, NC: SAS Institute Inc.

Schlotzhauer, Sandra D. and Littell, Ramon C. (1991), *SAS System for Elementary Statistical Analysis, Second Edition*, Cary, NC: SAS Institute Inc.

U.S. Bureau of the Census (1995), *Statistical Abstract of the United States*, Washington, D.C.

The correct bibliographic citation for this manual is as follows: SAS Institute Inc., *The Analyst Application, First Edition*, Cary, NC: SAS Institute Inc., 1999. 476 pp.

The Analyst Application, First Edition

ISBN 1–58025–446–2

SAS Institute Inc., SAS Campus Drive, Cary, North Carolina 27513.

1st printing, October 1999