



Published on *Plant Breeding E-Learning in Africa* (<https://pbea.agron.iastate.edu>)

[Home](#) > [Course Materials](#) > [Quantitative Methods](#) > Continuous Data

Continuous Data



By Ron Mowers, Ken Moore, M. L. Harbur, Kendra Meade, William Beavis, Laura Merrick (ISU)



Except otherwise noted, this work is licensed under a [Creative Commons Attribution-NonCommercial 4.0 International License](https://creativecommons.org/licenses/by-nc/4.0/).

Introduction

In this module we will learn that sample means from a population that has a Normal distribution are also distributed as Normal, but with a smaller variance than the population. We will then extend the concept to samples from two populations. For example, sometimes we want to identify the best management practice between two alternatives. Such as comparing the use of starter fertilizer to no starter fertilizer, or one corn hybrid against another, or a new insecticide against a common one. In all of these comparisons, the key question to be answered is "are the means of the populations representing the two treatments the same." In the language of statistics this is called a two-sample hypothesis; the most common method for testing it is the t-test.

Objectives

- How the distribution of a set of sample means around their population mean can be described in terms of t-values
- How to calculate the t-value for sample means as related to their population means
- How a confidence limit describes a range of possible values for the population mean, μ
- How a t-test can be used to determine whether the difference between two sample means is significant
- How to conduct t-tests using Excel

The t -Distribution

When the Variance is Unknown

THE t -DISTRIBUTION IS USED FOR SAMPLE MEAN WHEN THE VARIANCE IS NOT KNOWN.

Statistical standards must apply to a wide variety of research, from corn yield trials to soybean disease studies to nitrate sampling in groundwater. Even within one of these categories, the amount of variability may differ from one month or year to the next. Therefore, we need to describe the distribution of our sample means according to the distribution curve itself, and in a way that is unitless. We cannot use the normal distribution described in the module on Distributions and Probability because we do not know the variance. We define this distribution of sample means in terms of the t -distribution.

As we learned, the shape of the distribution curve for sample means is based on the number of observations per sample. Our values for t will also vary with the number of observations or replications, that occur. However, for normally distributed variables and a given number of **degrees of freedom** (df), we know that 95% of the sample means will lie within a particular t -value of the mean. This t -value remains the same, for the given df, whether we are studying swine, alfalfa, or soil moisture.

What are **degrees of freedom**? In general, the number of degrees of freedom associated with a set of sample means is the number of individuals in the sample used to calculate the mean minus 1.

A Sample *t*-Value Table

For example, look at the following *t*-value table. For a sample size of 5, 95% of the sample means will be within plus or minus 2.776 standard errors (S/\sqrt{n}) of the mean:

Table 1 *t*-values: Probability of obtaining a value as large or larger

Degrees of Freedom	Percentage Points in Top Tail				
	5	2.5	1	0.5	1
1	6.314	12.706	31.821	63.657	318.309
2	2.920	4.303	6.965	9.925	22.327
3	2.353	3.182	4.541	5.841	10.215
4	2.132	2.776	3.747	4.604	7.173
5	2.015	2.571	3.365	4.032	5.893
6	1.943	2.447	3.143	3.707	5.208
7	1.895	2.365	2.998	3.499	4.785
8	1.860	2.306	2.896	3.355	4.501
9	1.833	2.262	2.821	3.250	4.297
10	1.812	2.228	2.764	3.169	4.144
11	1.796	2.201	2.718	3.106	4.025
12	1.782	2.179	2.681	3.055	3.930
13	1.771	2.160	2.650	3.012	3.852
14	1.761	2.145	2.624	2.977	3.787
15	1.753	2.131	2.602	2.947	3.733

In this example, a sample with four degrees of freedom will have 5% of the distribution $2.776 \times t$ above the mean and another 5% of the distribution $2.776 \times t$ below the mean.

Study Questions 1

Fill in the missing words

For a sample size of 10, within how many standard errors would 99% of the sample means be contained?

 Check

Sample Means in Values of t

Now that we understand how t values relate to the distribution curve for a set of sample means, we must learn how to express our sample means in values of t . This is given by the following formula:


$$t = \frac{\bar{Y} - \mu}{S_{\bar{y}}}$$

Equation 1

where

\bar{Y} = sample mean

μ = population mean

 **Invalid Equation** = standard error of the sample mean = $\frac{S}{\sqrt{n}}$

We use the t -value rather than z because we are estimating σ^2 with s^2 . The t -value represents the difference between a particular sample mean and the mean of the whole population of sample means, as a function of the standard error for that population, $S_{\bar{y}} = S/\sqrt{n}$. Stated more simply, the t -value is the number of standard errors a particular sample mean is from the average value for all of the sample means. This tells us whether our difference is large in relation to the variation of our sample.

t-Value Scenarios

For example, a 500 kg/ha difference between the sample mean of a certain corn hybrid and the population mean of all corn hybrids would not be a large difference when we are dealing with corn (which might have a standard error of 750 kg/ha for a given number of plots) as when we are studying soybean (which could have a standard error of 250 kg/ha for the same number of plots). The *t*-values for each of these scenarios are given below:

Corn:

$$t = \frac{(\bar{Y} - \mu)}{\bar{S}_y} = \frac{250\text{kg}}{750\text{kg}} = 0.34$$

Equation 2

Soybean:

$$t = \frac{(\bar{Y} - \mu)}{\bar{S}_y} = \frac{500\text{kg}}{200\text{kg}} = 2$$

Equation 3

Study Questions 2 and 3

Confidence Limits

Confidence Intervals

Confidence intervals (CI) are used to describe a range of values for a set of sample means from multiple samples of the same population. A CI does not predict that the true population mean will be included in the interval.

We can set the ranges for our [confidence interval](#) using Excel. You will learn how to do this in the Excel lab for this module.

Try This! Calculation Exercise

We use the term confidence to remind us that the variation associated with this procedure is the result of our own experimental and sampling techniques. Stated otherwise, the population mean in which we are interested stays the same -- it is our confidence interval that shifts according to the number of and particular samples in our set. There is no probability or chance associated with the population mean, but rather with the value of the sample means.

Ex. 1: Calculating a Confidence Interval

Often we are asked to make decisions about some characteristic of a population based on samples selected from it. In this example, we will consider the protein value of two lots of alfalfa hay.

The data in the Alfalfa Quality worksheet were collected from two lots of hay that were offered for sale at auction. Each lot was first-cutting hay and weighed approximately 30 tons. Based upon a one percentage unit difference in the protein analysis, lot 2 was discounted \$10 per ton because the buyer claimed it was less than 15.5% protein. The buyer's method to see if a lot is less than 15.5% is to calculate a 95% confidence interval for protein and see if it includes 15.5%.

Our first question should be "how confident are we that the sample means represent the true mean." One approach to answering this question is to calculate a confidence interval for the mean of each sample population.

Exercise 1

- Calculate the 95% confidence interval of the sample mean for each lot. Download and open the Excel file [QM-mod6-ex1data.xls](#).
- There should be three columns: Lot, Rep, and Crude Protein.
- We will calculate a mean, standard deviation and standard error for each lot. Use the Average and STDEV.S to calculate the mean and standard deviation of each lot sample. Then use the standard deviation to calculate the standard error by dividing the standard deviation by the square root of the number of samples in a lot.

It is possible to use Excel to calculate a confidence interval from a t-distribution using the formula Confidence.T. The arguments are the alpha level, standard deviation, and the number of samples or replications.

- Now subtract and add the calculated value from the mean to find the confidence interval.
- Now, we will hand calculate the confidence interval to see that it matches that given by Excel. Use Equation 4 and the SE from Excel. Use the t-value for 5% probability in two tails (2.5% in one tail) from either Appendix 3 or the t-table given earlier in this unit. Remember, there are 7 df for 8 reps in each sample.
- Calculate the 95% confidence interval for each lot.

	Lot	Rep	Crude Protein	Average	Std Dev	Std Err
1						
2	1	1	16.08	15.57	0.481901	0.170378
3	1	2	14.77			
4	1	3	15.03			
5	1	4	15.95			
6	1	5	15.58			
7	1	6	15.39			
8	1	7	15.70			
9	1	8	16.06			
10	2	1	12.69	14.5025	1.369731	0.484273
11	2	2	15.86			
12	2	3	13.12			
13	2	4	16.31			
14	2	5	15.64			
15	2	6	14.62			
16	2	7	13.27			
17	2	8	14.51			
18						

Study Questions 4

Do the confidence intervals overlap?

No

Yes



Check

Discussion

Should Lot 2 have been discounted for a small difference? Support your assertions.

t-Tests For Significance

Comparing Two Treatment Means

t-tests can be used to compare two treatment means.

Confidence intervals may work well for single samples compared to the mean. But what can we use to compare two samples? We can also use *t*-values to determine whether the calculated mean values of two sample sets are significantly different. This difference can be expressed in terms of *t*:

$$t = \frac{\bar{d} - \mu_{\bar{d}}}{S_{\bar{d}}}$$

where \bar{d} = difference between sample means

$\mu_{\bar{d}}$ = difference between population means

$S_{\bar{d}}$ = standard error of the difference

Equation 4

In most cases, we are trying to identify any significant difference; the value we use for the mean difference is zero. A zero value for a mean difference therefore represents the null hypothesis that there is no treatment effect. Our *t*-value expresses the actual difference between the two sample means as a function of standard error of the difference. Again, it is this conversion that makes the difference meaningful regardless of what particular variable we are measuring.

Paired t-Test

It is at this point that we must know more about the experiment or how the populations were sampled in order to compute S_d . If both treatments of an experiment are applied to units that are paired in blocks, we use a paired t -test method. If the populations are independent or treatments are applied completely at random to the experimental units, we use what is called an independent samples t -test. Each of these has a different standard error of the difference S_d .


Paired t -test

Suppose first that the treatments are on units that are naturally paired (and thus not independent). An example is a test for difference in two wheat varieties each grown in the same block or area in a field for 5 different blocks. There is a natural pairing of the varieties because they each appear in each block even though they have been randomly assigned to either the left or right side of the block. Another example is a test of side dressing of nitrogen fertilizer vs the check of no additional fertilizer, each applied to half a field on fields from 12 different farms.

For the paired t -test, the computation of the t -value is easy. Just take differences of treatment 1 - treatment 2 for each pair. Then, use these differences as the data and compute the t value as done in equation 1. The d in equation 5 is just the average of the differences and the S_d is the square root of S_d^2/n . The variance of the differences is S_d^2 and number of pairs is n . The test of the null hypothesis that treatments have the same mean is just a test of differences being zero, $H_0 : \mu_d = 0$. We will see in the Excel exercises how this is done.

Independent Samples *t*-Test

The independent samples *t*-test, done for independent populations or samples, uses the $S_{\bar{d}}$ given in Equation 5 below. One independent sample example is a set of 24 pots of soybeans in a growth chamber, half of which receive a zinc foliar supplement, while the other half receive no supplement. The pots which receive the treatment are randomly chosen, for example, pots 2, 3, 7, 9, 10, 11, 13, 14, 16, 18, 21, 24. Sometimes samples are considered to be independent when they are drawn at random from two mutually-exclusive populations, for example a poll on farmers' opinions on a land-use policy, with 100 respondents in Illinois having average farm size over 1000 acres and another 100 with farms less than 1000 acres. It is important that these samples be chosen at random for a valid comparison.

When populations are independent, as in the 2-sample case, the variance of a difference is the sum of the two variances. In the 2-sample (independent samples) case, the variance of a difference in sample means $(\bar{x}_1 + \bar{x}_2)$ is the sum of their variances  Invalid Equation, where subscripts indicate the population. In equation 5, the standard error used to divide the mean difference is the square root of the sum of the variances associated with x_1 and x_2 :

$$S_{\bar{d}} = \sqrt{S_1^2 + S_2^2}$$

where $S_{\bar{d}}$ = standard error of the difference between sample means

S_1^2 = variance from sample population x_1

S_2^2 = variance from sample population x_2

Equation 5

Calculation Results

In many texts, the standard error of the difference is calculated as:

$$S_{\bar{d}} = \sqrt{\frac{S_1^2}{n_1} + \frac{S_2^2}{n_2}}$$

Equation 6 Standard error of the difference

This is because the variance of population 1 is calculated as:

$$S_1^2 = \frac{S_1^2}{n}$$

Equation 7

and likewise for population 2.

Since both samples have variability involved, both variables must be accounted for. The number of degrees of freedom associated with this difference is either calculated by the computer, or if done by hand, is the degrees of freedom associated with the smaller sample. For samples of equal size, and if we can assume $S_1^2 = S_2^2$, the number of degrees of freedom is $2(n-1)$.

After calculating the t -value of the difference in which we are interested, we then look up the critical table t -value, for which there is only a 5% chance of a larger value occurring. Generally we have critical t -values of about 2, or slightly higher, for a 0.05-level two-tail test. If the t -value of our difference exceeds the critical t -value, then we say that we have a significant difference between the two means. Would a significant t -test indicate that the two populations are truly different?

Try This!: t-Test Exercises

There are three classes of t-tests that are commonly used in agronomic experiments (which can be calculated using Excel):

Ex. 2: t-Test Assuming Equal Variances

t-test assuming equal variances

There is a more direct approach to comparing the means of two sample populations - the *t*-test. In our alfalfa example, suppose the populations have equal variances and what we really want to know is whether or not the two sample means are different from one another. This question can be stated as a simple statistical hypothesis: $H_0: \mu_1 = \mu_2$ vs. the alternative $H_a: \mu_1 \neq \mu_2$. In this exercise we will use the *t*-test procedure in Excel to test this hypothesis.

Exercise 2

- Perform a *t*-test to compare mean protein concentrations of the two alfalfa lots using the following steps:
Open the file created in [QM-mod6-ex1data.xls](#).
- Label a new column 't-test: Equal variance' and in the first row under the title enter the formula "`=T.TEST(C2:C9,C10:C17,2,2)`".

This dictates that the observations for the first sample are in C2:C9, the second set are in C10:C17, it is a two-tailed test, and that this is an independent sample *t*-test assuming equal variances.

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O
	Lot	Rep	Crude Protein		Average	Std Dev	Std Err	Confidence Interval Calculation	Lower Limit of CI	Upper Limit of CI	t-test: Equal Variances (Pr > t)				
1															
2	1	1	16.08		15.57	0.481901	0.170378	0.402879316	15.17	15.97	0.056441163				
3	1	2	14.77												
4	1	3	15.03												
5	1	4	15.95												
6	1	5	15.58												
7	1	6	15.39												
8	1	7	15.70												
9	1	8	16.06												
10	2	1	12.69		14.5025	1.369731	0.484273	1.145124162	13.36	15.65					
11	2	2	15.86												
12	2	3	13.12												
13	2	4	16.31												
14	2	5	15.64												
15	2	6	14.62												
16	2	7	13.27												
17	2	8	14.51												

- Read the results of the *t*-test (assuming equal variances) as Prob > |t|. This is the probability, assuming the null hypothesis is true, of observing the results in your data set, and is 0.0564.
- Note that the *t*-test is assuming equal variances and has a two-tailed alternative. To convert the

probability to a one-tailed alternative, divide the two-tailed probability by 2. It is $0.0564/2 = 0.0282$.

- Do the means differ? (Is the Lot 1 mean significantly lower than the Lot 2 mean based on these samples?)

Ex. 3: *t*-Test Assuming Unequal Variances

You may have noticed in the *t*-test output from the last exercise that the standard error for Lot 2 is about three times as great as that for Lot 1.

One of the assumptions we made for the *t*-test was that the two sample populations shared a common variance. Hence, a pooled estimate of the population variance was used in the test.

We can easily test the hypothesis that the two variances are equal by computing a simple F test:

- Open the [QM-mod6-ex1data.xls](#) file.
- Calculate the variance for each lot in a new column using the formula Var.S.
- Divide the larger of the two variances by the smaller one to calculate F.
- Determine the critical F value. (Numerator and denominator df are those from the two samples.) The alpha level is 0.05.
- If the calculated F is greater than the critical F, we conclude that the variances are different and use another *t*-test.
- If the calculated F is less than the critical F, we conclude that the variances are the same and that a pooled variance was appropriate for the *t*-test.

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O
	Lot	Rep	Crude Protein	Average	Std Dev	Std Err	Confidence Interval Calculation	Lower Limit of CI	Upper Limit of CI	t-test: Equal Variances (Pr > t)	Variance	F-value (ratio of variances)	Pr > F (probability of a value of F or greater by chance)	t-test: Unequal Variances	
1															
2	1	1	16.08	15.57	0.481901	0.170378	0.402879316	15.17	15.97	0.056441163	0.23	8.078955	0.003911	0.068367	
3	1	2	14.77												
4	1	3	15.03												
5	1	4	15.95												
6	1	5	15.58												
7	1	6	15.39												
8	1	7	15.70												
9	1	8	16.06												
10	2	1	12.69	14.5025	1.369731	0.484273	1.145124162	13.36	15.65		1.88				
11	2	2	15.86												
12	2	3	13.12												
13	2	4	16.31												
14	2	5	15.64												
15	2	6	14.62												
16	2	7	13.27												
17	2	8	14.51												
18															

For our example the calculated F is $(1.88/.23) = 8.17$ and the critical F = 3.79. Therefore, we conclude that the variances of the two sample populations are different and a different *t*-test is required.

- It is possible to find a p-value for this test using "`=F.DIST.RT(M2,8,8)`". M2 is the cell with the F-ratio in it.

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O
	Lot	Rep	Crude Protein	Average	Std Dev	Std Err	Confidence Interval Calculation	Lower Limit of CI	Upper Limit of CI	t-test: Equal Variances (Pr > t)	Variance	F-value (ratio of variances)	Pr > F (probability of a value of F or greater by chance)	t-test: Unequal Variances	
1															
2	1	1	16.08	15.57	0.481901	0.170378	0.402879316	15.17	15.97	0.056441163	0.23	8.078955	0.003911	0.068367	
3	1	2	14.77												
4	1	3	15.03												
5	1	4	15.95												
6	1	5	15.58												
7	1	6	15.39												
8	1	7	15.70												
9	1	8	16.06												
10	2	1	12.69	14.5025	1.369731	0.484273	1.145124162	13.36	15.65		1.88				
11	2	2	15.86												
12	2	3	13.12												
13	2	4	16.31												
14	2	5	15.64												
15	2	6	14.62												
16	2	7	13.27												
17	2	8	14.51												
18															

Note, however, that this test and other tests for equality of variance are not robust procedures. We saw earlier that the *t*-test is reasonably robust. It returns fairly accurate probability statements for samples from similarly shaped distributions when sample sizes are nearly equal. The F-test for variances, however, is very sensitive to non-normal distributions.

Exercise 3

Perform a t-test to compare mean protein concentrations of the two alfalfa lots. This is the same hypothesis as Exercise 2: $H_0: \mu_1 = \mu_2$ vs. the alternative $H_a: \mu_1 \neq \mu_2$ Use the following steps:

- Label a new column 't-test: Unequal Variances'.
- This time, we again look for the $P > |t|$ for the test of the null hypothesis that the means for the two lots are equal.
- Use the same formula from Exercise 2, but dictate that the required test assumes unequal variances.
- For the test with unequal variance, $P = 0.0684$. This is again for the two-tailed alternative, by default.

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O
	Lot	Rep	Crude Protein		Average	Std Dev	Std Err	Confidence Interval Calculation	Lower Limit of CI	Upper Limit of CI	t-test: Equal Variances	Variance	Bartlett's Test	F-test	t-test: Unequal Variances
1															
2	1	1	16.08		15.57	0.481901	0.170378	0.402879316	15.17	15.97	0.056441163	0.23	8.078955	0.003911	0.068367
3	1	2	14.77												
4	1	3	15.03												
5	1	4	15.95												
6	1	5	15.58												
7	1	6	15.39												
8	1	7	15.70												
9	1	8	16.06												
10	2	1	12.69		14.5025	1.369731	0.484273	1.145124162	13.36	15.65		1.88			
11	2	2	15.86												
12	2	3	13.12												
13	2	4	16.31												
14	2	5	15.64												
15	2	6	14.62												
16	2	7	13.27												
17	2	8	14.51												
18															
19															
20															

Ex. 4: Paired t-test

Sometimes data collected from two sample populations are related in some way. When this occurs we say that the observations are paired. A good example of this is the strip design that is often used in on-farm trials to compare two treatments. Treatments are applied to strips that run across a field in such a way that every adjacent pair of strips contains both treatments. Treatments are randomly allotted to strips within each pair. For a more complete description of this methodology read the paper by **Exner and Thompson**.

Exercise 4

- Perform a paired t-test to determine if starter fertilizer had an effect on corn yield using the following steps:
- Download and open the Excel file [QM-mod6-ex4data.xls](#).
- We will first calculate the paired t-test step-by-step.
- Calculate the difference between each pair of observations.
- Calculate the mean of the differences.
- Now calculate the standard deviation, using Stdev.S.
- The standard error is the standard deviation divided by the square root of the number of pairs.
- The t-statistic is the mean divided by the standard error.
- The calculated value is less than the tabular value, so we fail to reject the null. There is not a significant difference between the two treatments.

Study Questions 5 and 6

Two-Sample Hypothesis Testing

There are three sets of hypotheses that can be tested with a t-test (Table 2).

Table 2 Hypothesis tests

Set	Null Hypothesis	Alternative Hypothesis
1	$\mu_1 = \mu_2$	$\mu_1 \neq \mu_2$
2	$\mu_1 \geq \mu_2$	$\mu_1 < \mu_2$
3	$\mu_1 \leq \mu_2$	$\mu_1 > \mu_2$

The first set simply states that there is either a difference between the two means or there is not. The alternative hypothesis gives no regard to how the two means may differ, but simply states that they are different in some way. We use this hypothesis when we do not have a preconceived idea as to how the means will differ. For example, if we are comparing yields of two highly recommended varieties, we probably have no expectation which will be greater. This type of comparison is called a two-tailed test because we want to know if the observed t lies within either tail of the distribution. Assuming a mean difference for the two varieties of 250 kg/ha, a standard error of 120, and 8 df the test would look like this:

$$t = \frac{\bar{d}}{s_{\bar{d}}} = \frac{250}{120} = 2.083 < t_{.05, 8df} = 2.306$$

Equation 8

Test Conclusion

In this example, since our calculated t is less than the tabular t , we fail to reject the null hypothesis and conclude that the means are not significantly different at alpha 0.05.

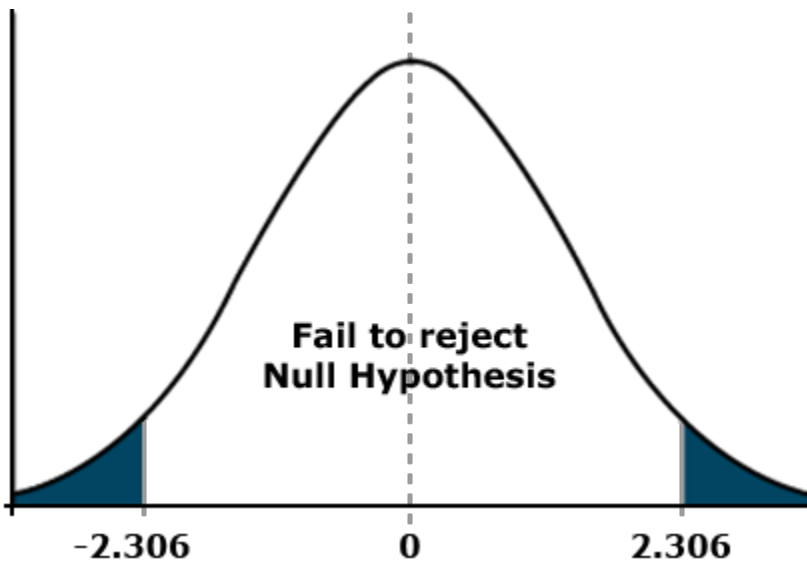


Fig. 1 2-tailed t-test acceptance region.

Study Questions 7

Using the given data but with a 300 kg/ha difference, what decision do we make about the samples?

Fail to reject the null hypothesis.

Reject the null hypothesis and accept the alternative hypothesis.

Check

Direction of the Mean Difference

With the other two sets of hypotheses we consider the nature or direction of the mean difference. For the alternative hypothesis we want to know if the mean for population 1 is significantly smaller (set 2) or larger (set 3) than the mean of population 2. We use one of these hypotheses when we have an expectation about how the means will differ. For example, if we are comparing the yield of a new variety with that of an older control variety, we expect the new one to out perform the old one. This type of comparison is called a one-tailed test because we want to know if the observed t lies within one tail of the distribution. Assuming a mean difference for the two varieties of 250 kg/ha, a standard error of 120, and 8 df, the one-tailed test would look like this:

$$t = \frac{\bar{d}}{s_{\bar{d}}} = \frac{250}{120} = 2.083 < t_{.05,8df} = 1.860$$

Equation 9

Calculate t Results

In this case the calculated t exceeds the tabular t and we conclude that the mean difference of 5 bu/acre is significant. The tabular t for a one-tailed test can be obtained from Appendix 3 in your text. However, the probability values listed in Table 1 are for a two-tailed test. To convert these values to probabilities for a one-tail test, divide them by 2. For our example above, you would use the t -values listed under 0.100 because 0.100 divided by 2 equals 0.05. The reason t -values differ for one and two-tailed tests has to do with the shaded areas of the curve. In the two-tailed test above, each shaded area represents 2.5% of the total area under the curve. For the one-tailed test, the shaded area under the right tail represents 5% of the total area.

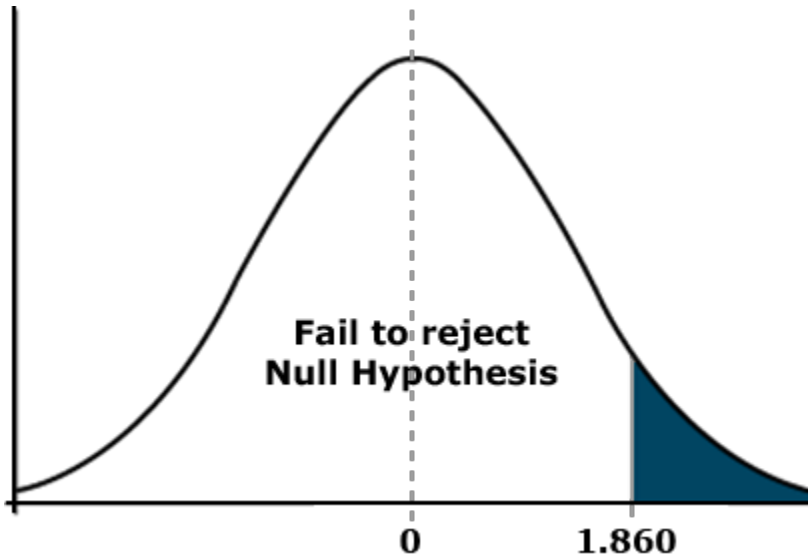


Fig. 2 One-tailed t -test acceptance region.

Linear Additive Model

Underlying the two sample experiment is a linear additive model. Linear models are mathematical representations of the variables contributing to the value of an observation. The observed response is partitioned into effects due to treatments and the experimental units. That is, each measured response is an additive function of the treatment effect and a random effect error associated with the experimental unit on which the measurement was made.

The linear model for the t -test is:

$$Y_{ij} = \mu + T_i + \epsilon_{(i)j}$$

where Y_{ij} = response observed for the ij^{th} experimental unit

μ = overall population mean

T_i = effect of the i^{th} treatment

$\epsilon_{(i)j}$ = effect associated with the ij^{th} experimental unit; commonly referred to as error

Equation 10

Assumptions for Linear Models

Now let's look at an example to see if we can make better sense of all this. A sales representative puts out a demonstration plot to compare a new herbicide product with an older one. Being a clever person, she has replicated each herbicide treatment four times. She therefore has eight plots and will have eight observations for each variable measured. The linear additive model (Fig. 4) for plot 3 of this experiment says the observed response in the plot equals the overall mean response for the population plus an effect due to the herbicide applied plus a random effect due to the plot itself. This latter effect acknowledges that there will be some differences in the measured response among plots receiving the same herbicide treatment.

$$\begin{aligned} \text{Number of weeds in plot 3} &= \text{Average of all plots} + \\ &\quad \text{Effect of Herbicide 2} + \\ &\quad \text{Plot 3 error} \end{aligned}$$

Equation 11 Linear additive response for herbicide response in an individual plot.

Essential assumptions for linear models:

- Additivity - the effect of each term is additive.
- Treatment effects are constant for all experimental units (plots).
- The effects of plots are independent and not related to treatment effects.

The linear additive model for an experiment determines how the data are to be analyzed. We will revisit this topic in later lessons as we discuss the analysis of variance.

Confidence Limits

Least Significant Difference

The least significant difference (LSD) allows an easy test for the difference in two treatment means.

Earlier, we established confidence limits for our sample mean. The same procedure can be applied to a difference between two means, using the t-value of the difference between the two means. The confidence interval for the difference between two means is:

$$CL = \bar{d} \pm ts_{\bar{d}}$$

where \bar{d} = difference between sample means

$s_{\bar{d}}$ = standard error of the difference

Equation 12

Confidence intervals (CI) for the true mean difference μ_d are obtained by computing the difference in means, d , and adding and subtracting t standard errors of d . To test the null hypothesis that μ_d is zero, we just see if zero is within the CL, and fail to reject the null if it is. This actually gives us an easy way to test the null hypothesis of zero difference: subtract the means and see if the difference is larger than ts_d .

The least significant difference (LSD) between two sample means is simply this difference:

$$LSD = ts_{\bar{d}}$$

where $s_{\bar{d}}$ = standard error of the difference

t = t-value for appropriate df and probability

Equation 13

Comparing Several Means

But what if we are comparing not one or two means, but several? For example, what if we were comparing the grain yield of five corn varieties. We would then have to calculate:

Table 3 Potential comparisons between five corn varieties.

Variety 1 vs. Variety 2	Variety 1 vs. Variety 3	Variety 1 vs. Variety 4	Variety 1 vs. Variety 5
	Variety 2 vs. Variety 3	Variety 2 vs. Variety 4	Variety 2 vs. Variety 5
		Variety 3 vs. Variety 4	Variety 3 vs. Variety 5
			Variety 4 vs. Variety 5

That's ten difference comparisons, from a relatively simple experiment! Wouldn't it be nice if we could know whether there were any significant differences to be found before we conducted this series of comparisons? As we will learn later, there is such a way.

Summary

T-Distribution

Used for sample means when variance is not known, but estimated.

Depends on degrees of freedom (df).

As df gets large, t -distribution becomes same as normal.

T-Values

Like z - value, but used when we estimate variance.

Tells how many standard errors from the mean.

Confidence Intervals (CI)

Give range for the population mean μ .

Interval is $\bar{Y} \pm t$ standard errors.

t - value depends on df (usually $n-1$) and on probability in the tails.

T-Tests

Can test for difference in two treatment means.

Test is $t = (\bar{d} - \mu_d) / s_d$.

Paired t - test has s_d computed from differences.

Independent samples t - test has $S_d^2 = (S_1^2/n_1) + (S_2^2/n_2)$.

Least Significant Difference

Allows easy test for difference in two treatment means.

LSD = $t \sqrt{(2s^2/n)}$ for independent samples of size n .

Reflection

The **Module Reflection** appears as the last "task" in each module. The purpose of the Reflection is to enhance your learning and information retention. The questions are designed to help you reflect on the module and obtain instructor feedback on your learning. Submit your answers to the following questions to your instructor.

1. In your own words, write a short summary (< 150 words) for this module.
2. What is the most valuable concept that you learned from the module?
3. Why is this concept valuable to you? What concepts in the module are still unclear/the least clear to you?

Acknowledgements

This module was developed as part of the Bill & Melinda Gates Foundation Contract No. 24576 for Plant Breeding E-Learning in Africa.

Quantitative Methods Continuous Data Author: Ron Mowers, Ken Moore, M. L. Harbur, Kendra Meade, William Beavis, and Laura Merrick (ISU)

Multimedia Developers: Gretchen Anderson, Todd Hartnell, and Andy Rohrback (ISU)

How to cite this module: Mowers, R., K. Moore, M.L. Harbur, K. Meade, W. Beavis, and L. Merrick. 2016. Continuous Data. *In* Quantitative Methods, interactive e-learning courseware. Plant Breeding E-Learning in Africa. Retrieved from <https://pbea.agron.iastate.edu>.

Source URL: <https://pbea.agron.iastate.edu/course-materials/quantitative-methods/continuous-data-0?cover=1>