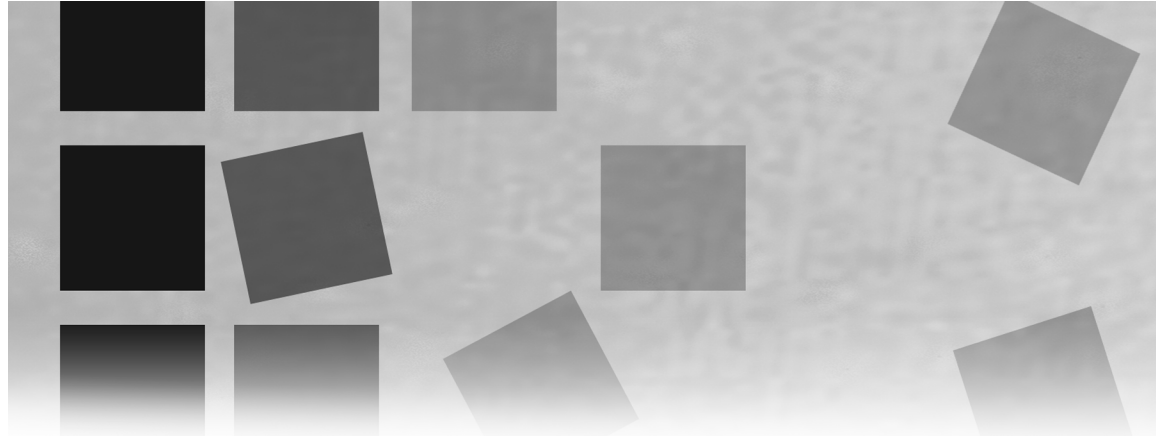


Part 3

Comparing Groups

Chapter 7	Comparing Paired Groups	189
Chapter 8	Comparing Two Independent Groups	217
Chapter 9	Comparing More Than Two Groups	257



Chapter 7

Comparing Paired Groups

Are students' grades on an achievement test higher after they complete a special course on how to take tests? Do employees who follow a regular exercise program for a year have lower resting pulse rates than they had when they started the program?

These questions involve comparing paired groups of data. The variable that classifies the data into two groups should be nominal. The response variable can be ordinal or continuous, but must be numeric. This chapter discusses the following topics:

- deciding whether you have independent or paired groups
- summarizing data from paired groups
- building a statistical test of hypothesis to compare the two groups
- deciding which statistical test to use
- performing statistical tests for paired groups

For the various tests, this chapter shows how to use SAS to perform the test, and how to interpret the test results.

Deciding between Independent and Paired Groups	191
Independent Groups	191
Paired Groups	191
Summarizing Data from Paired Groups	192
Finding the Differences between Paired Groups	192
Summarizing Differences with PROC UNIVARIATE	196
Summarizing Differences with Other Procedures	198
Building Hypothesis Tests to Compare Paired Groups	198
Deciding Which Statistical Test to Use	199
Understanding Significance	200
Performing the Paired-Difference <i>t</i>-test	201
Assumptions	201
Using PROC UNIVARIATE to Test Paired Differences	204
Using PROC TTEST to Test Paired Differences	206
Identifying ODS Tables	209
Performing the Wilcoxon Signed Rank Test	210
Finding the <i>p</i> -value	211
Understanding Other Items in the Output	211
Identifying ODS Tables	211
Summary	212
Key Ideas	212
Syntax	213
Example	214

Deciding between Independent and Paired Groups

When comparing two groups, you want to know whether the means for the two groups are different. The first step is to decide whether you have independent or paired groups. This chapter discusses paired groups. Chapter 8 discusses independent groups.

Independent Groups

Independent groups of data contain measurements for two unrelated samples of items. Suppose you have random samples of the salaries for male and female accountants. The salary measurements for men and women form two distinct and separate groups. The goal of analysis is to compare the average salaries for men and women, and decide whether the difference between the average salaries is greater than what could happen by chance. As another example, suppose a researcher selects a random sample of children, some who use fluoride toothpaste, and some who do not. There is no relationship between the children who use fluoride toothpaste and the children who do not. A dentist counts the number of cavities for each child. The goal of analysis is to compare the average number of cavities for children who use fluoride toothpaste and for children who do not.

Paired Groups

Paired groups of data contain measurements for one sample of items, but there are two measurements for each item. A common example of paired groups is before-and-after measurements, where the goal of analysis is to decide whether the average change from before to after is greater than what could happen by chance. For example, a doctor weighs 30 people before they begin a program to quit smoking, and weighs them again six months after they have completed the program. The goal of analysis is to decide whether the average weight change is greater than what could happen by chance.

Summarizing Data from Paired Groups

With paired groups, the first step is to find the paired differences for each observation in the sample. The second step is to summarize the differences.

Finding the Differences between Paired Groups

Sometimes, the difference between the two measurements for an observation is already calculated. For example, you might be analyzing data from a cattle-feeding program, and the data gives the weight change of each steer. If you already have the difference in the weight of each steer, you can skip the rest of this section.

Other times, only the before-and-after data is available. Before you can summarize the data, you need to find the differences. Instead of finding the differences with a calculator, you can compute them in SAS. To do this, you add *program statements* to the **DATA** step. There are many types of program statements in SAS. The simplest program statement creates a new variable.

Table 7.1 shows data from an introductory statistics class, STA6207.¹ The table shows scores from two exams for the 20 students in the class. Both exams covered the same material. Each student has a pair of scores, one score for each exam. The professor wants to find out if the exams appear to be equally difficult. If they are, the average difference between the two exam scores should be small.

¹ Data is from Dr. Ramon Littell, University of Florida. Used with permission.

Table 7.1 Exam Data

Student	Exam 1 Score	Exam 2 Score
1	93	98
2	88	74
3	89	67
4	88	92
5	67	83
6	89	90
7	83	74
8	94	97
9	89	96
10	55	81
11	88	83
12	91	94
13	85	89
14	70	78
15	90	96
16	90	93
17	94	81
18	67	81
19	87	93
20	83	91

This data is available in the STA6207 data set in the sample data for the book.

The following SAS statements create a data set and use a program statement to create a new variable:

```
data STA6207;
  input student exam1 exam2 @@;
  scorediff = exam2 - exam1 ;
  label scorediff='Differences in Exam Scores';
  datalines;
1 93 98 2 88 74 3 89 67 4 88 92 5 67 83 6 89 90
7 83 74 8 94 97 9 89 96 10 55 81 11 88 83 12 91 94
13 85 89 14 70 78 15 90 96 16 90 93 17 94 81
18 67 81 19 87 93 20 83 91
;
run;
```

The spacing that was used for the program statement in the example was chosen for readability. None of the spaces are required. The following statement produces the same results as the previous statement:

```
scorediff=exam2-exam1;
```

Table 7.2 outlines the steps for creating new variables with program statements.

Table 7.2 Creating New Variables with Program Statements

To create a new variable in the DATA step, follow these steps:

1. Type the **DATA** and **INPUT** statements for the data set.
2. Choose the name of the new variable. (Use the rules for SAS names in Table 2.3.) Type the name of the variable, and follow the name with an equal sign. If the new variable is **DIFF**, you would type the following:

```
diff =
```

3. Decide what combination of other variables will be used to create the new variable. You must have defined these other variables in the **INPUT** statement. Type the combination of variables after the equal sign. For paired groups, you are usually interested in the difference between two variables. If the variable **DIFF** is the difference between **AFTER** and **BEFORE**, you would type the following:

```
diff = after - before
```


4. End the program statement with a semicolon. For this example, here is the complete program statement:

```
diff = after - before;
```

5. Follow the program statement with the **DATALINES** statement, and then with the lines of data to complete the **DATA** step. Or, if you are including data from an external source, use the **INFILE** statement.

You can create several variables in a single **DATA** step. See “Technical Details: Program Statements” for examples.

Technical Details: Program Statements

You can create several variables in the same **DATA** step by including a program statement for each new variable. In addition, many combinations of variables can be used to create variables. Here are some examples:

Total cavities is the sum of cavities at the beginning and new cavities:

```
totcav = begcav + newcav;
```

The new variable is the old variable plus 5:

```
newvar = oldvar + 5;
```

Dollars earned is the product of the number sold and the cost for each one:

```
dollars = numsold*unitcost;
```

Yield per acre is the total yield divided by the number of acres:

```
yldacre = totyield/numacres;
```

Credit limit is the square of the amount in the checking account:

```
credlimt = checkamt ** 2;
```

In general, use a plus sign (+) to indicate addition, a minus sign (–) for subtraction, an asterisk (*) for multiplication, a slash (/) for division, and two asterisks and a number (***n*) to raise a variable to the *n*-th power.

Summarizing Differences with PROC UNIVARIATE

After creating the differences, you can summarize them using **PROC UNIVARIATE**. Because the difference is a single variable, you can use the statements and options discussed in Chapters 4 and 5. (This procedure provides an extensive summary, histograms, and box plots. Later on, the section “Summarizing Differences with Other Procedures” discusses how you can also use the **MEANS**, **FREQ**, **GCHART**, and **CHART** procedures.)

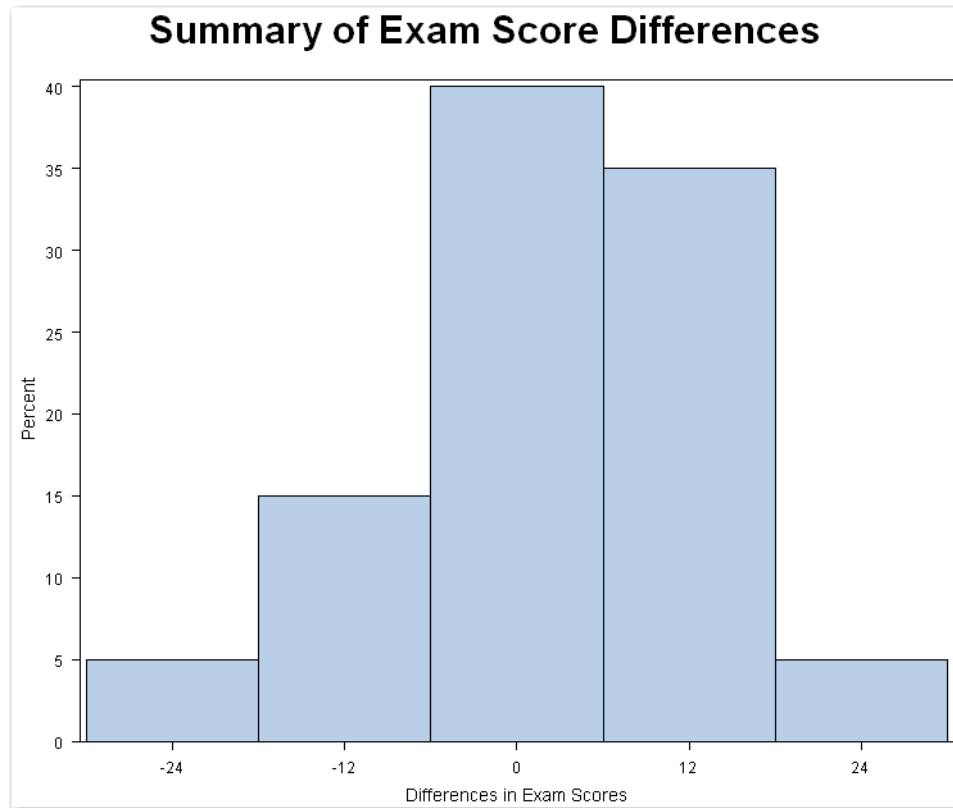
```
proc univariate data=STA6207;
  var scorediff;
  histogram;
  title 'Summary of Exam Score Differences';
run;
```

Figure 7.1 shows the Basic Statistical Measures and the Extreme Observations tables from PROC UNIVARIATE output. The statements above also create other tables in the output, which are not shown in Figure 7.1. Figure 7.1 also shows the histogram.

Figure 7.1 Summary of Exam Score Differences

Summary of Exam Score Differences			
The UNIVARIATE Procedure			
Variable: scorediff (Differences in Exam Scores)			
Basic Statistical Measures			
Location		Variability	
Mean	2.550000	Std Deviation	10.97113
Median	4.000000	Variance	120.36579
Mode	3.000000	Range	48.00000
		Interquartile Range	9.50000
Extreme Observations			
----Lowest----		----Highest---	
Value	Obs	Value	Obs
-22	3	8	14
-14	2	8	20
-13	17	14	18
-9	7	16	5
-5	11	26	10

The average difference between exam scores is 2.55. The minimum difference is –22, and the maximum difference is 26.

Figure 7.1 Summary of Exam Score Differences (continued)

The histogram shows that many of the exam score differences are between -6 and 6 (the bar labeled **0** in the figure). The table below summarizes the midpoints and range of values for each bar in Figure 7.1.

Midpoint	Range of Values
-24	$-30 \leq \text{scorediff} < -18$
-12	$-18 \leq \text{scorediff} < -6$
0	$-6 \leq \text{scorediff} < 6$
12	$6 \leq \text{scorediff} < 18$
24	$18 \leq \text{scorediff} < 30$

However, some of the score differences are large, in both the positive and negative directions. You might think that the average difference in exam scores is not statistically significant. However, you cannot be certain from simply looking at the data.

Use the PROC UNIVARIATE output to check for errors in the new difference variable. In general, check the minimum and maximum values to see whether these values make sense for your data. For example, with the exam scores, the minimum cannot be less than -100 , and cannot be more than 100 (in other words, no points on one exam, and a perfect score on the other exam). The count of differences should match the number of pairs in your data. If the count is different, then check for an error in the program statement that creates the difference variable. You might have an idea of the expected average difference—in that case, check the average value to see whether it makes sense.

Summarizing Differences with Other Procedures

Because the difference is a single variable, you can use the other procedures and options discussed in Chapter 4.

Add the **PLOT** option to PROC UNIVARIATE to create a box plot and a stem-and-leaf plot. Add the **FREQ** option to create a frequency table. Or, use PROC FREQ to create the frequency table. Use PROC MEANS for a concise summary. You can use the GCHART or CHART procedure to create bar charts, but the histogram is a better choice for the difference variable. Because the difference is a continuous variable, using a histogram is more appropriate.

Building Hypothesis Tests to Compare Paired Groups

So far, this chapter has discussed how to summarize differences between paired groups. Suppose you want to know how important the differences are, and if the differences are large enough to be significant. In statistical terms, you want to perform a hypothesis test. This section discusses hypothesis testing when comparing paired groups. (See Chapter 5 for the general idea of hypothesis testing.)

In building a test of hypothesis, you work with two hypotheses. When comparing paired groups, the null hypothesis is that the mean difference is 0 , and the alternative hypothesis is that the mean difference is different from 0 . In this case, the notation is the following:

$$H_0: \mu_D = 0$$

$$H_a: \mu_D \neq 0$$

μ_D indicates the mean difference.

In statistical tests that compare paired groups, the hypotheses are tested by calculating a test statistic from the data, and then comparing its value to a reference value that would be the result if the null hypothesis were true. The test statistic is compared to different reference values that are based on the sample size. In part, this is because smaller differences can be detected with larger sample sizes. As a result, if you use the reference value for a large sample size on a smaller sample, you are more likely to incorrectly conclude the mean difference is zero when it is not.

This concept is similar to the concept of confidence intervals, in which different t -values are used (discussed in Chapter 6). With confidence intervals, the t -value is based on the degrees of freedom, which are determined by the sample size.

Deciding Which Statistical Test to Use

Because you test different hypotheses for independent and paired groups, you use different tests. In addition, there are parametric and nonparametric tests for each type of group. When deciding which statistical test to use, first, decide whether you have independent groups or paired groups.

Then, decide whether you should use a parametric or a nonparametric test. This second decision is based on whether the assumptions for the parametric test seem reasonable. In general, if the assumptions for the parametric test seem reasonable, use the parametric test. Although nonparametric tests have fewer assumptions, these tests typically are less powerful in detecting differences between groups.

The rest of this chapter describes the test to use in the two situations for paired groups. See Chapter 8 for details on the tests for the two situations for independent groups. Table 7.3 summarizes the tests.

Table 7.3 Statistical Tests for Comparing Two Groups

Type of Test	Groups	
	Independent ²	Paired
Parametric	Two-sample <i>t</i> -test	Paired-difference <i>t</i> -test
Nonparametric	Wilcoxon Rank Sum test	Wilcoxon Signed Rank test

The steps for performing the analyses to compare paired groups are the following:

1. Create a SAS data set.
2. Check the data set for errors.
3. Choose the significance level for the test.
4. Check the assumptions for the test.
5. Perform the test.
6. Make conclusions from the test results.

The rest of this chapter uses these steps for all analyses.

Understanding Significance

For each of the four tests in Table 7.3, there are two possible results: the *p*-value is less than the reference probability, or it is not. (See Chapter 5 for discussions on significance levels and on understanding practical significance and statistical significance.)

Groups Significantly Different

If the *p*-value is less than the reference probability, the result is statistically significant, and you reject the null hypothesis. For paired groups, you conclude that the mean difference is significantly different from 0.

² Chapter 8 discusses comparing independent groups. For completeness, this table shows the tests for comparing both types of groups.

Groups Not Significantly Different

If the p -value is greater than the reference probability, the result is not statistically significant, and you fail to reject the null hypothesis. For paired groups, you conclude that the mean difference is not significantly different from 0.

Do not conclude that the mean difference is 0. You do not have enough evidence to conclude that the mean difference is 0. The results of the test indicate only that the mean difference is not significantly different from 0. You never have enough evidence to conclude that the population is exactly as described by the null hypothesis. In statistical hypothesis testing, do not accept the null hypothesis. Instead, either reject or fail to reject the null hypothesis. For a more detailed discussion of hypothesis testing, see Chapter 5.

Performing the Paired-Difference t -test

The paired-difference t -test is a parametric test for comparing paired groups.

Assumptions

The two assumptions for the paired-difference t -test are the following:

- each pair of measurements is independent of other pairs of measurements
- differences are from a normal distribution

Because the test analyzes the differences between paired observations, it is appropriate for continuous variables.

To illustrate this test, consider an experiment in liquid chromatography.³ A chemist is investigating synthetic fuels produced from coal, and wants to measure the naphthalene values by using two different liquid chromatography methods. Each of the 10 fuel samples is divided into two units. One unit is measured using standard liquid chromatography. The other unit is measured using high-pressure liquid chromatography. The goal of analysis is to test whether the mean difference between the two methods is different from 0 at the 5% significance level. The chemist is willing to accept a 1-in-20 chance of saying that the mean difference is significantly different from 0, when, in fact, it is not. Table 7.4 shows the data.

³ Data is from C. K. Bayne and I. B. Rubin, *Practical Experimental Designs and Optimization Methods for Chemists* (New York: VCH Publishers, 1986). Used with permission.

Table 7.4 Liquid Chromatography Data

Sample	High-Pressure	Standard
1	12.1	14.7
2	10.9	14.0
3	13.1	12.9
4	14.5	16.2
5	9.6	10.2
6	11.2	12.4
7	9.8	12.0
8	13.7	14.8
9	12.0	11.8
10	9.1	9.7

This data is available in the **chromat** data set in the sample data for the book.

Applying the six steps of analysis.

1. Create a SAS data set.
2. Check the data set for errors. Use PROC UNIVARIATE to check for errors. Based on the results, assume that the data is free of errors.
3. Choose the significance level for the test. Choose a 5% significance level, which requires a p -value less than 0.05 to conclude that the groups are significantly different.
4. Check the assumptions for the test. The first assumption is independent pairs of measurements. The assumption seems reasonable because there are 10 different samples that are divided into two units. The second assumption is that the differences are from a normal distribution. Use program statements to create a new variable named **methdiff** (which is **std** subtracted from **hp**, or the standard method subtracted from the high-pressure method). Use PROC UNIVARIATE to check for normality. Although there are only a few data points, the second assumption seems reasonable.

What if the second assumption was not reasonable? With only 10 samples, you do not have a very complete picture of the distribution. Think of weighing 10 people, and then deciding whether the weights are normally distributed based only on the weights of these 10 people. Sometimes you have additional knowledge about the population that enables you to use tests that require a normally distributed population. For example, you might know that the weights are generally normally distributed, and you can use that information. However, this approach should be used with caution. Use it only when you have substantial additional knowledge about the population. Otherwise, if you are concerned about the assumption of normality, or if a test for normality indicates that your data is not a sample from a normally distributed population, consider performing a Wilcoxon Signed Rank test, which is discussed later in this chapter.

5. Perform the test. See “Building Hypothesis Tests to Compare Paired Groups” earlier in this chapter. For paired groups, the null hypothesis is that the mean difference (μ_D) is 0. The alternative hypothesis is that the mean difference (μ_D) is different from 0.

SAS provides two ways to test these hypotheses. In some cases, you have only the difference between the pairs. See “Using PROC UNIVARIATE to Test Paired Differences” for this case. In other cases, you have the raw data with the values for each pair. You can still create the difference variable and use PROC UNIVARIATE, or you can use **PROC TTEST**. See “Using PROC TTEST to Test Paired Differences” for this case.

6. Make conclusions from the test results. See the section “Finding the p -value” for each test.

Using PROC UNIVARIATE to Test Paired Differences

For the `chromat` data, the difference variable is available. PROC UNIVARIATE automatically performs a test that the mean of this variable is 0. The SAS statements below create the data set and perform the statistical test. These statements use the **ODS** statement discussed at the end of Chapter 4. The ODS statement limits the output to the test results. (Assume that you have checked the difference variable for errors and tested it for normality.)

```
data chromat;
  input hp std @@;
  methdiff=hp-std;
  datalines;
12.1 14.7 10.9 14.0 13.1 12.9 14.5 16.2 9.6 10.2 11.2 12.4
9.8 12.0 13.7 14.8 12.0 11.8 9.1 9.7
;
run;

ods select TestsForLocation;
proc univariate data=chromat;
  var methdiff;
  title 'Testing for Differences between Chromatography Methods';
run;
```

Figure 7.2 shows the output. This output shows the results from three statistical tests. The heading in the procedure identifies the variable being tested. The **Tests for Location** table identifies the hypothesis being tested. PROC UNIVARIATE automatically tests the hypothesis that the mean difference is 0, shown by the heading **Mu0=0**. The **Test** column identifies the statistical test. This particular section discusses the paired-difference *t*-test. The nonparametric Wilcoxon Signed Rank test is discussed later in the chapter.

Figure 7.2 Comparing Paired Groups with PROC UNIVARIATE

Testing for Differences between Chromatography Methods				
The UNIVARIATE Procedure				
Variable: methdiff				
Tests for Location: Mu0=0				
Test		-Statistic-		-----p Value-----
Student's t	t	-3.56511	Pr > t	0.0061
Sign	M	-3	Pr >= M	0.1094
Signed Rank	S	-24.5	Pr >= S	0.0078

Finding the p -value

Look for the **Student's t** row in the Tests for Location table. The **Statistic** column identifies the t -statistic of **-3.565** (rounded). The **p Value** column identifies the p -value as **Pr > |t|** of **0.0061**, which is less than the significance level of 0.05. You conclude that the mean difference in naphthalene values from the two methods is significantly different from 0.

In general, to interpret SAS results, look at the p -value that appears to the right of the **Pr > |t|** heading. If the p -value is less than the significance level, then you can conclude that the mean difference between the paired groups is significantly different from 0. If the p -value is greater than the significance level, then you can conclude that the mean difference is not significantly different from 0. Do not conclude that the mean difference is 0. (See the discussion earlier in this chapter.)

Technical Details: Formula for Student's t

The test statistic for the paired-difference t -test is calculated as the following:

$$t = \frac{\bar{X} - \mu_0}{s / \sqrt{n}}$$

\bar{X} is the average difference, s is the standard deviation of the difference, and n is the sample size. For testing paired groups, μ_0 is 0 (testing the null hypothesis of no difference between groups).

The general form of the statements to perform the paired-difference *t*-test using PROC UNIVARIATE is shown below:

```
PROC UNIVARIATE DATA=data-set-name;  
VAR difference-variable;
```

data-set-name is the name of a SAS data set, and *difference-variable* is the difference between the pairs of values. If you omit the *difference-variable* from the VAR statement, then the procedure uses all numeric variables. You can specify the TestsForLocation table in an ODS statement, which limits the output to the test results.

Using PROC TTEST to Test Paired Differences

So far, this section has discussed how to test paired differences using a difference variable. This approach is important because sometimes you receive data and you don't have the original pairs of information. When you do have the original pairs, you do not need to create a difference variable. PROC TTEST provides another way to test paired differences.

```
proc ttest data=chromat;  
  paired hp*std;  
  title 'Paired Differences with PROC TTEST';  
run;
```

The **PAIRED** statement identifies the two variables that contain the original pairs of information. In the **PAIRED** statement, an asterisk separates the two variables. This special use of the asterisk is not interpreted as multiply. Instead, the asterisk tells the procedure to subtract the second variable from the first variable. Figure 7.3 shows the output.

Figure 7.3 Comparing Paired Groups with PROC TTEST

Paired Differences with PROC TTEST					
The TTEST Procedure					
Difference: hp - std					
N	Mean	Std Dev	Std Err	Minimum	Maximum
10	-1.2700	1.1265	0.3562	-3.1000	0.2000
Mean	95% CL Mean	Std Dev	95% CL Std Dev		
-1.2700	-2.0758 -0.4642	1.1265	0.7748	2.0565	
DF	t Value	Pr > t			
9	-3.57	0.0061			

This output shows how SAS creates the difference variable, gives descriptive statistics for the difference variable, and performs the paired-difference *t*-test.

Finding the *p*-value

Look below the heading **Pr > |t|** in Figure 7.3. The value is **0.0061**, which is less than the significance level of 0.05. This test has the same result as the analysis of the difference variable. You conclude that the mean difference in naphthalene values from the two methods is significantly different from 0. You interpret the *p*-value for this test the exact same way as described for PROC UNIVARIATE.

The subheading **Difference: hp - std** indicates how PROC TTEST calculates the paired difference. **Difference** is the definition of the paired difference. SAS subtracts the second variable in the PAIRED statement from the first variable in the statement.

The list below summarizes the **T tests** section of output for the paired-difference *t*-test.

DF	Degrees of freedom for the test. In general, this is the number of pairs minus 1. For the chromat data, DF is 9 , which is 10–1.
-----------	---

t Value

Test statistic for the t -test, calculated as the following:

$$t = \frac{\bar{X} - \mu_0}{s / \sqrt{n}}$$

\bar{X} is the average difference, s is the standard deviation the difference, and n is the sample size. For testing paired groups, μ_0 is 0 (testing the null hypothesis of no difference between groups). PROC TTEST and PROC UNIVARIATE perform the same test.

Pr > |t|

Gives the p -value associated with the test. This p -value is for a two-sided test, which tests the alternative hypothesis that the mean difference between the paired groups is significantly different from 0.

To perform the test at the 5% significance level, you conclude that the mean difference between the paired groups is significantly different from zero if the p -value is less than 0.05. To perform the test at the 10% significance level, you conclude that the mean difference between the paired groups is significantly different from zero if the p -value is less than 0.10. Conclusions for tests at other significance levels are made in a similar manner.

Understanding Other Items in the Output

The list below summarizes the Statistics and Confidence Limits tables in the output in Figure 7.3.

N	Sample size.
Mean	Average of the difference variable. This statistic appears in both output tables.
Std Dev	Standard deviation of the difference variable. This statistic appears in both output tables.

Std Err	Standard error of the difference variable. It is calculated as s/\sqrt{n} , where s is the standard deviation of the difference variable and n is the sample size.
Minimum Maximum	Minimum and maximum values of the difference variable.
95% CL Mean⁴	Upper and lower 95% confidence limits for the mean. If the confidence interval includes 0, then the test will not be significant. If the confidence interval does not include 0, then the test will be significant. See Chapter 6 for more discussion of confidence limits.
95% CL Std Dev⁵	Upper and lower 95% confidence limits for the standard deviation of the difference variable.

Caution: Use the six steps of analysis regardless of the SAS procedure. Although both PROC UNIVARIATE and PROC TTEST perform the paired-difference t -test, the test is the last step in the analysis. Before performing the test, be sure to check the data set for errors, choose the significance level, and check the assumptions for the test.

The general form of the statements to perform the paired-difference t -test using PROC TTEST is shown below:

```
PROC TTEST DATA=data-set-name;  
    PAIRED first-variable * second-variable;
```

data-set-name is the name of a SAS data set. SAS subtracts the *second-variable* from the *first-variable* to create the differences. The PAIRED statement is required, and the asterisk in the statement is required. The variables in the PAIRED statement must be numeric and should be continuous variables for a valid t -test.

Identifying ODS Tables

PROC UNIVARIATE and PROC TTEST automatically create the output tables that test for differences in paired groups. Table 7.5 identifies the ODS table names.

⁴ If you are running a version of SAS earlier than 9.2, you will see headings of Upper CL Mean and Lower CL Mean. The descriptive statistics appear in a different order.

⁵ If you are running a version of SAS earlier than 9.2, you will see headings of Lower CL Std Dev and Upper CL Std Dev. The descriptive statistics appear in a different order.

Table 7.5 ODS Table Names for Testing Paired Groups

SAS Procedure	ODS Name	Output Table
UNIVARIATE	TestsForLocation	Tests for Location table, which contains both the <i>t</i> -test and Wilcoxon Signed Rank test
TTEST	Statistics	Row with descriptive statistics
	ConfLimits	Rows showing mean, standard deviation, and associated confidence limits
	TTests	Row with test statistic, degrees of freedom, and <i>p</i> -value

Performing the Wilcoxon Signed Rank Test

The Wilcoxon Signed Rank test is a nonparametric analogue to the paired-difference *t*-test. The Wilcoxon Signed Rank test assumes that each pair of measurements is independent of other pairs of measurements. This test can be used with ordinal and continuous variables. This section uses the `chromat` data as an example.

Apply the six steps of analysis to the `chromat` data:

1. Create a SAS data set.
2. Check the data set for errors. Assume that you checked and found no errors.
3. Choose the significance level for the test. Choose a 5% significance level.
4. Check the assumptions for the test. Check the assumption of independent pairs of measurements. (See the earlier discussion for the paired-difference *t*-test.)
5. Perform the test. For this test, you need to create the difference variable.
6. Make conclusions from the test results.

PROC UNIVARIATE automatically performs the Wilcoxon Signed Rank test. The following statements were shown earlier in the chapter. Figure 7.2 shows the output.

```
ods select TestsForLocation;
proc univariate data=chromat;
    var methdiff;
    title 'Testing for Differences between Chromatography Methods';
run;
```

The next two topics discuss the output.

Finding the p -value

In Figure 7.2, look for the **Signed Rank** row in the Tests for Location table. The **p Value** column identifies the p -value as **Pr >= |S|** of **0.0078**, which is less than the significance level of 0.05. You conclude that the mean difference in naphthalene values from the two methods is significantly different from 0. In general, you interpret the p -value for the Wilcoxon Signed Rank test the exact same way as described for the paired-difference t -test.

Understanding Other Items in the Output

The Signed Rank row in the Tests for Location table shows the value of the Wilcoxon Signed Rank test statistic, **S**. For the chromat data, this value is **-24.5**.

The **Sign** row in the Tests for Location table shows the results of the sign test. The sign test tests the null hypothesis that the median of the difference variable is 0 against the alternative hypothesis that it is not. Interpret the p -value the exact same way as described for the other two tests. The sign test is generally less likely to detect differences between paired groups than the Wilcoxon Signed Rank test.

Identifying ODS Tables

PROC UNIVARIATE creates the Tests for Location table, which contains both the t -test and the Wilcoxon Signed Rank test. Table 7.5 identifies the ODS name for this table as TestsForLocation.

Summary

Key Ideas

- Paired groups contain paired measurements for each item.
- Use program statements in the **DATA** step to create a difference variable. Use the same SAS procedures that summarize a single variable to summarize the new difference variable.
- To choose a statistical test, first, decide whether the data is from independent or paired groups. Then, decide whether to use a parametric or a nonparametric test. The second decision is based on whether the assumptions for a parametric test seem reasonable. Tests for the four cases are:

Type of Test	Groups	
	Independent ⁶	Paired
Parametric	Two-sample <i>t</i> -test	Paired-difference <i>t</i> -test
Nonparametric	Wilcoxon Rank Sum test	Wilcoxon Signed Rank test

- Regardless of the statistical test you choose, the six steps of analysis are:
 1. Create a SAS data set.
 2. Check the data set for errors.
 3. Choose the significance level for the test.
 4. Check the assumptions for the test.
 5. Perform the test.
 6. Make conclusions from the test results.

⁶ Chapter 8 discusses comparing independent groups. For completeness, this table shows the tests for comparing both types of groups.

- Regardless of the statistical test, to make conclusions, compare the p -value for the test with the significance level.
 - If the p -value is less than the significance level, then you reject the null hypothesis and conclude that the mean difference between the paired groups is significantly different from 0.
 - If the p -value is greater than the significance level, then you do not reject the null hypothesis. You conclude that the mean difference between the paired groups is not significantly different from 0. (Remember, do not conclude that the mean difference is 0.)
- Test for normality using PROC UNIVARIATE. For paired groups, check the difference variable.
- When comparing paired groups, PROC UNIVARIATE provides both parametric and nonparametric tests for the difference variable. PROC TTEST can provide the paired-difference t -test using the two original pairs.
- SAS creates a confidence interval for the mean difference with PROC TTEST. PROC MEANS can also display this confidence interval. (See Chapter 6 for details.) If the confidence interval includes 0, then the test will not be significant. If the confidence interval does not include 0, then the test will be significant.

Syntax

To summarize data from paired groups

- To summarize data from paired groups, use PROC UNIVARIATE with the new difference variable just as you would for a single variable. For a concise summary, use PROC MEANS. To create a frequency table, use PROC FREQ.
- To plot the difference variable, use a HISTOGRAM statement in PROC UNIVARIATE. Or, use PROC GCHART or PROC CHART for a bar chart.

To check the assumptions (step 4)

- To check the assumption of independent observations, you need to think about your data and whether this assumption seems reasonable. This is not a step where using SAS will answer this question.
- To check the assumption of normality, use PROC UNIVARIATE. For paired groups, check the difference variable.

To perform the test (step 5)

- To perform the paired-difference *t*-test or the Wilcoxon Signed Rank test on a difference variable, use PROC UNIVARIATE.

```
PROC UNIVARIATE DATA=data-set-name;  
VAR difference-variable;
```

data-set-name is the name of a SAS data set, and *difference-variable* is the difference between the pairs of values. If you omit the *difference-variable* from the VAR statement, then the procedure uses all numeric variables. You can specify the TestsForLocation table in an ODS statement, which limits the output to the test results.

- To perform the paired-difference *t*-test with the original pairs, use PROC TTEST and identify the pair of variables in the PAIRED statement.

```
PROC TTEST DATA=data-set-name;  
PAIRED first-variable * second-variable;
```

data-set-name is the name of a SAS data set. SAS subtracts the *second-variable* from the *first-variable* to create the difference. The PAIRED statement is required, and the asterisk in the statement is required. The variables in the PAIRED statement must be numeric and should be continuous variables for a valid *t*-test.

Example

The program below produces the output shown in this chapter:

```
data STA6207;  
  input student exam1 exam2 @@;  
  scorediff = exam2 - exam1 ;  
  label scorediff='Differences in Exam Scores';  
  datalines;  
1 93 98 2 88 74 3 89 67 4 88 92 5 67 83 6 89 90  
7 83 74 8 94 97 9 89 96 10 55 81 11 88 83 12 91 94  
13 85 89 14 70 78 15 90 96 16 90 93 17 94 81  
18 67 81 19 87 93 20 83 91  
;  
run;
```

```
proc univariate data=STA6207;
    var scorediff;
    histogram;
    title 'Summary of Exam Score Differences';
run;

data chromat;
    input hp std @@;
    methdiff=hp-std;
    datalines;
12.1 14.7 10.9 14.0 13.1 12.9 14.5 16.2 9.6 10.2 11.2 12.4
 9.8 12.0 13.7 14.8 12.0 11.8 9.1 9.7
;
run;

ods select TestsForLocation;
proc univariate data=chromat;
    var methdiff;
    title 'Testing for Differences between Chromatography Methods';
run;

proc ttest data=chromat;
    paired hp*std;
    title 'Paired Differences with PROC TTEST';
run;
```

