

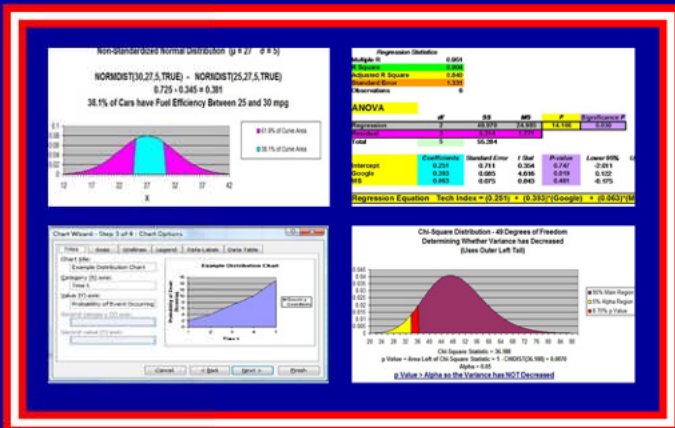
Excel **MASTER** Series

t-Tests

in Excel

The Complete Guide

Excel Statistical Master



Clear and Simple
yet **THOROUGH**
Statistical Instruction for the
Graduate Student and
Business Manager
with
LOTS of Worked-Out Problems
and Screen Shots

Mark Harmon MBA

t-Tests in Excel

The Excel Statistical Master

By Mark Harmon

Copyright © 2011 Mark Harmon

No part of this publication may be reproduced
or distributed without the express permission
of the author.

mark@ExcelMasterSeries.com

www.ExcelMasterSeries.com

ISBN: 978-0-9833070-3-7

Table of Contents

(Click On Chapters To Go To Them)

Excel Hypothesis Tools	4
One Caution About Using t-Tests on Small Samples.....	5
t-Test: Paired Two Sample for Means.....	6
t-Test: Two Sample Assuming Unequal Variance.....	9
t-Test: Two Sample Assuming Equal Variances.....	13
z-Test: Two Sample for Means.....	19
ZTEST.....	25
TTEST.....	27
t Tests in Excel	29
Using the t Test In Excel To Find Out If Your New Marketing Is Working	29
Using the t Test To Find Out If Your New Marketing Is Working.....	29
t Test - General Description.....	29
Two-Sample, Paired t Test.....	29
The t Test - What Is It?.....	30
The Higher the t Value - The More Likely the Groups Are Different.....	30
The Lower the Combined Variance, the Higher the t Value.....	31
T-Test Paired Two Sample for Means.....	32
A Little Bit More About This t Test.....	32
Two-Tailed t Test Is More Stringent.....	34
One-Tailed t Test Is Less Stringent.....	34

Doing The Paired Two-Sample t Test in Excel.....	35
Input the data as follows:.....	36
One-tailed Test.....	37
Two-Tailed Test.....	37
Hand Calculation of the t Value and p Value.....	39
The t Test in Excel Can Determine What Your Best Sales Days Are.....	41
t Test - General Description.....	42
t Test for Two Samples Having Unequal Sizes and Variances.....	42
The t Test - What Is It?.....	43
The Higher the t Value - The More Likely the Groups Are Different.....	43
The Lower the Combined Variance, the Higher the t Value.....	44
T-Test for Two Samples Having Unequal Sizes and Variances.....	45
A Little Bit More About This t Test.....	45
Doing The Two-Sample t Test for Unequal Variances in Excel.....	47
Input the data as follows:.....	52
Hand Calculation of the t Value and p Value.....	54
Correctable Reasons Why Your Data Is not Normally Distributed.....	57
What can you do if your data does not appear to be normally distributed.....	57
The Biggest 7 Correctable Causes of Non-Normality in Data Samples.....	57
1) Outliers	57
2) Data has been affected by more than one process	58
3) Not enough data	58
4) Measuring devices that have poor resolution.....	58
5) A different distribution describes the data	58
6) Data approaching zero or a natural limit	59

7) Only a subset of process' output is being analyzed 59

Excel Hypothesis Tools

Excel offers several data analysis tools that are designed to be used for Hypothesis testing and are described in this module. They are:

t-Test: Paired Two Sample for Means

t-Test: Two-Sample Assuming Equal Variances

t-Test: Two-Sample Assuming Unequal Variances

z-Test: Two Sample for Means

ZTEST

TTEST

The first four tools can be found in Excel 2003 at:

Tools / Data Analysis

(in Excel 2010, Data Analysis tools are found in the Data tab).

ZTEST and **TTEST** are statistical functions found at:

Insert / Function / Statistical

One Caution About Using t-Tests on Small Samples

We will explain exactly how each of these tools is used below. One caution with using any of the t-Test tools:

t-Tests are often performed when only small samples ($n < 30$) are available. This should only be done if it can be proven that the underlying population from which the samples are drawn is Normally distributed. In the real world, the distribution of the underlying population is usually not known. If any t-Test is applied to a small sample drawn from a population of unknown distribution, the answer can be ***totally*** wrong.

Statistical analysis can be applied to large samples using t-Tests. The t distribution approaches the Normal distribution as sample size becomes larger. Statistical tests using the Normal distribution can be applied to large samples ($n > 30$) due to the Central Limit Theorem. Further explanation of this theorem can be found in the course module entitled Confidence Intervals.

Having said all that, here are the three t-Tests and one z-Test that are convenient, built-in statistical analysis tools in Excel:

t-Test: Paired Two Sample for Means

Problem: Google and Yahoo each provide two types of networks that advertisers can run pay-per-click ads on. One type of network is the search network and the other is a content network. One advertiser is running all of his ads exclusively on the search networks of both Google and Yahoo. This advertiser wants to determine with 95% accuracy whether advertising on the content networks of Google and Yahoo will produce more clicks than advertising on the search networks. The advertiser simultaneously ran the same ad on the search and content networks of both Google and Yahoo. This procedure was repeated with 15 different ads.

This is a Hypothesis test because we are determining whether "After Data" is different in either direction or different in only one direction than the "Before" data. We know that this will require a one-tailed test because we are being asked to determine if the "After Data" is different in just one direction (greater) than the "Before" data.

This is classified by Excel as "paired data" because the test is being duplicated on both search engines.

This is a two-sample test because no initial data was available. Samples had to be taken for each of the two means being tested.

The problem requests a 5% Level of Significance which equates to a 95% Level of Certainty. The Level of Significance is often referred to as α (alpha) so in this case, $\alpha = 0.05$.

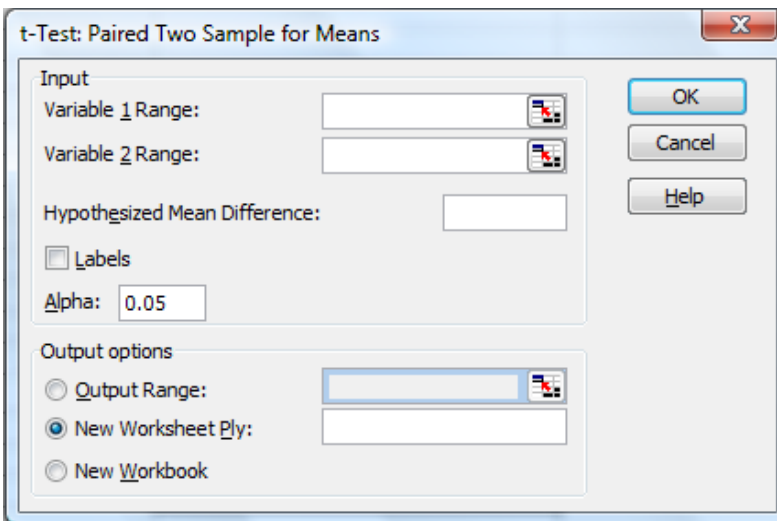
The Null Hypothesis for this problem would state that there is no difference between the number of clicks that the ads generate on the search or content networks.

Expressed another way, the mean difference between the number of clicks on both types of networks is 0.

The Alternate Hypothesis states that ads on the content networks generate more clicks.

This Excel tool is located at:

Tools / Data Analysis / t-Test: Paired Two Sample for Means



The tool's dialog box asks for the following 6 pieces of data:

- 1) **Variable 1 Range:** Variable 1 represents the "Before Data." In this case, the "Before Data" is the Search Network data column highlighted in yellow. "Before Data" would have to be arranged in a column and highlighted in this manner, including the column label on the top row.
- 2) **Variable 2 Range:** Variable 2 represents the "After Data." In this case, the "After Data" is the Content Network data column highlighted in tan. "After Data" would have to be arranged in a column and highlighted in this manner, including the column label on the top row.
- 3) **Hypothesized Mean Difference:** This is normally 0 for a Hypothesis test of paired data. The Hypothesized Mean Difference will correspond to whatever is on the right side of the equal sign for the Null Hypothesis.
- 4) **Includes Labels?** Yes, if you included them when the columns were highlighted, No, if they were not included when the data columns were highlighted. You must either highlight and include all labels or no labels.
- 5) **Alpha** - The Level of Significance equals alpha, so $\alpha = 0.05$

6) **Output Range** - The output range must have at least 14 rows and 3 columns free.

Click OK and the following output will be generated:

t-Test: Paired Two Sample for Means		
	<i>Search Network</i>	<i>Content Network</i>
Mean	37,058	40,267
Variance	22,513,133	41,031,253
Observations	30	30
Pearson Correlation	0.76	
Hypothesized Mean Difference	0	
df	29	
t Stat	-4.25	
P(T<=t) one-tail	0.00010	
t Critical one-tail	1.70	
P(T<=t) two-tail	0.00	
t Critical two-tail	2.05	

Interpretation of the data is very simple: Compare the p-Value with the Level of Significance (α). If the p-Value is less than α , then the Null Hypothesis is rejected and the Alternative Hypothesis is accepted.

This is a one-tailed test so the p-value shown for the one-tailed test is the correct p-Value. In this case, the one-tailed p-Value is 0.00010, which is less than the level of significance ($\alpha = 0.05$) so the Null Hypothesis is rejected, and the Alternate Hypothesis is accepted, which states that the content network produces more clicks than the search network.

t-Test: Two Sample Assuming Unequal Variance

Problem: Evaluate the returns of these two stocks to determine if there is a real difference. Use a 0.05 level of significance. The returns are shown here:

<u>IBM</u>	<u>Apple</u>
0.7541	-4.6296
14.9701	18.986
11.9792	-1.7226
7.907	-0.5535
-5.1724	6.679
3.4091	1.8261
0.7541	-4.6296
14.9701	18.986
11.9792	-1.7226
7.907	-0.5535
-5.1724	6.679
3.4091	1.8261
0.7541	-4.6296
14.9701	18.986
11.9792	-1.7226
7.907	-0.5535
-5.1724	6.679
3.4091	1.8261
3.4091	1.8261
0.7541	-4.6296
14.9701	18.986
11.9792	-1.7226
7.907	-0.5535
-5.1724	6.679
3.4091	1.8261
0.7541	-4.6296
14.9701	18.986
11.9792	-1.7226
7.907	-0.5535
-5.1724	6.679

This is a Hypothesis test because we are determining whether the mean of one population is different in either direction or different in only one direction than the mean of another population. We know that this will require a two-tailed test because we are being asked to determine if the population means are simply different.

This is a two-sample test because no initial data was available. Samples had to be taken for each of the two means being tested.

We do not initially know the variances of each data set so we Assume them to be unequal. For that reason, this is classified as

An Unequal Variance test.

The problem requests a 5% Level of Significance which equates to a 95% Level of Certainty. The Level of Significance is often referred to as α (alpha) so in this case, $\alpha = 0.05$.

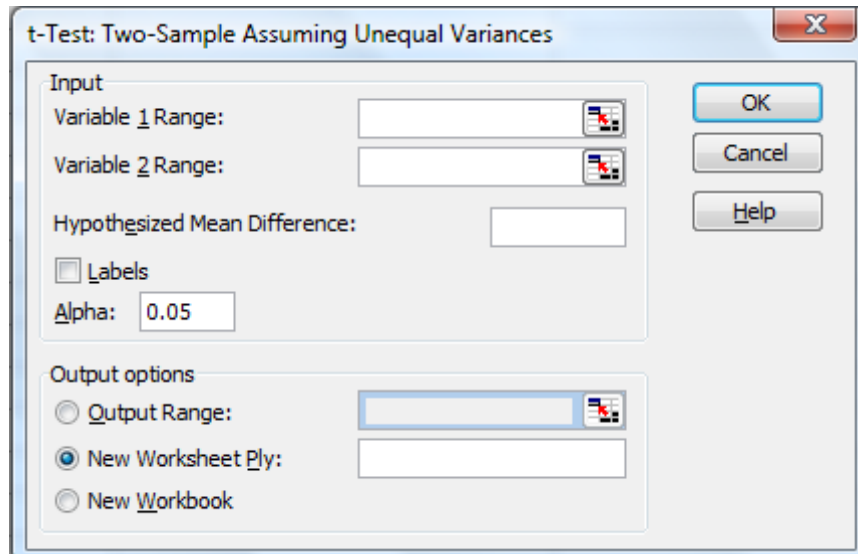
The Null Hypothesis for this problem would state that there is no difference between the average means of the two independent populations. Expressed another way, the mean difference between the means = 0.

The Alternate Hypothesis states that the difference between the average means does not equal 0.

On the following page is the Excel dialogue box used to solve the problem:

The tool is located at:

Tools / Data Analysis / t-Test: Two Sample Assuming Unequal Variances



The tool's dialog box asks for the following 6 pieces of data:

- 1) **Variable 1 Range:** Variable 1 represents Population 1 Data. In this case, the Population 1 Data is the IBM return column highlighted in yellow. Population 1 Data would have to be arranged in a column and highlighted in this manner, including the column label on the top row.
- 2) **Variable 2 Range:** Variable 2 represents Population 2 Data. In this case, the Population 2 Data is the Apple return column highlighted in tan. Population 2 Data would have to be arranged in a column and highlighted in this manner, including the column label on the top row.
- 3) **Hypothesized Mean Difference:** This is normally 0 for a Hypothesis test of paired data. The Hypothesized Mean Difference will correspond to whatever is on the right side of the equal sign for the Null Hypothesis.
- 4) **Includes Labels?** Yes, if you included them when the columns were highlighted, No, if they were not included when the data columns were highlighted. You must either highlight and include all labels or no labels.
- 5) **Alpha** - The Level of Significance equals alpha, so $\alpha = 0.05$
- 6) **Output Range** - The output range must have at least 14 rows and 3 columns free.

Click OK and the following output will be generated:

t-Test: Two-Sample Assuming Unequal Variances		
	<i>IBM</i>	<i>Apple</i>
Mean	5.64	3.43
Variance	47.95	62.50
Observations	30	30
Hypothesized Mean Difference	0	
df	57	
t Stat	1.15	
P(T<=t) one-tail	0.13	
t Critical one-tail	1.67	
P(T<=t) two-tail	0.25	
t Critical two-tail	2.00	

Interpretation of the data is very simple: Compare the p-Value with the Level of Significance (α). If the p-Value is less than α , then the Null Hypothesis is rejected and the Alternate Hypothesis is accepted.

This is a two-tailed test so the p-value shown for the two-tailed test is the correct p-Value. In this case, the two-tailed p-Value is 0.25, which is greater than the level of significance ($\alpha = 0.05$) so the Null Hypothesis is not rejected and it hasn't been disproven that the means of the two populations are the same.

t-Test: Two Sample Assuming Equal Variances

This test assumes that the variances of the populations being compared are equal. Usually it is the case that the variances are not known. This test would only be valid for comparing populations of known variances.

Problem: A company is testing batteries from 2 suppliers. Below are listed the hours of usage before each sample ran out. Determine using a 0.05 level of significance whether the new supplier's batteries really last longer than the old supplier's. The variances of battery life of batteries from both suppliers are known to be equal. The following page provides the sample data:

Battery Suppliers

Old Supplier	New Suppliers
61	51
49	42
56	37
51	45
38	65
44	47
61	51
51	75
50	49
60	56
39	52
51	49
43	69
37	53
45	51
65	61
51	49
46	56
50	49
44	58
26	52
52	48
43	79
46	42
64	46
53	52
38	51
44	61
61	49
44	56
42	52
46	49
64	
53	
38	

This is a Hypothesis test because we are determining whether "After Data" is different in either direction or different in only one direction than the "Before" data. We know that this will require a one-tailed test because we are being asked to determine if the "After Data" is different in just one direction (greater) than the "Before" data.

This is a two-sample test because no initial data was available. Samples had to be taken for each of the two means being tested.

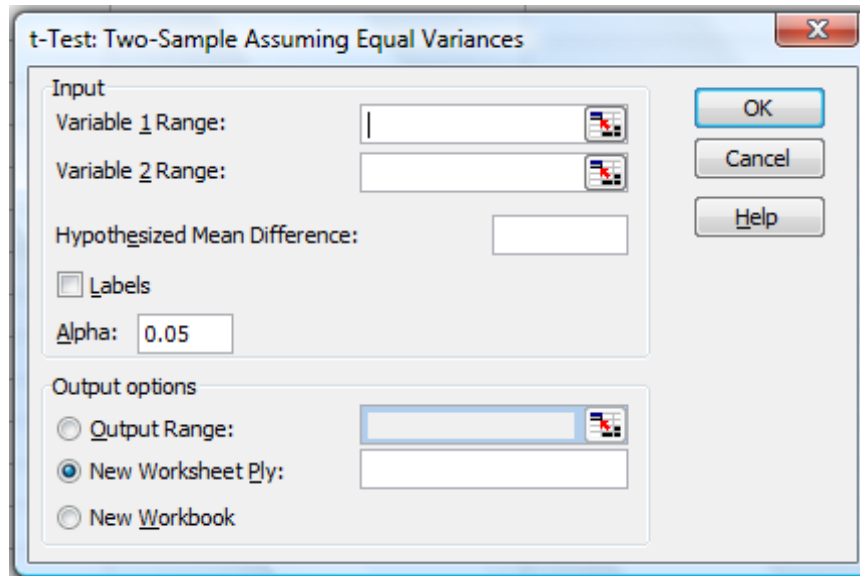
The problem requests a 5% Level of Significance which equates to a 95% Level of Certainty. The Level of Significance is often referred to as α (alpha) so in this case, $\alpha = 0.05$.

The Null Hypothesis for this problem would state that there is no difference between the mean lifetimes of batteries from the old or new suppliers. Expressed another way, the mean difference between mean battery lifetime from the old supplier and from the new supplier is 0.

The Alternate Hypothesis states that batteries from the new supplier last longer. The following page provides the dialogue box use to solve the problem.

The tool is located at:

Tools / Data Analysis / t-Test: Two-Sample Assuming Equal Variances



The tool's dialog box asks for the following 6 pieces of data:

- 1) **Variable 1 Range:** Variable 1 represents the "Before Data." In this case, the "Before Data" is the Old Supplier data column highlighted in yellow. "Before Data" would have to be arranged in a column and highlighted in this manner, including the column label on the top row.
- 2) **Variable 2 Range:** Variable 2 represents the "After Data." In this case, the "After Data" is the New Supplier data column highlighted in tan. "After Data" would have to be arranged in a column and highlighted in this manner, including the column label on the top row.
- 3) **Hypothesized Mean Difference:** This is normally 0 for a Hypothesis test of paired data. The Hypothesized Mean Difference will correspond to whatever is on the right side of the equal sign for the Null Hypothesis.
- 4) **Includes Labels?** Yes, if you included them when the columns were highlighted, No, if they were not included when the data columns were highlighted. You must either highlight and include all labels or no labels.
- 5) **Alpha** - The Level of Significance equals alpha, so $\alpha = 0.05$
- 6) **Output Range** - The output range must have at least 14 rows and 3 columns free.

Click OK and the following output will be generated:

t-Test: Two-Sample Assuming Equal Variances		
	<i>Old Suppliers</i>	<i>New Suppliers</i>
Mean	48.74	53.19
Variance	81.31	81.51
Observations	35	32
Pooled Variance	81.41	
Hypothesized Mean μ	0	
df	65	
t Stat	-2.01	
P(T<=t) one-tail	0.02	
t Critical one-tail	1.67	
P(T<=t) two-tail	0.05	
t Critical two-tail	2.00	

Interpretation of the data is very simple: Compare the p-Value with the Level of Significance (α). If the p-Value is less than α , then the Null Hypothesis is rejected and the Alternate Hypothesis is accepted.

This is a one-tailed test so the p-value shown for the one-tailed test is the correct p-Value. In this case, the one-tailed p-Value is 0.02, which is less than the level of significance ($\alpha = 0.05$) so the Null Hypothesis is rejected and the Alternate Hypothesis is accepted, which states that the batteries from the new supplier have a longer average life.

t-Test: Two-Sample Assuming Equal Variances

	<i>Old Suppliers</i>	<i>New Suppliers</i>
Mean	48.74	53.19
Variance	81.31	81.51
Observations	35	32
Pooled Variance	81.41	
Hypothesized Mean μ	0	
df	65	
t Stat	-2.01	
P(T<=t) one-tail	0.02	
t Critical one-tail	1.67	
P(T<=t) two-tail	0.05	
t Critical two-tail	2.00	

z-Test: Two Sample for Means

We recommend the use of this z-Test data analysis tool, provided that all sample sizes are large ($n > 30$). As a reminder of why large samples are necessary, statistics' most basic theorem, the Central Limit Theorem, requires large sample sizes in order to plot sample averages on a Normal curve, regardless of the underlying population distribution. Another way of expressing that would be to say that sample means will be Normally distributed no matter how the underlying population is distributed, as long as sample size is large ($n > 30$).

Here we will apply the Excel data analysis tool z-Test: Two Sample for Means to the same problem as above, since each of the two samples is large ($n > 30$).

Problem: A company is testing batteries from 2 suppliers. Below are listed the hours of usage before each sample stopped working. Determine using a 0.05 level of significance whether the new supplier's batteries really last longer than the old supplier's. The data is located on the next page:

Battery Suppliers

Old Supplier	New Suppliers
61	51
49	42
56	37
51	45
38	65
44	47
61	51
51	75
50	49
60	56
39	52
51	49
43	69
37	53
45	51
65	61
51	49
46	56
50	49
44	58
26	52
52	48
43	79
46	42
64	46
53	52
38	51
44	61
61	49
44	56
42	52
46	49
64	
53	
38	

This is a Hypothesis test because we are determining whether "After Data" is different in either direction or different in only one direction than the "Before" data. We know that this will require a one-tailed test because we are being asked to determine if the "After Data" is different in just one direction (greater) than the "Before" data.

This is a two-sample test because no initial data was available. Samples had to be taken for each of the two means being tested.

The problem requests a 5% Level of Significance which equates to a 95% Level of Certainty. The Level of Significance is often referred to as α (alpha) so in this case, $\alpha = 0.05$.

The Null Hypothesis for this problem would state that there is no difference between the mean lifetimes of batteries from the

old or new suppliers. Expressed another way, the mean difference between mean battery lifetime from the old supplier and from the new supplier is 0.

The Alternate Hypothesis states that batteries from the new supplier last longer.

This is unpaired data with a large number of samples (>30). The Excel data analysis tool for this would be the z-Test:

Two-Sample for Means. Use this Excel data analysis tool only when sample size is greater than 30. The samples do not have to be the same size. On the following page is the dialogue to solve the problem.

The tool is located at:

Tools / Data Analysis / z-Test: Two-Sample for Means

The Excel data analysis tool requires us to determine the variances of each sample first. Do that with the following Excel function:

= VAR (Yellow Highlighted Data) = 81.3

= VAR (Tan Highlighted Data) = 81.5

The tool's dialog box asks for the following 8 pieces of data:

- 1) **Variable 1 Range:** Variable 1 represents the "Before Data." In this case, the "Before Data" is the Old Supplier data column highlighted in yellow. "Before Data" would have to be arranged in a column and highlighted in this manner, including the column label on the top row.
- 2) **Variable 2 Range:** Variable 2 represents the "After Data." In this case, the "After Data" is the New Supplier data column highlighted in tan. "After Data" would have to be arranged in a

column and highlighted in this manner, including the column label on the top row.

- 3) **Hypothesized Mean Difference:** This is normally 0 for a Hypothesis test of paired data. The Hypothesized Mean Difference will correspond to whatever is on the right side of the equal sign for the Null Hypothesis.
- 4) **Variance of Variable 1** - We have calculated this above to be: 81.3
- 5) **Variance of Variable 2** - We have calculated this above to be: 81.5
- 6) **Includes Labels?** Yes, if you included them when the columns were highlighted, No, if they were not included when the data columns were highlighted. You must either highlight and include all labels or no labels.
- 7) **Alpha** - The Level of Significance equals alpha, so $\alpha = 0.05$
- 8) **Output Range** - The output range must have at least 12 rows and 3 columns free.

Click OK and the following output will be generated:

z-Test: Two Sample for Means		
	<i>Old</i>	<i>New</i>
Mean	48.74	53.19
Known Variance	81.3	81.5
Observations	35	32
Hypothesized Mean Difference	0	
z	-2.01	
P(Z<=z) one-tail	0.02199871	
z Critical one-tail	1.64	
P(Z<=z) two-tail	0.04	
z Critical two-tail	1.96	

Interpretation of the data is very simple: Compare the p-Value with the Level of Significance (α). If the p-Value is less than α , then the Null Hypothesis is rejected and the Alternate Hypothesis is accepted.

This is a one-tailed test so the p-value shown for the one-tailed

test is the correct p-Value. In this case, the one-tailed p-Value is 0.02199871, which is less than the level of significance ($\alpha = 0.05$) so the Null Hypothesis is rejected and the Alternate Hypothesis is accepted, which states that the batteries from the new supplier have a longer average lifetime.

The previous four Excel Hypothesis Test tools are located at:

Tools / Data Analysis (Data Analysis Tools in Excel 2010 Are Found in the Data tab)

We will now discuss two Hypothesis test tools that are statistical functions. They can be found at:

Insert / Function / Statistical

The z-Test and t-Test are found here and are the Excel functions ZTEST and TTEST. Each will be discussed on the following pages:

ZTEST

The ZTEST function should not be used because it will produce an incorrect answer in certain circumstances. The ZTEST attempts to perform a Hypothesis test to determine whether a sample mean is statistically different from a population mean. The ZTEST attempts to calculate a p Value for sample mean. The user highlights a sample that is the first input of the function. ZTEST then calculates the mean of the sample and uses that information along with the other user inputs to the function (the population mean and the population or sample standard deviation) to calculate a p Value for a one-tailed, one-sample Hypothesis test.

Beware that the ZTEST only calculates a correct p Value when the sample mean is greater than the population mean.

The ZTEST calculates an incorrect p Value if the sample mean is less than the population mean.

It is highly recommended to follow instructions in the Hypothesis Testing module and NOT use the ZTEST function when performing a Hypothesis test.

Here is an explanation of why the ZTEST produces an incorrect p Value when the sample mean is less than the population mean:

The format of ZTEST is:

ZTEST (Array, μ , (σ or S))

The ZTEST function calculates the one-tailed p-Value for a one-sample Hypothesis test. The Array is the sample drawn from a population whose mean is hypothesized to be μ , The third parameter is the population standard deviation, σ , or the sample standard deviation, s, if the population standard deviation is not known.

The ZTEST then calculates a Z Score (z) by the following method:

$$\text{Z Score} = z = [(X_{\text{avg}} - \mu) / ((\sigma \text{ or } s) / \text{SQRT} (n))]$$

The Z Score, z, is the number of standard errors that X_{avg} , the sample average, is from the population mean, μ .

The ZTEST is the equivalent of the following:

$$\begin{aligned} \text{ZTEST} (\text{Array}, \mu, (\sigma \text{ or } s)) &= 1 - \text{NORMSDIST} [(\bar{X}_{\text{avg}} - \mu) / ((\sigma \text{ or } s) / \text{SQRT} (n))] \\ &= 1 - \text{NORMSDIST}(z) \end{aligned}$$

$1 - \text{NORMSDIST}(z)$ calculates the value of the area under the Normal curve that is to the **RIGHT** of z . This would be the correct p Value if the sample mean is greater than the population mean and the p Value therefore occurs in the outer right tail.

If the sample mean is less than the population mean, the p Value will occur in the outer left tail. In this case, the p Value would be the area to the LEFT of z , which would be calculated as

$\text{NORMSDIST}(z)$, not $1 - \text{NORMSDIST}(z)$ that is calculated by ZTEST.

You are therefore recommended to not use ZTEST when performing a Hypothesis test. If you follow the Hypothesis testing instructions provided in the Hypothesis Testing modules in this course, you will correctly perform a Hypothesis test and obtain correct p Values in all cases.

TTEST

The TTEST function duplicates the one or two-tailed p Values that are part of the output of each of the three t-Test Data Analysis tools that are described previously. There is really no reason to use the TTEST function because the TTEST output is just one part of much more complete output that each of the t-Test Data Analysis tools provides. The t-Test Data Analysis tools described previously provide a much more complete output, but require nearly the same amount of input information as the TTEST. This will be shown below:

The TTEST has the following format:

TTEST (Array 1, Array 2, Number of Tails, Type of Test)

The arrays are highlighted by the user, just as they are for the t-Test Data Analysis tools above.

The Number of Tails can be 1 or 2.

The Type of Test has the following three choices:

- 1) t-Test: Paired Two Sample for Means
- 2) t-Test: Two-Sample Assuming Equal Variances
- 3) t-Test: Two-Sample Assuming Unequal Variances

It is shown on the next page that the TTEST output is only a small part of the output of any of the t-Test Data Analysis tools.

	A	B	C	D	E	F	G	H	I	J	K	L
1												
2												
3												
4		Array 1	Array 2		t-Test: Paired Two Sample for Means				TTEST Outputs			
5		3	6			Array 1	Array 2					
6		4	19		Mean	4.555555556	7.888888889					
7		5	3		Variance	6.777777778	47.11111111					
8		8	2		Observations	9	9					
9		9	14		Pearson Correlation	0.101820087						
10		1	4		Hypothesized Mean Difference	0						
11		2	5		df	8						
12		4	17		t Stat	-1.410691232						
13		5	1		P(T<=t) one-tail	0.098007892			0.098007892	=TTEST(B5:B13,C5:C13,1,1)		
14					t Critical one-tail	1.859548033						
15					P(T<=t) two-tail	0.196015785			0.196015785	=TTEST(B5:B13,C5:C13,2,1)		
16					t Critical two-tail	2.306004133						
17												
18												
19												
20												
21												
22					t-Test: Two-Sample Assuming Equal Variances							
23		Array 1	Array 2			Array 1	Array 2					
24		3	6		Mean	4.555555556	7.888888889					
25		4	19		Variance	6.777777778	47.11111111					
26		5	3		Observations	9	9					
27		8	2		Pooled Variance	26.94444444						
28		9	14		Hypothesized Mean Difference	0						
29		1	4		df	16						
30		2	5		t Stat	-1.362229828						
31		4	17		P(T<=t) one-tail	0.095997943			0.095997943	=TTEST(B24:B32,C24:C32,1,2)		
32		5	1		t Critical one-tail	1.745883669						
33					P(T<=t) two-tail	0.191995887			0.191995887	=TTEST(B24:B32,C24:C32,2,2)		
34					t Critical two-tail	2.119905285						
35												
36												

t Tests in Excel – What They Are and How To Do Them

Using the t Test To Find Out If Your New Marketing Is Working - Simplified and Done in Excel

If you had recently launched a new marketing campaign, you would want to know as soon as possible whether the campaign was working. If you are able to take a large sample of before and after measurements (for example, sales in all of the sales territories), Excel has the perfect tool for you - the two-sample paired t-test for means. It is very simple to use and the output is straight-forward and easy to interpret.

t Test - General Description

This test will tell you whether the difference between the before and after numbers is genuine or whether this difference could merely have been the result of chance. Overall a t-test compares two means and determines within a specified degree of certainty whether the two means really are different, or whether the difference might have occurred by chance.

Two-Sample, Paired t Test

The two-sample paired t-test for means evaluates whether the average difference between the before and after measurements is greater than zero or not. In other words, this test evaluates within a specified degree of certainty whether the average measured difference between before and

after is real or could have occurred merely by chance.

Before we start discussing this specific test in detail, The t-test needs to be generally explained. The basic question to be answered is:

The t Test - What Is It?

The t test is a statistics test generally used to test whether means of populations are different. In the t test, a t value is calculated based upon the difference in the means and variances of the two populations. The greater the t value, the more certain it is that the means are different.

The t value can be generally described as follows:

t value = (Difference between the group means) / (Variability of the groups)

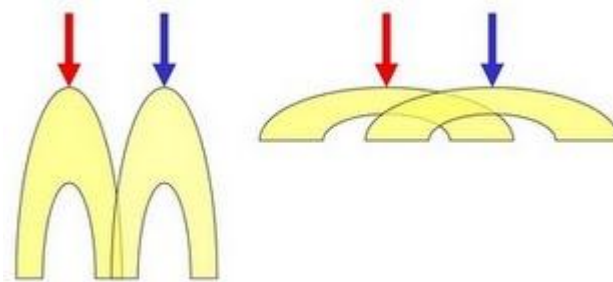
There are many variations of the t test. Each has its own specific formula for calculating a t value for the sampled data sets. All of the t value formulas can be described by the above formula.

The Higher the t Value - The More Likely the Groups Are Different

The higher the t value is, the more likely it is that the two means are different. If the two groups being compared have a high degree of variance (t value has a high denominator), it is much harder to tell them apart. On the other hand, if the two groups being compared have a low degree of variance (the t value has a low denominator), it is much easier to tell the two groups apart.

The Lower the Combined Variance, the Higher the t Value

The illustrations below should clarify how the degree of variance in the two groups determines how easy or difficult it is to state that the means of the two groups are really different. The t test quantifies this relationship and provides a way to determine whether the measured difference between two means can be considered real or not based upon the amount of variance in both groups. Here are illustrations that should clarify this relationship.



We can see that pair of data sets on the right are much easier to differentiate because they have much less overlap than the pair of data sets on the left. The overlap represents the overall variability between the two data sets in each pair. The higher the total variability within the pair of data sets, the higher will be the denominator in the t value formula. The higher the denominator, the lower the t value for the pair of data sets. The lower the t value, the less likely it is that the two data sets are separate data sets with different means.

T-Test Paired Two Sample for Means

A paired t test or paired difference t test is use to determine whether the average of the "before" and "after" measurements taken of a single set of objects is the same. The Null Hypothesis being tested states that there is no difference between the average "before" and "after" measurements. Specifically, the Null Hypothesis states that the mean of all "after" measurements minus the mean of all "before" measurements taken of the same objects equals 0.

We are going to use the paired t test to determine within 95% certainty whether the average sales from a group of sales territories increased after a new marketing program was implemented. We will simply measure the before and after sales from each territory and apply this t test using Excel to get our result.

A Little Bit More About This t Test

The t Test in general is a special case of one-way (sometimes called "single factor") ANOVA. This paired two-sample student's t test is applied when there is a natural pairing of samples. It is most often used to determine whether "before" and "after" means of a sample of the same objects have changed during an experiment. One really great thing about this t test is that the paired two-sample t test does not require that the variances of both populations to be the same.

To sum up the paired two-sample student's t test, a single t value is calculated from data from both samples. Here is the formula to calculate the t value for a paired two-sample student's t test if you are testing to

determine whether the difference between two means is greater than zero:

$$t \text{ value} = \frac{\text{Average Difference Between Each Pair}}{[\text{Stan. Dev. Of Average Differences} / \text{SQRT}(n)]}$$

You can see that this follows the general formula for calculating the t value in a t test, which is:

$$t \text{ value} = (\text{Difference between the group means}) / (\text{Variability of the groups})$$

The t value is a specific point on the x-axis in the t distribution

(student's t distribution). If this t value falls outside the region of required certainty, it can be stated that the two means are probably different. If this t value falls within the region of required certainty, it cannot be stated that the two means are probably different.

The required region of certainty depends upon the degree of certainty required in the test. If 95% certainty is required, then the required region of certainty consists of 95% of the area under the student's t distribution. The outer 5% is the region of uncertainty. This is also referred to as α (alpha) or the degree of significance. If the t value is large enough to be located all the way out on the x-axis in the 5% region of uncertainty, it can be stated within 95% certainty that the two means are different.

A t test can be a one-tailed test or a two-tailed test. A one-tailed test determines whether the means are different in one specific direction. For example, a one-tailed test could be used to determine only if the mean of the "after" measurements is greater than the mean of the "before" measurements. A two-tailed test determines whether the two means are merely different.

Two-Tailed t Test Is More Stringent

The two-tailed test is more stringent because the area in the outer tails outside of the region of required degree of certainty is split into two tails. For example, if the required degree of certainty is 95% on a two-tailed test, the calculated t value must be all the way out in the outer 2.5% of either tail for the t test to conclude within 95% certainty that the means are different.

One-Tailed t Test Is Less Stringent

A one-tailed test is less stringent. If the required degree of certainty is 95% on a one-tailed test, the calculated t value only has to be within the outer 5% of whatever tail is being tested to be able to state the two means are probably different.

Doing The Paired Two-Sample t Test in Excel

We are testing to determine whether a new marketing campaign has increased sales in a group of six sales territories. In this case the sample size (n) equals 6. For this type of t test, the degrees of freedom = $n - 1 = 5$.

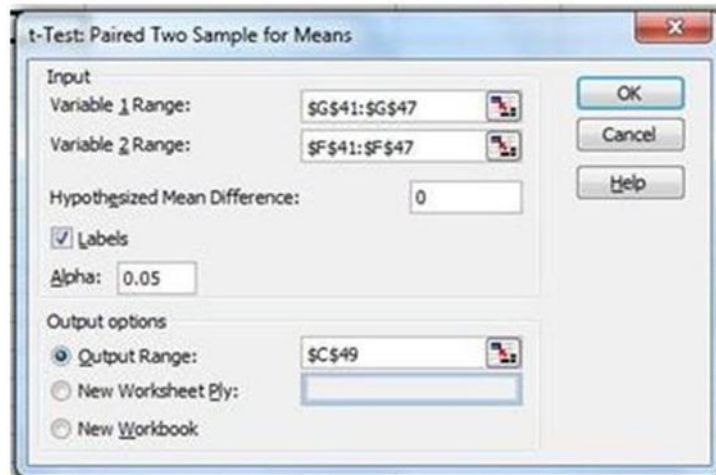
The data need to be arranged in Excel as follows:

Sales Territory	Variable 2 Range	Variable 1 Range
	Sales Before New Ads	Sales After New Ads
1	10	12
2	3	15
3	5	9
4	6	7
5	3	9
6	5	6

Now, access this Excel t Test as follows (this is Excel 2003):

Tools / Data Analysis / t-Test: Paired Two Sample for Means
(Data Analysis Tools in Excel 2010 Are Found in the Data tab)

This following dialogue box will appear:



Input the data as follows:

Variable 1 Range: Select everything that is highlighted light blue, including the label “Sales After New Ads.” If you are trying to determine whether the “after” measurements have gone up, the “after” data is Input Variable 1. If you are trying to determine whether the “after” measurements have gone down, the “after” data is Input Variable 2.

Variable 2 Range: Select everything that is highlighted in yellow, including the label “Sales Before New Ads.”

Hypothesized Mean Difference: 0

Labels: Check the box because you included the labels for Variables 1 and 2.

Alpha: This depends on your desired degree of certainty. 0.05, if you desired 95% certainty. 0.20 if you desire 80% certainty.

Output Range: Select the cell that you want the upper left corner of the output to appear in.

Hit “OK” to run the analysis and the following Excel output appears:

t-Test: Paired Two Sample for Means		
	<i>Sales After New Ads</i>	<i>Sales Before New Ads</i>
Mean	9.666666667	5.333333333
Variance	11.06666667	6.666666667
Observations	6	6
Pearson Correlation	-0.007761505	
Hypothesized Mean Difference	0	
df	5	
t Stat	2.511	
P(T<=t) one-tail	0.027	
t Critical one-tail	2.015	
P(T<=t) two-tail	0.054	
t Critical two-tail	2.571	

This output can be interpreted as follows:

The t value is 2.511.

One-tailed Test

This t value is greater than the critical t value for a one-tailed test (2.015). We can therefore state with 95% certainty that the mean sales has increased as a result of the new marketing campaign.

The above conclusion can also be reached because the p Value for the one-tailed test (highlighted in light blue on the Excel output) is 0.027. This is less than alpha (0.05). The p Value being less than alpha is an equivalent result to the t value being greater than the t critical value.

Two-Tailed Test

A different result is arrived at for the two-tailed test. The two-tailed test is more stringent because the alpha region of uncertainty (5% of the area

under the student's t distribution curve) is now divided between both outer tails. The t value needs to be larger for the two-tailed test to wind up in the outer 2.5% area of either outer tail.

In this case, the t value was not large enough to be positioned in the outer 2.5% of either outer tail. The t value (2.511) is smaller than the critical t value for the two-tailed test (2.571). This indicates that it cannot be stated with 95% certainty that there has been a change in the mean from before to after.

The p value calculated for the two-tailed test (0.054) is larger than alpha (0.05). This is an equivalent result to the above.

Hand Calculation of the t Value and p Value

Let's calculate the t value and p values for the one and two-tailed tests by hand to make sure that Excel has done a correct job. The t value is stated as the t statistic.

Here is the original test data:

Sales Territory	Variable 2 Range	Variable 1 Range	Difference	
	Sales Before New Ads	Sales After New Ads		
1	10	12	2	
2	3	15	12	
3	5	9	4	
4	6	7	1	
5	3	9	6	
6	5	6	1	
			4.333	Mean of Differences
			4.227	Standard Deviation of Differences
			6	n = Count
			5	Degrees of Freedom = n - 1

Here is the hand calculation of the t value and p values for the one and two-tailed tests for this Paired Two-Sample t Test. The hand calculation agrees with the Excel outputs. There are very slight differences due to rounding differences:

4.333	Mean of Differences
4.227	Standard Deviation of Differences
6	n = Count
5	Degress of Freedom = n - 1

$$[\text{Mean of Difference}] / [\text{Stan. Dev. of Difference} / \text{SQRT}(n)] = t \text{ Statistic}$$

$$[4.333] / [4.227 / \text{SQRT}(6)] = t \text{ Statistic}$$

$$2.5209 = t \text{ Statistic}$$

$$\text{TDIST}(t \text{ Statistic}, \text{Degrees of Freedom}, \# \text{ tails}) = P(T \leq t) \text{ one-tail}$$

$$\text{TDIST}(2.5209, 5, 1) =$$

$$0.027 = P(T \leq t) \text{ one-tail}$$

$$\text{TDIST}(t \text{ Statistic}, \text{Degrees of Freedom}, \# \text{ tails}) = P(T \leq t) \text{ two-tail}$$

$$\text{TDIST}(2.5209, 5, 2) =$$

$$0.053 = P(T \leq t) \text{ two-tail}$$

The Paired Two-Sample t Test is a very simple test to run and can be applied to nearly any aspect of your marketing program to see if a single change affected a large number of elements whose before and after measurements can be taken. One note: the before and after measurements must be continuous and using the same scale.

The t Test in Excel

Can Determine What Your Best Sales Days Are

It's always great to know what day of the week you can expect to get peak sales. The t test in Excel can provide that information for you. It's quite a simple test to run, as you will see. This blog article will walk you step-by-step through a t Test in Excel. The t Test compares two groups of samples and determines whether the mean of one sample is different than the other. Most types of t Tests require that each sample group have the same number of samples and also have the same variance. This Excel t Test has neither of those requirements.

Here is the scenario we are going to test: Suppose that you have been monitoring your daily sales for about a year. Your two best sales during the week are normally Monday and Wednesday. You would like to know which of those two days really does produce the best sales. You've tracked Monday sales for the last 40 weeks and Wednesday sales for the last 42 weeks. Mean sales for Wednesday is a bit higher than mean sales for Monday, but you would like to know with 95% certainty whether the difference in means is not just by chance and that Wednesday really is a better day for selling.

You can run your data through an Excel t Test and know within a minute whether Wednesday really is the best sales day. Excel has several built-in t Tests. The specific Excel t Test we will use is called the "t Test: Two Sample Assuming Unequal Variances." This t Test allows for two samples that have unequal sizes and variances. The only requirement is that both samples are Normally distributed. This will be discussed shortly.

In Excel 2003, this test can be accessed through this menu path: **Tools / Data Analysis / t Test: Two Sample Assuming Unequal Variances** (**Data Analysis tools in Excel 2010 are located in the Data tab**). Before we perform this t Test, we need to have a discussion of what the t Test is.

t Test - General Description

This test will tell you whether the difference between the before and after numbers is genuine or whether this difference could merely have been the result of chance. Overall a t-test compares two means and determines within a specified degree of certainty whether the two means really are different, or whether the difference might have occurred by chance.

t Test for Two Samples Having Unequal Sizes and Variances

The t Test that can be applied to two samples with unequal sizes and unequal variances determines whether the means of both samples are the same. In other words, this test evaluates within a specified degree of certainty whether the measured difference between the means is real or could have occurred merely by chance.

Before we start discussing this specific test in detail, The t-test needs to be generally explained. The basic question to be answered is:

The t Test - What Is It?

The t test is a statistics test generally used to test whether means of populations are different. In the t test, a t value is calculated based upon the difference in the means and variances of the two populations. The greater the t value, the more certain it is that the means are different.

The t value can be generally described as follows:

t value = (Difference between the group means) / (Variability of the groups)

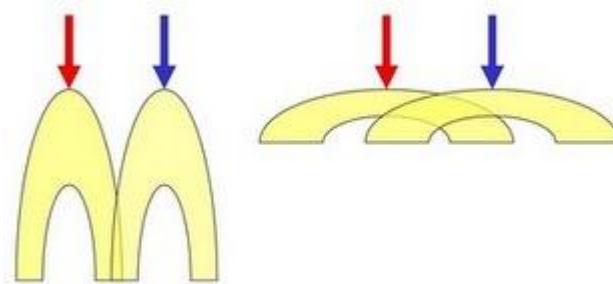
There are many variations of the t test. Each has its own specific formula for calculating a t value for the sampled data sets. All of the t value formulas can be described by the above formula.

The Higher the t Value - The More Likely the Groups Are Different

The higher the t value is, the more likely it is that the two means are different. If the two groups being compared have a high degree of variance (t value has a high denominator), it is much harder to tell them apart. On the other hand, if the two groups being compared have a low degree of variance (the t value has a low denominator), it is much easier to tell the two groups apart.

The Lower the Combined Variance, the Higher the t Value

The illustrations below should clarify how the degree of variance in the two groups determines how easy or difficult it is to state that the means of the two groups are really different. The t test quantifies this relationship and provides a way to determine whether the measured difference between two means can be considered real or not based upon the amount of variance in both groups. Here are illustrations following that should clarify this relationship:



We can see that pair of data sets on the right are much easier to differentiate because they have much less overlap than the pair of data sets on the left. The overlap represents the overall variability between the two data sets in each pair. The higher the total variability within the pair of data sets, the higher will be the denominator in the t value formula. The higher the denominator, the lower the t value for the pair of data sets. The lower the t value, the less likely it is that the two data sets are separate data sets with different means.

T-Test for Two Samples Having Unequal Sizes and Variances

The t Test for comparing two samples with unequal sizes and variance is a variation of the t Test called Welch's t Test. It is not the classic Student's t Test, which does not allow for samples having unequal variances.

We are going to use this t test to determine within 95% certainty whether the means sales from Wednesdays is different than the mean sales from Monday. We have measured from the last 42 Wednesdays and the last 40 Mondays and we will apply this Excel t test to determine whether the measured difference between the means is real or not.

A Little Bit More About This t Test

The t Test in general is a special case of one-way (sometimes called “single factor”) ANOVA. This paired two-sample student's t test is applied when there is a natural pairing of samples. It is most often used to determine whether “before” and “after” means of a sample of the same objects have changed during an experiment. One really great thing about this t test is that the paired two-sample t test does not require that the variances of both populations to be the same.

Here is the formula to calculate the t value for a two-sample t test of unequal variances if you are testing to determine whether there difference between the two samples:

$$\mathbf{t\ value} = [X1 - X2] / [\text{SQRT}((s1^2 / n1) + (s2^2/n2))]$$

$$\mathbf{Degree\ of\ Freedom} = df =$$

$$[[(s1^2 / n1) + (s2^2 / n2)]^2] / [([(s1^2 / n1)^2] / [n1 - 1]) + ([(s2^2 / n2)^2] / [n2 - 1])]$$

X1 and X2 are the sample means. s1 and s2 are the sample standard deviations, and n1 and n2 are sample sizes.

You can see that this follows the general formula for calculating the t value in a t test, which is:

$$\mathbf{t\ value} = (\text{Difference between the group means}) / (\text{Variability of the groups})$$

The t value is a specific point on the x-axis in the t distribution

(student's t distribution). If this t value falls outside the region of required certainty, it can be stated that the two means are probably different. If this t value falls within the region of required certainty, it cannot be stated that the two means are probably different.

The required region of certainty depends upon the degree of certainty required in the test. If 95% certainty is required, then the required region of certainty consists of 95% of the area under the student's t distribution. The outer 5% is the region of uncertainty. This is also referred to as α (alpha) or the degree of significance. If the t value is large enough to be located all the way out on the x-axis in the 5% region of uncertainty, it can be stated within 95% certainty that the two means are different.

A t test can be a one-tailed test or a two-tailed test. A one-tailed test determines whether the means are different in one specific direction. For example, a one-tailed test could be used to determine only if the mean of the “after” measurements is greater than the mean of the “before” measurements. A two-tailed test determines whether the two means are merely different.

Two-Tailed t Test Is More Stringent

The two-tailed test is more stringent because the area in the outer tails outside of the region of required degree of certainty is split into two tails. For example, if the required degree of certainty is 95% on a two-tailed test, the calculated t value must be all the way out in the outer 2.5% of either tail for the t test to conclude within 95% certainty that the means are different.

One-Tailed t Test Is Less Stringent

A one-tailed test is less stringent. If the required degree of certainty is 95% on a one-tailed test, the calculated t value only has to be within the outer 5% of whatever tail is being tested to be able to state the two means are probably different.

Doing The Two-Sample t Test for Unequal Variances in Excel

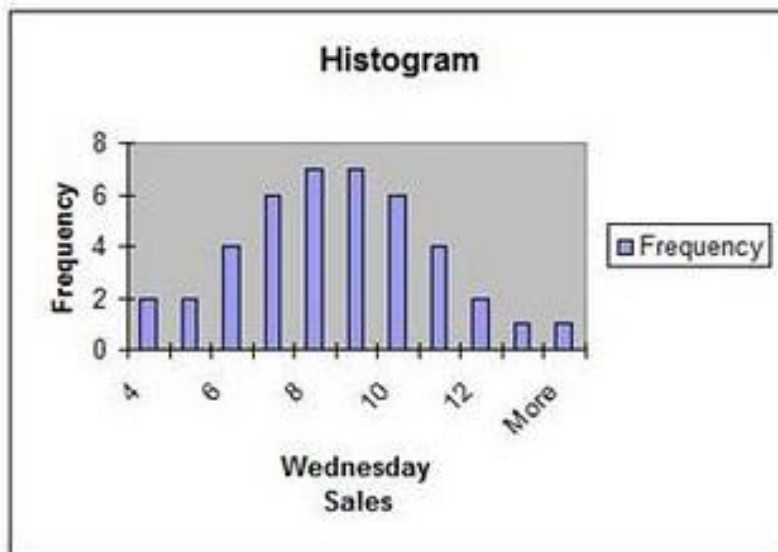
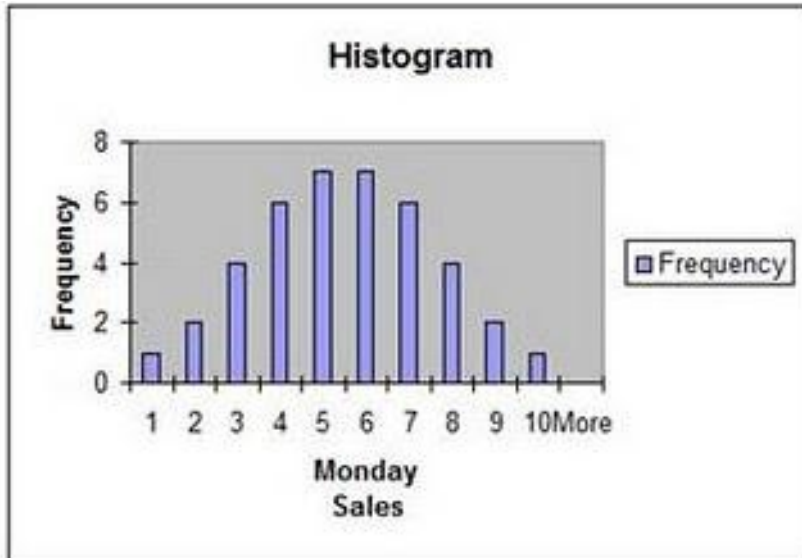
We are testing to determine whether there really is a difference between mean sales on Monday and mean sales on Wednesday.

The data need to be arranged in Excel as follows:

Monday Sales	Wednesday Sales
1	3
2	4
3	5
4	6
5	7
6	8
7	9
8	10
9	11
10	12
2	13
3	14
4	5
5	6
6	7
7	8
8	9
9	10
3	11
4	12
5	6
6	7
7	8
8	9
3	10
4	11
5	6
6	7
7	8
8	9
4	10
5	11
6	7
7	8
4	9
5	10
6	7
7	8
5	9
6	10
	8
	9

The t Test we are about to use allows for different sample sizes and different variances, but that standard requirement for all t Tests is that both samples being compared are Normally distributed. There are a number of different ways of doing this. For brevity, we are going to do it the simplest possible way. We will make an Excel histogram of each sample's data and simply eyeball the shape of the histogram. If the shape of each histogram resembles the Normal curve, we will go with it. The Excel histogram is a simple thing to construct.

Completed histograms for each of the two samples are as follows:



Both histograms appear to be Normally distributed so we can use t Test

to compare the two samples. If either sample is not Normally distributed, the t test cannot be used because the output is likely to be **totally** incorrect. If either sample is not Normally distributed, we must use a nonparametric test such as the Mann-Whitney U Test to compare the samples.

Before we run the t Test, we would like to take a look at a description of each sample. In Excel 2003, this can be quickly done by the following tool

Tools / Data Analysis / Descriptive Statistics. The Descriptive Statistics for each sample are as follows:

Monday Sales	
Mean	5.5 = X1
Standard Error	0.33
Median	5.5
Mode	5
Standard Deviation	2.09 = s1
Sample Variance	4.36
Kurtosis	-0.42
Skewness	0.00
Range	9
Minimum	1
Maximum	10
Sum	220
Count	40 = n1

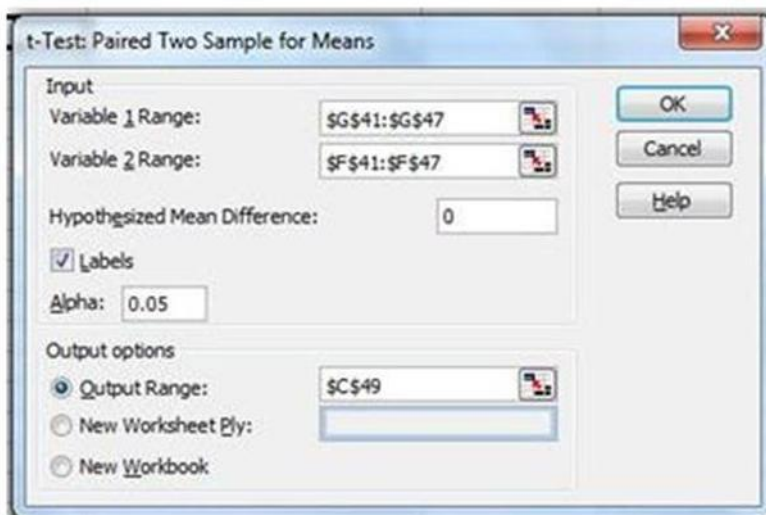
Wednesday Sales		
Mean	8.5	= X2
Standard Error	0.37	
Median	8.5	
Mode	8	
Standard Deviation	2.37	= s2
Sample Variance	5.62	
Kurtosis	0.01	
Skewness	0.00	
Range	11	
Minimum	3	
Maximum	14	
Sum	357	
Count	42	= n2

Now, access this Excel t Test as follows (this is Excel 2003):

Tools / Data Analysis / t-Test: Two Sample Assuming Unequal Variances

(Data Analysis Tools in Excel 2010 Are Found in the Data tab)

This following dialogue box will appear:



Input the data as followings:

Variable 1 Range: Select everything that is highlighted yellow, including the label “Monday Sales.”

Variable 2 Range: Select everything that is highlighted tan, including the label “Wednesday Sales.”

Hypothesized Mean Difference: 0

Labels: Check the box because you included the labels for Variables 1 and 2.

Alpha: This depends on your desired degree of certainty. 0.05, if you desired 95% certainty. 0.20 if you desire 80% certainty.

Output Range: Select the cell that you want the upper left corner of the output to appear in.

Hit “OK” to run the analysis and the following Excel output appears:

t-Test: Two-Sample Assuming Unequal Variances		
	<i>Monday Sales</i>	<i>Wednesday Sales</i>
Mean	5.5	8.5
Variance	4.36	5.62
Observations	40	42
Hypothesized Mean Difference	0	
df	80	
t Stat	-6.088	
P(T<=t) one-tail	1.88E-08	
t Critical one-tail	1.664	
P(T<=t) two-tail	3.76E-08	
t Critical two-tail	1.990	

This output can be interpreted as follows:

The t value is -6.088.

$\alpha = 0.05 = 1 - \text{Required Degree of Certainty} = 1 - 95\%$

p Value (1-Tailed) = 1.88E-08

p Value (2-Tailed) = 3.77E-08

One-tailed Test

This t value has a greater absolute value (6.088) than the critical t value for a one-tailed test (1.664). We can therefore state with 95% certainty that there really is a difference between Wednesday sales and Monday sales.

The above conclusion can also be reached because the p Value for the one-tailed test (highlighted in light red on the Excel output) is 1.88E-08. This is much less than alpha (0.05). The p Value being less than alpha is an equivalent result to the t value being greater than the t critical value.

Two-Tailed Test

The same result is arrived at for the two-tailed test. The two-tailed test is more stringent because the alpha region of uncertainty (5% of the area under the student's t distribution curve) is now divided between both outer tails. The t value needs to be larger for the two-tailed test to wind up in the outer 2.5% area of either outer tail.

In this case, the t value was large enough to be positioned in the outer 2.5% of either outer tail. The absolute value of t value (6.088) is much larger than the critical t value for the two-tailed test (1.990). This indicates that it can be stated with 95% certainty that there really is a difference between Wednesday sales and Monday sales in this case as well. The p value calculated for the two-tailed test (3.77E-08) is much smaller than alpha (0.05). This is an equivalent result to the above.

Hand Calculation of the t Value and p Value

Let's calculate the t value and p values for the one and two-tailed tests by hand to make sure that Excel has done a correct job. The t value is stated as the t statistic.

Here is the original test data Excel Descriptive Statistics:

Monday Sales	
Mean	5.5 = X1
Standard Error	0.33
Median	5.5
Mode	5
Standard Deviation	2.09 = s1
Sample Variance	4.36
Kurtosis	-0.42
Skewness	0.00
Range	9
Minimum	1
Maximum	10
Sum	220
Count	40 = n1

Wednesday Sales	
Mean	8.5 = X2
Standard Error	0.37
Median	8.5
Mode	8
Standard Deviation	2.37 = s2
Sample Variance	5.62
Kurtosis	0.01
Skewness	0.00
Range	11
Minimum	3
Maximum	14
Sum	357
Count	42 = n2

Here is the hand calculation of the t value and p values for the one and two-tailed tests for this Two-Sample t Test Assuming Unequal Variance. The hand calculations below of the t Value and p Values agree with the Excel outputs. There are very slight differences due to rounding differences:

$$\mathbf{t\ value} = [X1 - X2] / [\text{SQRT}((s1^2 / n1) + (s2^2/n2))]$$

$$\mathbf{Degree\ of\ Freedom} = df =$$

$$[[(s1^2 / n1) + (s2^2 / n2)]^2] / [([(s1^2 / n1)^2] / [n1 - 1]) + ([(s2^2 / n2)^2] / [n2 - 1])]$$

The Degrees of Freedom calculation must be rounded to the nearest whole number, which in this case is 80.

X1 and X2 are the sample means. s1 and s2 are the sample standard deviations, and n1 and n2 are sample sizes.

$$p\ \text{Value} = \text{TDIST} (T\ \text{Statistic}, df, \text{Number of Tails})$$

Here are the actual calculations done by hand in Excel on the next page:

$$df = \frac{[(s1^2 / n1) + (s2^2 / n2)]^2}{\left[\frac{(s1^2 / n1)^2}{n1 - 1} + \frac{(s2^2 / n2)^2}{n2 - 1} \right]}$$

$$df = 80$$

$$t = \frac{X1 - X2}{\sqrt{(s1^2 / n1) + (s2^2 / n2)}}$$

$$t = -6.088$$

$$p \text{ value (1 tail)} = \text{TDIST (T Statistic, df, \# tails)} = \text{TDIST (-6.088, 80, 1)}$$

$$p \text{ value - 1 tail} = 1.87843E-08$$

$$p \text{ value (2 tail)} = \text{TDIST (T Statistic, df, \# tails)} = \text{TDIST (-6.088, 80, 2)}$$

$$p \text{ value - 2 tail} = 3.75687E-08$$

The Two Sample t Test Assuming Unequal Variances is a very simple test to run in Excel and can be applied to nearly any aspect of your marketing program to see if one group of samples is different from another group of samples. One note: both sample groups must be continuous and measured using the using the same scale.

Correctable Reasons Why Your Data Is Not Normally Distributed

In the ideal world, all of your data samples are normally distributed. In this case you can usually apply the well-known parametric statistical tests such as ANOVA, the t Test, and regression to the sampled data.

What can you do if your data does not appear to be normally distributed?

You can either:

- Apply nonparametric tests to the data. Nonparametric tests do not rely on the underlying data to have any specific distribution
- Evaluate whether your “non-normal” data was really normally-distributed before it was affected by one of the seven correctable causes listed as follows:

The Biggest 7 Correctable Causes of Non-Normality in Data Samples

1) Outliers – Too many outliers can easily skew normally-distributed data. If you can identify and remove outliers that are caused by error in measurement or data entry, you might be able to obtain normally-distributed data from your skewed data set. Outliers should only be removed if a specific cause of their extreme value is identified. The nature of the normal distribution is that some outliers will occur. Outliers should be examined carefully if there are more than would be expected.

2) Data has been affected by more than one process

– It is very important to understand all of the factors that can affect data sample measurement. Variations to process inputs might skew what would otherwise be normally-distributed output data. Input variation might be caused by factors such as shift changes, operator changes, or frequent changes in the underlying process. A common symptom that the output is being affected by more than one process is the occurrence of more than one mode (most commonly occurring value) in the output. In such a situation, you must isolate each input variation that is affecting the output. You must then isolate the overall effect which that variation had on the output. Finally, you must remove that input variation's effect from output measurement. You may find that you now have normally-distributed data.

3) Not enough data – A normal process will not look normal at all until enough samples have been collected. It is often stated that 30 is the where a “large” sample starts. If you have collected 50 or fewer samples and do not have a normally-distributed sample, collect at least 100 samples before re-evaluating the normality of the population from which the samples are drawn.

4) Measuring devices that have poor resolution

– Devices with poor resolution may round off incorrectly or make continuous data appear discrete. You can, of course, use a more accurate measuring device. A simpler solution is to use a much larger sample size to smooth out sharp edges.

5) A different distribution describes the data

– Some forms of data inherently follow different distributions. For example, radioactive decay is described by the exponential distribution. The Poisson distribution describes events that tend to occur at predictable intervals over time, such as calls over a switchboard, number of defects, or demand for services. The lengths of time between occurrences of Poisson-distributed processes are described by the exponential distribution. The uniform distribution describes events that

have an equal probability of occurring. Application of the Gamma distribution often based on intervals between Poisson-distributed events, such as queuing models and the flow of items through a manufacturing process. The Beta distribution is often used for modeling planning and control systems such as PERT and CPM. The Weibull distribution is used extensively to model time between failure of manufactured items, finance, and climatology. It is important to become familiar with the applications of other distributions. If you know that the data is described by a different distribution than the normal distribution, you will have to apply the techniques of that distribution or use nonparametric analysis techniques.

6) Data approaching zero or a natural limit –

If the data has a large number of values that are near zero or a natural limit, the data may appear to be skewed. In this case, you may have to adjust all data by adding a specific value to all data being analyzed. You need to make sure that all data being analyzed is “raised” to the same extent.

7) Only a subset of process’ output is being analyzed –

If you are sampling only a specific subset of the total output of a process, you are likely not collecting a representative sample from the process and therefore will not have normally distributed samples. For example, if you are evaluating manufacturing samples that occur between 4 and 6AM and not an entire shift, you might not obtain the normally-distributed sample that a whole shift would provide. It is important to ensure that your sample is representative of an entire process.

If you are unable to obtain a normally-distributed data sample, you can usually apply non-parametric tests to the data.



Meet the Author

Mark Harmon is a master number cruncher. Creating overloaded Excel spreadsheets loaded with complicated statistical analysis is his idea of a good time. His profession as an Internet marketing manager provides him with the opportunity and the need to perform plenty of meaningful statistical analysis at his job.

Mark Harmon is also a natural teacher. As an adjunct professor, he spent five years teaching more than thirty semester-long courses in marketing and finance at the Anglo-American College in Prague and the International University in Vienna, Austria. During that five-year time period, he also worked as an independent marketing consultant in the Czech Republic and performed long-term assignments for more than one hundred clients. His years of teaching and consulting have honed his ability to present difficult subject matter in an easy-to-understand way.

Harmon received a degree in electrical engineering from Villanova University and MBA in marketing from the Wharton School.

