

Chapter 405

Hotelling's One-Sample T^2

Introduction

The one-sample Hotelling's T^2 is the multivariate extension of the common one-sample or paired Student's t -test. In a one-sample t -test, the mean response is compared against a specific value. Hotelling's one-sample T^2 is used when the number of response variables is two or more, although it can be used when there is only one response variable.

T^2 makes the usual assumption that the data are approximately multivariate normal. *Randomization tests* are provided that do not rely on this assumption. These randomization tests should be used whenever you want exact results that do not rely on several assumptions.

One-Sample Case

The one-sample T^2 is used to test hypotheses about a set of means simultaneously. Specifically, suppose a set of p response variables Y_1, Y_2, \dots, Y_p is measured. Assume that the population is distributed as $N_p(\mu, \Sigma)$, where $N_p(\mu, \Sigma)$ is the p -variable multivariate normal distribution with mean vector μ and covariance matrix Σ . The null hypothesis that $\mu = \mu_0$, where μ_0 is a vector of p constants (often 0's), can be tested using the test statistic

$$T^2 = n(\bar{y} - \mu_0)'S^{-1}(\bar{y} - \mu_0)$$

where \bar{y} is the sample mean vector, n is the sample size, and S^{-1} is the inverse of the sample covariance matrix.

If the null hypothesis that $\mu = \mu_0$ is true, then T^2 follows Hotelling's T^2 distribution. That is, $T^2 \sim T_{p,n-1}^2$. Reject the null hypothesis if $T^2 \geq T_{1-\alpha,p,n-1}^2$. Note that rejecting the null hypothesis concludes that at least one of the p means is not equal to its hypothesized value.

Equality of Means

A second null hypothesis may be of interest. This hypothesis is that all means are equal to each other. This hypothesis also tested using the one-sample T^2 value calculated using the formula

$$T^{2'} = n(C\bar{y})'(CSC')^{-1}(C\bar{y})$$

where C is a contrast matrix of the form

$$C = \begin{bmatrix} 1 & -1 & 0 & 0 \\ 0 & 1 & -1 & 0 \\ 0 & 0 & 1 & -1 \end{bmatrix}$$

Our choice of C tests the hypothesis that all p means are equal. In this case, $T^{2'} \sim T_{p-1,n-1}^2$.

Paired-Sample Case

The one-sample T^2 test may also be applied to the situation in which two samples are to be compared that had a natural pairing between two observation vectors. An example of this pairing occurs when responses are measured on each experimental subject before and after a treatment is administered. Thus, the one-sample T^2 test may be applied in the one-factor repeated measures design.

When such pairing exists, the differences between the first and second measurements are formed—replacing the two observation vectors with one difference vector. This difference vector may then be used in the one-sample T^2 test as described above.

Randomization Test

Because of the strict assumptions that must be made when using this procedure, **NCSS** also includes a randomization test as outlined by Edgington (1987). Randomization tests are becoming more and more popular as the speed of computers allows them to be computed in seconds rather than hours.

A randomization test is conducted by enumerating all possible permutations of the sample data, calculating the test statistic for each permutation, and counting the number of permutations that result in a T^2 value greater than or equal to the actual T^2 value. Dividing this count by the number of permutations tried gives the significance level of the test. Each permutation is found by randomly multiplying each observation by a plus or a minus.

For even moderate sample sizes, the total number of permutations is in the trillions, so a Monte Carlo approach is used in which the permutations are found by random selection rather than complete enumeration. Edgington suggests that at least 1,000 permutations be selected. We suggest that this be increased to 10,000.

Permutation results are provided for the equal covariance case, the unequal covariance case, and for all individual t tests.

Assumptions

The following assumptions are made when using T^2 .

1. The population follows the multivariate normal distribution.
2. The members of the sample are independent.

Data Structure

The data must be entered in a format that places the response variables side by side. An example of the data structure for a paired Hotelling's T^2 design is shown below. In this example, measurements were taken at three points in time before and after a certain drug was administered. Each subject performed strenuous exercise between the first and second measurements of the before set and the after set. This dataset is stored in the file, T2.

T2 Dataset

Before1	Before2	Before3	After1	After2	After3
36	34	30	38	35	29
36	36	28	38	37	27
41	32	29	43	31	25
11	10	8	14	11	10
17	15	13	19	14	12
21	20	18	24	25	17
36	33	30	40	34	28
36	35	34	41	36	30
37	33	28	36	37	29
31	28	25	31	25	26

Example 1 – Paired T^2 Test

This section presents an example of how to run a paired T^2 analysis of the T2 dataset shown earlier. In this analysis, the before and after variables will be compared.

Setup

To run this example, complete the following steps:

1 Open the T2 example dataset

- From the File menu of the NCSS Data window, select **Open Example Data**.
- Select **T2** and click **OK**.

2 Specify the Hotelling's One-Sample T^2 procedure options

- Find and open the **Hotelling's One-Sample T^2** 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 Variables.....	After1-After3
Paired Variables	Before1-Before3
Run Randomization Tests	Checked
Monte Carlo Samples	1000
Random Seed.....	3218505 (for reproducibility)

Reports Tab

All Reports	Checked
SD's and Variances Decimal Places.....	4
Means Decimal Places	4

3 Run the procedure

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

Descriptive Statistics

Descriptive Statistics

Difference	Mean	Standard Deviation
After1 - Before1	2.2	1.7512
After2 - Before2	0.9	2.331
After3 - Before3	-1	2
Count	10	10

This report provides the means, standard deviations, and counts of each difference (or variable). Look the values over to be certain that the right variables were selected.

Hotelling's T^2 Tests

Hotelling's T^2 Tests

Null Hypothesis	T^2	Degrees of Freedom		P-Value	
		DF1	DF2	Parametric Test	Randomization Test
Means All Zero	17.034	3	9	0.0483	0.0630
Means All Equal	9.321	2	9	0.0582	0.0711

The randomization test results are based on 1000 Monte Carlo samples with a random seed of 3218505.

This report gives the results of the two T^2 tests.

Hypothesis

This option indicates the null hypothesis that is tested on this output line. For the case of the paired observations, the first line (Means All Zero) is of most interest.

T^2

The values of T^2 are given here.

DF1

This is the number of response variables.

DF2

This is the degrees of freedom of the covariance matrix which is $n - 1$.

Parametric Test P-Value

This is the p -value of the T^2 test. If this value is less than 0.05, we say that the test was significant at the 0.05 level and at least one pair of means are significantly different. If the value is less than 0.01, we say that the test was significant at the 0.01 level. This result is accurate if the assumptions are met.

Randomization Test P-Value

This is the p -value of the randomization test. The result is accurate even if the response variables were binary.

Individual Difference Influence Tests

Individual Difference Influence Tests

Difference Omitted	Others		Change		Alone	
	T^2	P-Value	T^2	P-Value	T^2	P-Value
After1 - Before1	4.399	0.2036	12.635	0.0424	15.783	0.0032
After2 - Before2	16.484	0.0156	0.550	0.7110	1.491	0.2531
After3 - Before3	16.554	0.0154	0.480	0.7294	2.500	0.1483

This report provides information about the influence of each of the individual response variable differences on the overall T^2 value. This is accomplished by calculating the change in T^2 when a response variable difference is omitted.

Difference Omitted

This is the variable difference shown on this line of the report.

T^2 Others

This is the value of T^2 calculated with all response variables except the variable listed to its left.

P-Value Others

This is the p -value of the associated T^2 test.

T^2 Change

This is the amount that T^2 is reduced when the response variable shown on this line is omitted.

P-Value Change

This is the p -value of the T^2 change test. It is computed using the fact that the change in T^2 is related to an F distribution using the formula

$$F_{\alpha,1,n-p} = \frac{T_p^2 - T_{p-1}^2}{n - 1 + T_{p-1}^2}$$

Note that this quantity tests the drop in T^2 when a variable is removed, conditional on the other response variables that are included. Another way of looking at this quantity is that it tests whether the variable omitted significantly increases the distance between the two populations. See Rencher (1998) page 68 for further details.

Hotelling's One-Sample T^2 T^2 Alone

This is the value of T^2 calculated when only this response variable is used. It is the square of the common one-sample t test. It is the two-sided test of the null hypothesis that the mean for this variable is equal to the hypothesized value (which is usually zero), ignoring all other variables.

P-Value Alone

This is the p -value of the associated T^2 test.

Student's T-Test Section

Student's T-Tests

Difference	T^2 or Student's T	P-Value	
		Parametric Test	Randomization Test
All (T^2)	17.034	0.0483	0.0630
After1 - Before1	3.973	0.0032	0.0130
After2 - Before2	1.221	0.2531	0.2990
After3 - Before3	1.581	0.1483	0.1960

The randomization test results are based on 1000 Monte Carlo samples with a random seed of 3218505.
These individual t -test significance levels should only be used when the overall T^2 value is significant.

This report provides the results of individually conducting a two-sided, paired t -test on each pair of response variables. You might think that since there are a series of p t -tests being employed, a Bonferroni adjustment should be applied to the significance levels. However, if these individual tests are only considered when the overall T^2 is significant at the same level, such as 0.05, then their significance levels are "protected" by the T^2 test and the unadjusted significance levels given here can be used.

Difference

The difference whose results are presented on this line. Note that the first line gives the overall results for T^2 .

 T^2 or |Student's T|

The first line is the value of T^2 . The other lines are the two-sided Student's t -test values.

Parametric Test P-Value

These are the p -values of the test statistics given to the left. Note that if the individual tests are only used when the overall test is significant, these p -values are accurate even though several individual tests are made. The multivariate T^2 test is said to "protect" the significance levels of the individual tests.

Randomization Test P-Value

These are the results of randomization tests that are run on each of the differences. These tests are exact when the Monte Carlo sample size is large, say over 5000. These tests should be used when there is even a hint that the regular assumptions of the t -tests are not valid. For example, this p -value is accurate even when the response variable takes on binary values (the t -test assumes a continuous, normal response variable).

Hotelling's One-Sample T²

Note that these values will change from run to run. As you increase the number of Monte Carlo samples, these values will become more and more stable. You may have to go as large as 100,000 before the results remain the same from run to run. This instability is due to the use of a random sample of all the trillions of permutations that are possible. As you increase the Monte Carlo sample size, you reduce the sampling error (and greatly increase the time it takes to generate the results).

Confidence Intervals for the Mean Differences

Confidence Intervals for the Mean Differences

Difference	Mean	95% Bonferroni Confidence Interval Limits		95% Simultaneous Confidence Interval Limits	
		Lower	Upper	Lower	Upper
After1 - Before1	2.2	0.5756	3.8244	-0.0675	4.4675
After2 - Before2	0.9	-1.2622	3.0622	-2.1182	3.9182
After3 - Before3	-1.0	-2.8552	0.8552	-3.5897	1.5897

This report provides confidence intervals for the mean differences (or the means) for each response variable. Two intervals are provided: Bonferroni and simultaneous.

Difference

The variable difference whose results are presented on this line.

Mean

The actual mean difference for the corresponding response variable(s).

Bonferroni Confidence Interval Limits

Bonferroni confidence intervals are based on the formula

$$\bar{d}_j \pm t_{\alpha/(2p), n-1} \sqrt{\left(\frac{S_{jj}}{n}\right)}$$

This formula is derived by applying a Bonferroni adjustment to the regular univariate confidence interval. This adjustment is made by dividing the alpha level by p , the number of such intervals to be created. These intervals are usually not as wide as the simultaneous intervals, yet still have an appropriate adjustment because of the multiple intervals that are being created.

Simultaneous Confidence Interval Limits

Simultaneous confidence intervals are based on the formula

$$\bar{d}_j \pm \sqrt{T_{1-\alpha, p, n-1}^2 \left(\frac{S_{jj}}{n}\right)}$$

This formula is derived from a formula for confidence intervals for *any* linear combination of the mean differences, including those that are generated after looking at the data. Because of this, these confidence intervals are extra wide and may not be of must use.

Correlation\Covariance Matrix

Correlation\Covariance Matrix

Difference	Difference		
	After1 - Before1	After2 - Before2	After3 - Before3
After1 - Before1	3.0667	0.3556	-2.0000
After2 - Before2	0.0871	5.4333	0.4444
After3 - Before3	-0.5710	0.0953	4.0000

This report displays correlations and covariances of the variables, or differences, analyzed. The correlations are shown in the lower-left half of the matrix and the covariances are shown on the diagonal and in the upper-right half of the matrix.

Example 2 – One-Sample T^2 Test

This section presents an example of how to run a one-sample T^2 analysis of the T2 dataset shown earlier. In this analysis, the analyst wants to test the null hypothesis that the three measurements conform to the response pattern: 30, 33, 30. These values are entered into each row of three new columns: H01, H02, and H03. The analysis will proceed as in the paired test of Example 1.

Setup

To run this example, complete the following steps:

1 Open the T2 example dataset

- From the File menu of the NCSS Data window, select **Open Example Data**.
- Select **T2** and click **OK**.

2 Specify the Hotelling's One-Sample T^2 procedure options

- Find and open the **Hotelling's One-Sample T^2** 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 Variables.....**After1-After3**
Paired Variables**H01-H03**
Run Randomization Tests**Checked**
Monte Carlo Samples**1000**
Random Seed.....**3229881** (for reproducibility)

Reports Tab

All Reports**Checked**
SD's and Variances Decimal Places.....**4**
Means Decimal Places**4**

3 Run the procedure

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

Hotelling's One-Sample T^2

Output

Hotelling's T^2 Tests

Null Hypothesis	T^2	Degrees of Freedom		P-Value	
		DF1	DF2	Parametric Test	Randomization Test
Means All Zero	98.484	3	9	0.0004	0.0030
Means All Equal	51.939	2	9	0.0005	0.0020

The randomization test results are based on 1000 Monte Carlo samples with a random seed of 3229881.

Individual Difference Influence Tests

Difference Omitted	Others		Change		Alone	
	T^2	P-Value	T^2	P-Value	T^2	P-Value
After1 - H01	22.002	0.0071	76.482	0.0060	0.569	0.4698
After2 - H02	98.226	0.0000	0.259	0.9009	2.221	0.1703
After3 - H03	33.000	0.0021	65.485	0.0163	8.079	0.0193

Student's T-Tests

Difference	T^2 or Student's T	P-Value	
		Parametric Test	Randomization Test
All (T^2)	98.484	0.0004	0.0030
After1 - H01	0.755	0.4698	0.4880
After2 - H02	1.490	0.1703	0.2240
After3 - H03	2.842	0.0193	0.0050

The randomization test results are based on 1000 Monte Carlo samples with a random seed of 3229881.
These individual t-test significance levels should only be used when the overall T^2 value is significant.

The significance of the T^2 value indicates that at least one mean does not equal the hypothesized value. A look at the individual t -tests indicates that the significance occurs with the third variable: After3. After1 and After2 are not significantly different from 30 and 33, respectively.