LECTURE 5:
CONFIDENCE INTERVALS

VICTORIA UNIVERSITY OF WELLINGTON
TE HERENGA WAKA

Introduction

- We have discussed point estimates:
  - as an estimate of population mean (point estimate is sample mean)
  - as an estimate of a success probability (point estimate is sample proportion)
- These point estimates are almost never exactly equal to the true values they are estimating because they are single points.
- In order for the point estimate to be useful, it is necessary to describe just how far off from the true value it is likely to be.
- One way to estimate how far our estimate is from the true value is to report an estimate of the standard deviation, or uncertainty.

Example 1

Assume that a large number of independent measurements from a normal population, all using the same procedure, are made on the diameter of a piston. The sample mean of the measurements is 14.0 cm, and the uncertainty in this quantity, which is the standard deviation of the sample mean, is 0.1 cm.

So, we have a high level of confidence that the true diameter is in the interval (13.7, 14.3). This is because it is highly unlikely that the sample mean will differ from the true diameter by more than three standard deviations.

Section 5.1: Large-Sample Confidence Interval for a Population Mean

Recall the previous example: Since the population mean will not be exactly equal to the sample mean of 14, it is best to construct a **confidence interval** around 14 that is likely to cover the population mean.

We can then quantify our level of confidence that the population mean is actually covered by the interval.

Constructing a CI

To see how to construct a confidence interval, let μ represent the unknown population mean and let σ^2 be the unknown population variance. Let $X_1, \ldots$ $\bar{X}$, X_{100} be the 100 diameters of the pistons. The observed value of is the mean of a large sample, and the Central Limit Theorem specifies that it comes from a normal distribution with mean μ and whose standard deviation is

$$\sigma_{\bar{X}} = \sigma / \sqrt{100}$$

Illustration of Capturing True Mean

Here is a normal curve, which represents the distribution of $\overline{X}$. The middle 95% of the curve, extending a distance of $1.96\sigma_{\overline{X}}$ on either side of the population mean μ, is indicated. The following illustrates what happens if $\overline{X}$ lies within the middle 95% of the distribution:

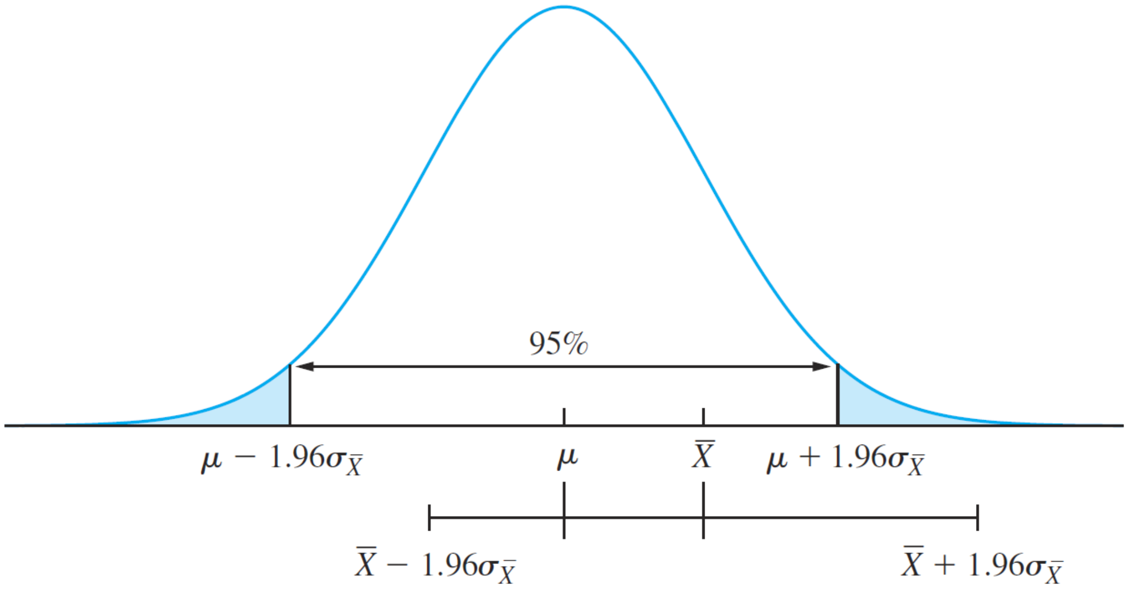

Illustration of Not Capturing True Mean

If the sample mean lies outside the middle 95% of the curve: Only 5% of all the samples that could have been drawn fall into this category. For those more unusual samples, the 95% confidence interval $\bar{X} \pm 1.96\sigma_{\bar{X}}$ fails to cover the true population mean μ.

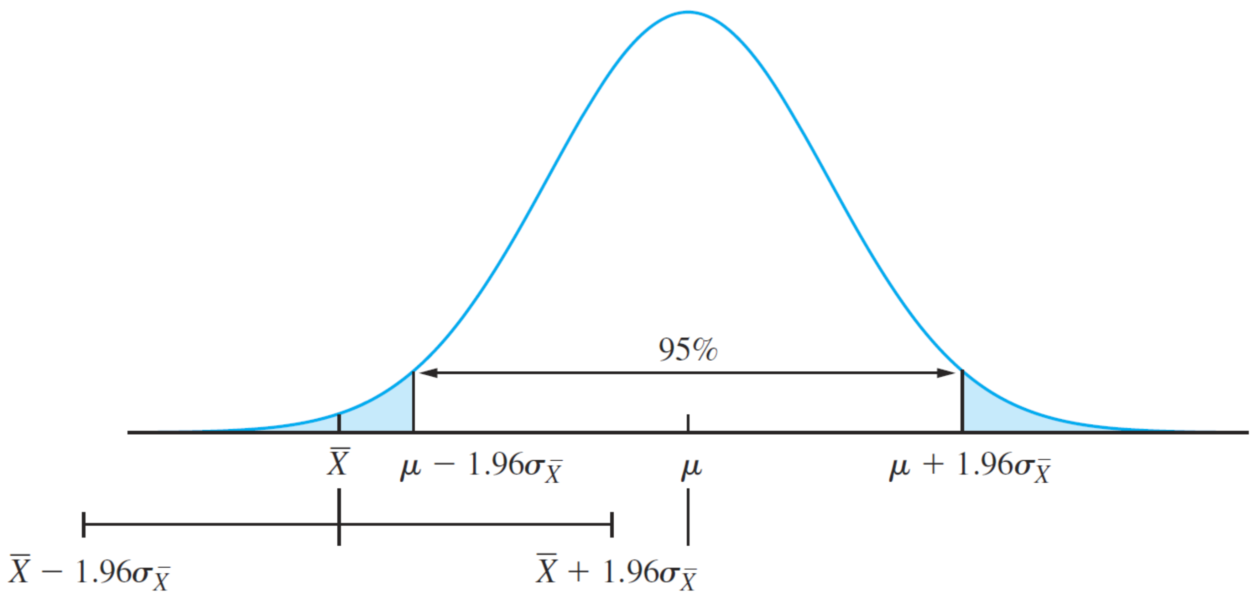

Computing a 95% Confidence Interval

The 95% confidence interval (CI) is $\overline{X} \pm 1.96\sigma_{\overline{X}}$.

So, a 95% CI for the mean is 14 $\pm$ 1.96(0.1). We can use the sample standard deviation as an estimate for the population standard deviation, since the sample size is large.

We can say that we are **95% confident**, or **confident at the 95% level**, that the population mean diameter for pistons lies, between 13.804 and 14.196.

Warning: The methods described here require that the data be a random sample from a population. When used for other samples, the results may not be meaningful.

Question?

- Does this 95% confidence interval actually cover the population mean μ?

  - It depends on whether this particular sample happened to be one whose mean came from the middle 95% of the distribution, or whether it was a sample whose mean was unusually large or small, in the outer 5% of the population.

  - There is no way to know for sure into which category this particular sample falls.

  - In the long run, if we repeatedly constructed these confidence intervals, then 95% of the samples will have means in the middle 95% of the population and 95% of the confidence intervals will cover the population mean.

Extension

- We are not always interested in computing 95% confidence intervals. Sometimes, we would like to have a different level of confidence.
- We can use the same reasoning we did for 95% confidence intervals to compute confidence intervals with various confidence levels.

Other Confidence Levels

- Suppose we are interested in 68% confidence intervals, then we know that the middle 68% of the normal distribution is in an interval that extends $1.0\sigma_{\bar{X}}$ on either side of the population mean.
- It follows that an interval of the same length around $\bar{X}$ specifically, will cover the population mean for 68% of the samples that could possibly be drawn.
- For our example, a 68% CI for the diameter of pistons is $14 \pm 1.0(0.1)$, or (13.9, 14.1).

100(1 − α)% CI

Let $X_1, \ldots, X_n$ be a *large* (n > 30) random sample from a population with mean μ and standard deviation σ, so that $\overline{X}$ is approximately normal. Then a level 100(1 − α)% confidence interval for μ is

$$\overline{X} \pm z_{\alpha/2}\sigma_{\overline{X}}$$

Where $\sigma_{\overline{X}} = \sigma / \sqrt{n}$. When the value of σ is unknown, it can be replaced with the sample standard deviation *s*.

Specific Confidence Intervals for μ

- $\bar{X} \pm \frac{s}{\sqrt{n}}$ is a 68% interval for μ.
- $\bar{X} \pm 1.645\frac{s}{\sqrt{n}}$ is a 90% interval for μ.
- $\bar{X} \pm 1.96\frac{s}{\sqrt{n}}$ is a 95% interval for μ.
- $\bar{X} \pm 2.58\frac{s}{\sqrt{n}}$ is a 99% interval for μ.
- $\bar{X} \pm 3\frac{s}{\sqrt{n}}$ is a 99.7% interval for μ.

Example 2

The sample mean for the fill weights of 100 boxes is 12.05 ounces, and the standard deviation is $s = 0.1$. Find an 85% confidence interval for the mean fill weight of the boxes.

Example 3

There is a sample of 50 microdrills with an average lifetime (expressed as the number of holes drilled before failure) of 12.68 and a standard deviation of 6.83. Suppose an engineer reported a confidence interval of (11.09, 14.27) but neglected to specify the level. What is the level of this confidence interval?

More About Confidence Levels

- The confidence level of an interval measures the reliability of the method used to compute the interval.
- A level $100(1 - \alpha)\%$ confidence interval is one computed by a method that in the long run will succeed in covering the population mean a proportion $1 - \alpha$ of all the times that it is used.
- In practice, there is a decision about what level of confidence to use.
- This decision involves a trade-off, because intervals with greater confidence are less precise.

Probability vs. Confidence

- In computing a CI, such as the one of diameter of pistons: (13.804, 14.196), it is tempting to say that the probability that μ lies in this interval is 95%.
- The term *probability* refers to random events, which can come out differently when experiments are repeated.
- The numbers 13.804 and 14.196 are fixed, not random. The population mean is also fixed. The mean diameter is either in the interval or not.
- There is no randomness involved.
- So, we say that we have 95% *confidence* (not probability) that the population mean is in this interval.

Example 4

A 90% confidence interval for the mean diameter (in cm) of steel rods manufactured on a certain extrusion machine is computed to be (14.73, 14.91). True or false: The probability that the mean diameter of rods manufactured by this process is between 14.73 and 14.91 is 90%.

Determining Sample Size for a Confidence Interval of Specified Width

Back to the example of diameter of pistons: We had a CI of (13.804, 14.196).

This interval specifies the mean to within ± 0.196. Now assume that the interval is too wide to be useful.

Assume that it is desirable to produce a 95% confidence interval that specifies the mean to within $\pm$ 0.1.

Determining Sample Size for a Confidence Interval of Specified Width

To do this, the sample size must be increased. The width of a CI is specified by $\pm z_{\alpha/2}\sigma/\sqrt{n}$.

If we know α and σ is specified, then we can find the *n* needed to get the desired width.

For our example, the $z_{\alpha/2} = 1.96$ and the estimated standard deviation is 1. So,

$$0.1 = 1.96(1)\big/\sqrt{n}$$

then the *n* that accomplishes this is 385 (always round up).

One-Sided Confidence Intervals

- We are not always interested in CIs with both an upper and lower bound (two-sided CIs).
- For example, we may only be interested in a lower bound for the mean crushing strength of a certain type of concrete block.
- With the same conditions as with the two-sided CI, the level $100(1-\alpha)\%$ lower confidence bound for μ is

$$\bar{X} - z_{\alpha}\sigma_{\bar{X}}.$$

and the level $100(1-\alpha)\%$ upper confidence bound for μ is

$$\bar{X} + z_{\alpha}\sigma_{\bar{X}}.$$

Example 2 cont.

Find both a 95% lower confidence bound and a 99% upper confidence bound for the mean fill weight of the boxes.

Section 5.2: Confidence Intervals for Proportions

- The method that we discussed in the last section was for a mean from any population from which a large sample has been drawn.
- When the population has a Bernoulli distribution, this expression takes on a special form.

95% CI for *p*

When *n* is large, the probability is 0.95 that the sample proportion is within 1.96 standard deviations of the true proportion:

$$p - 1.96\sqrt{\frac{p(1-p)}{n}} < \hat{p} < p + 1.96\sqrt{\frac{p(1-p)}{n}}$$

It is then also true that for 95% of all possible samples,

$$\hat{p} - 1.96\sqrt{\frac{p(1-p)}{n}} < p < \hat{p} + 1.96\sqrt{\frac{p(1-p)}{n}}$$

Comments

- This expression is not a confidence interval, because it contains the unknown population proportion *p* in the margin of error.
- We have to find an appropriate estimate of *p*.
- The traditional approach is to replace *p* with $\hat{p}$.
- Recent research shows that a slight modification of *n* and the following estimate of *p* provide a good confidence interval:
- Define $\tilde{n} = n + 4$ and $\tilde{p} = \dfrac{X + 2}{\tilde{n}}$

Confidence Interval for *p*

Let X be the number of successes in n independent Bernoulli trials with success probability p, so that

$X \sim \text{Bin}(n, p)$.

Then a $100(1 - \alpha)\%$ confidence interval for p is

$$\tilde{p} \pm z_{\alpha/2} \sqrt{\frac{\tilde{p}(1 - \tilde{p})}{\tilde{n}}}.$$

If the lower limit is less than 0, replace it with 0.

If the upper limit is greater than 1, replace it with 1.

Example 5

Interpolation methods are used to estimate heights above sea level for locations where direct measurements are unavailable. In an article in the *Journal of Survey Engineering,* a weighted-average method of interpolation for estimating heights from GPS measurements is evaluated. The method made "large" errors (errors whose magnitude was above a commonly accepted threshold) at 26 of the 74 sample test locations. Find a 90% confidence interval for the proportion of locations at which this method will make large errors.

More Intervals

- A level 100(1 − α)% lower confidence bound for p is

$$\tilde{p} - z_{\alpha}\sqrt{\frac{\tilde{p}(1-\tilde{p})}{\tilde{n}}}.$$

- A level 100(1 − α)% upper confidence bound for p is

$$\tilde{p} + z_{\alpha}\sqrt{\frac{\tilde{p}(1-\tilde{p})}{\tilde{n}}}.$$

Example 5 cont.

What sample size is needed to obtain a 95% confidence interval with width ± 0.08?

The Traditional Method

- Let $\hat{p}$ be the proportion of successes in a *large* number of independent Bernoulli trials with success probability *p*.
- Then the traditional level $100(1 - \alpha)\%$ confidence interval for *p* is

$$\hat{p} \pm z_{\alpha/2}\sqrt{\frac{\hat{p}(1-\hat{p})}{n}}.$$

- The method cannot be used unless the sample contains at least 10 successes and 10 failures.

Section 5.3: Small Sample CIs for a Population Mean

- The methods that we have discussed for a population mean previously require that the sample size be large.
- When the sample size is small, there are no good general methods for finding CIs.
- However, when the population is approximately normal, a probability distribution called the Student's *t* distribution can be used to compute confidence intervals for a population mean.

Small-Sample Confidence Intervals for the Mean

- What can we do if $\overline{X}$ is the mean of a *small* sample?

- If the sample size is small, *s* may not be close to σ, and X may not be approximately normal. If we know nothing about the population from which the small sample was drawn, there are no easy methods for computing CIs.

Small-Sample Confidence Intervals for the Mean

- However, if the population is approximately normal, it will be approximately normal even when the sample size is small. It turns out that we can use the quantity $(\bar{X} - \mu)/(s/\sqrt{n})$, but since s may not be close to σ, this quantity has a Student's *t* distribution.

Student's *t* Distribution

- Let $X_1,\ldots,X_n$ be a *small* ($n < 30$) random sample from a *normal* population with mean μ. Then the quantity

$$\frac{(\bar{X} - \mu)}{s / \sqrt{n}}$$

has a **Student's *t* distribution** with $n - 1$ degrees of freedom (denoted by t_{n-1}).

- When n is large, the distribution of the above quantity is very close to normal, so the normal curve can be used, rather than the Student's *t*.

More on Student's *t*

- The probability density of the Student's *t* distribution is different for different degrees of freedom.
- The *t* curves are more spread out than the standard normal distribution.
- Table A.3, called a ***t* table**, provides probabilities associated with the Student's *t* distribution.

Example 6

A random sample of size 10 is to be drawn from a normal distribution with mean 4. The Student's *t* statistic $t = (\overline{X} - 4)/(s/\sqrt{10})$ is to be computed. What is the probability that $t > 1.833$?

This *t* statistic has $10 - 1 = 9$ degrees of freedom. From the *t* table, $P(t > 1.833) = 0.05$.

Example 7

Find the value for the t_{14} distribution whose lower-tail probability is 0.01.

Look down the column headed with "0.01" to the row corresponding to 14 degrees of freedom. The value for $t = 2.624$. This value cuts off an area, or probability, of 1% in the upper tail. The value whose lower-tail probability is 1% is -2.624.

Student's *t* Confidence Interval

Let $X_1, \ldots, X_n$ be a *small* random sample from a *normal* population with mean μ. Then a level $100(1-\alpha)\%$ CI for μ is

$$\bar{X} \pm t_{n-1,\alpha/2} \frac{s}{\sqrt{n}}.$$

To be able to use the Student's *t* distribution for calculation and confidence intervals, you must have a sample that comes from a population that it approximately normal. Samples such as these rarely contain outliers. So if a sample contains outliers, this CI should not be used.

Other CIs

Let $X_1, \ldots, X_n$ be a *small* random sample from a *normal* population with mean μ.

- Then a level 100(1 - α)% upper confidence bound for μ is

$$\bar{X} + t_{n-1,\alpha} \frac{s}{\sqrt{n}}.$$

- Then a level 100(1 - α)% lower confidence bound for μ is

$$\bar{X} - t_{n-1,\alpha} \frac{s}{\sqrt{n}}.$$

Other CIs

- Occasionally a small sample may be taken from a normal population whose standard deviation σ is known. In these cases, we do not use the Student's *t* curve, because we are not approximating σ with *s*. The CI to use here is the one using the *z* table, which we discussed in the first section.

Example 8

An engineer reads a report that states that a sample of 11 concrete beams has an average compressive strength of 38.45 MPa with standard deviation 0.14 MPa. Should the *t* curve be used to find a confidence interval for the mean compressive strength?

Example 9

The article "Direct Strut-and-Tie Model for Prestressed Deep Beams" presents measurements of the nominal shear strength (in kN) for a sample of 15 prestressed concrete beams. The results are

580	400	428	825	850	875	920	550
575	750	636	360	590	735	950	

Assume that on the basis of a very large number of previous measurements of other beams, the population of shear strengths in known to be approximately normal, with standard deviation $\sigma = 180.0$ kN. Find a 99% confidence interval for the mean shear strength.

Section 5.4: Confidence Intervals for the Difference in Two Means

Set-Up:

Let X and Y be independent, with $X \sim N(\mu_X, \sigma_X^2)$ and $Y \sim N(\mu_Y, \sigma_Y^2)$. Then

$$X + Y \sim N(\mu_X + \mu_Y, \sigma_X^2 + \sigma_Y^2)$$

$$X - Y \sim N(\mu_X - \mu_Y, \sigma_X^2 + \sigma_Y^2).$$

A Confidence Interval for the Difference Between Two Means

- Let $X_1, X_2, \dots, X_{n_X}$ be a *large* random sample of size n_X from a population with mean μ_X and standard deviation σ_X, and let $Y_1, Y_2, \dots, Y_{n_Y}$ be a *large* random sample of size n_Y from a population with mean μ_Y and standard deviation σ_Y.

A Confidence Interval for the Difference Between Two Means

- If the two samples are independent, then a level

  $100(1 - \alpha)\%$ CI for $\mu_X - \mu_Y$ is

$$\bar{X} - \bar{Y} \pm z_{\alpha/2}\sqrt{\frac{\sigma_X^2}{n_X} + \frac{\sigma_Y^2}{n_Y}}.$$

- When the values of σ_X and σ_Y are unknown, they can be replaced with the sample standard deviations s_X and s_Y.

Example 10

The chemical composition of soil varies with depth. An article in *Communications in Soil Science and Plant Analysis* describes chemical analyses of soil taken from a farm in Western Australia. Fifty specimens were each taken at depths 50 and 250 cm. At a depth of 50 cm, the average NO_3 concentration (in mg/L) was 88.5 with a standard deviation of 49.4. At a depth of 250 cm, the average concentration was 110.6 with a standard deviation of 51.5. Find a 95% confidence interval for the difference in NO_3 concentrations at the two depths.

Section 5.5: Confidence Intervals for the Difference Between Two Proportions

Set-Up:

Let X be the number of successes in n_X independent Bernoulli trials with success probability p_X, and let Y be the number of successes in n_Y independent Bernoulli trials with success probability p_Y, so that $X \sim \text{Bin}(n_X, p_X)$ and $Y \sim \text{Bin}(n_Y, p_Y)$.

Define

$$\tilde{n}_X = n_X + 2, \tilde{n}_Y = n_Y + 2$$
$$\tilde{p}_X = (X + 1)/\tilde{n}_X, \tilde{p}_Y = (Y + 1)/\tilde{n}_Y$$

The Confidence Interval Formula

Given the set-up just described, the 100(1 − α)% CI for the difference $p_X - p_Y$ is

$$\tilde{p}_X - \tilde{p}_Y \pm z_{a/2}\sqrt{\frac{\tilde{p}_X(1-\tilde{p}_X)}{\tilde{n}_X} + \frac{\tilde{p}_Y(1-\tilde{p}_Y)}{\tilde{n}_Y}}$$

- If the lower limit of the confidence interval is less than −1, replace it with −1.
- If the upper limit of the confidence interval is greater than 1, replace it with 1.
- There is a traditional confidence interval as well. It is a generalization of the one for a single proportion.

Example 11

Methods for estimating strength and stiffness requirements should be conservative in that they should overestimate rather than underestimate. The success rate of such a method can be measured by a probability of an overestimate. An article in *Journal of Structural Engineering* presents the results of an experiment that evaluated a standard method for estimating the brace force for a compression web brace. In a sample of 380 short test columns (4 to 6 ft in length), the method overestimated the force for 304 of them, and in a sample of 394 long test columns (8 to 10 ft in length), the method overestimated the force for 360 of them. Find a 95% confidence interval for the difference between the success rates for long columns and short columns.

Section 5.6: Small-Sample Confidence Intervals for the Difference Between Two Means

Let $X_1, X_2, \ldots, X_{n_X}$ be a random sample of size n_X from a *normal* population with mean μ_X and standard deviation σ_X, and let $Y_1, Y_2, \ldots, Y_{n_Y}$ be a random sample of size n_Y from a *normal* population with mean μ_Y and standard deviation σ_Y.

Assume that the two samples are independent.

Section 5.6: Small-Sample Confidence Intervals for the Difference Between Two Means

If the populations do not necessarily have the same variance, a level

$100(1 - \alpha)\%$ CI for $\mu_X - \mu_Y$ is

$$\overline{X} - \overline{Y} \pm t_{\nu,\alpha/2}\sqrt{\frac{s_X^2}{n_X} + \frac{s_Y^2}{n_Y}}.$$

The number of degrees of freedom, ν, is given by (rounded down to the nearest integer):

$$\nu = \frac{\left(\frac{s_X^2}{n_X} + \frac{s_Y^2}{n_Y}\right)^2}{\frac{(s_X^2/n_X)^2}{n_X - 1} + \frac{(s_Y^2/n_Y)^2}{n_Y - 1}}$$

Example 12

Resin-based composites are used in restorative dentistry. An article presents a comparison of the surface hardness of specimens cured for 40 seconds with constant power with that of specimens cured for 40 seconds with exponentially increasing power. Fifteen specimens were cured with each method. Those cured with constant power had an average surface hardness (in N/mm^2) of 400.9 with a standard deviation of 10.6. Those cured with an exponentially increasing power had an average surface hardness of 367.2 with a standard deviation of 6.1. Find a 98% confidence interval for the difference in mean hardness between specimens cured by the two methods.

When the Populations Have Equal Variances

Suppose we have the same set-up as before, but the populations are known to have nearly the same variance. Then a 100(1 − α)% CI for

$\mu_X - \mu_Y$ is

$$\bar{X} - \bar{Y} \pm t_{n_X + n_Y - 2, \alpha/2} \cdot s_p \sqrt{\frac{1}{n_X} + \frac{1}{n_Y}}.$$

The quantity s_p is the pooled standard deviation, given by

$$s_p = \sqrt{\frac{(n_X - 1)s_X^2 + (n_Y - 1)s_Y^2}{n_X + n_Y - 2}}.$$

Example 13

A machine is used to fill plastic bottles with bleach. A sample of 18 bottles had a mean fill volume of 2.007 L and a standard deviation of 0.010 L. The machine is then moved to another location. A sample of 10 bottles filled at the new location had a mean fill volume of 2.001 L and a standard deviation of 0.012 L. It is believed that moving the machine may have changed the mean fill volume, but it is unlikely to have changed the standard deviation. Assume that both samples come from approximately normal populations. Find a 99% confidence interval for the difference between the mean fill volumes at the two locations.

Don’t Assume the Population Variances Are Equal…

- The assumption that the population variances are equal is very strict.
- The method can be quite unreliable if it is used when the population variances are not equal.
- Since we typically do not know the population variances, it is usually impossible to be sure that they are equal.

Don’t Assume the Population Variances Are Equal…

- When the *sample* variances are nearly equal, it is tempting to assume the population variances are nearly equal as well.
- However, with small sample sizes, the sample variances may not approximate the population variances well.
- Solution: The best practice is to assume the variances are unequal unless it is quite certain that they are equal.

Section 5.7: Confidence Intervals with Paired Data

- In some cases, we may wish to design an experiment so that each item in one sample is paired with an item in the other.
- When this is done, the samples are no longer independent.

Paired Data Example

- A tire manufacturer wishes to compare the tread wear of tires made of a new material with that of tires made of a conventional material. One tire of each type is place on each front wheel of each of 10 front-wheel-drive automobiles. The choice as to which type of tire goes on the right wheel and which on the left is made with the flip of a coin. Each car is driven 40,000 miles, then the tires are removed, and the depth of tread on each is measured.

The Set-Up for Paired Data

Let $(X_1,Y_1),\ldots, (X_n, Y_n)$ be sample pairs. Let $D_i = X_i - Y_i$. Let μ_X and μ_Y represent the population means for X and Y, respectively. We wish to find a CI for the difference $\mu_X - \mu_Y$. Let μ_D represent the population mean of the differences, then $\mu_D = \mu_X - \mu_Y$. It follows that a CI for μ_D will also be a CI for $\mu_X - \mu_Y$.

Now, the sample $D_1,\ldots, D_n$ is a random sample from a population with mean μ_D, we can use one-sample methods to find CIs for μ_D.

Confidence Interval

Let $D_1,\ldots, D_n$ be a *small* random sample ($n \leq 30$) of differences of pairs. If the population of differences is approximately normal, then a level $100(1-\alpha)\%$ CI for μ_D is

$$\overline{D} \pm t_{n-1,\alpha/2}\frac{s_D}{\sqrt{n}}.$$

If the sample size is large, a level $100(1-\alpha)\%$ CI for μ_D is

$$\overline{D} \pm z_{\alpha/2}\sigma_{\overline{D}}.$$

In practice, $\sigma_{\overline{D}}$ is approximated with $s_D/\sqrt{n}$.

Section 5.8:

Let $X_1, \ldots, X_n$ be a random sample from a normal population with variance σ^2. The sample variance is

$$s^2 = \frac{\sum_{i=1}^{n}(X_i - \overline{X})^2}{n-1}.$$

The quantity $\dfrac{(n-1)s^2}{\sigma^2} = \dfrac{\sum_{i=1}^{n}(X_i - \overline{X})^2}{\sigma^2}$

Has a chi-square distribution with n-1 degrees of freedom, denoted χ^2_{n-1}

Chi-Square Distribution

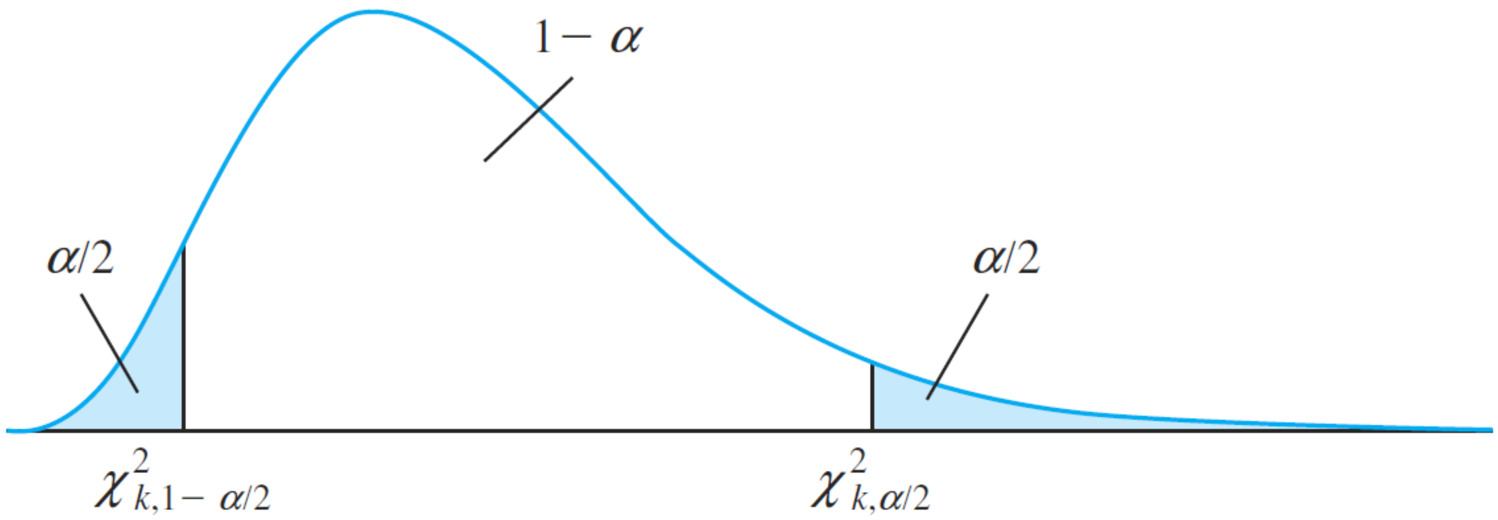

Confidence Interval

Let $X_1, \ldots, X_n$ be a random sample from a *normal* population with variance σ^2. Let s^2 be the sample variance. A level $100(1-\alpha)\%$ confidence interval for σ^2 is

$$\left(\frac{(n-1)s^2}{\chi^2_{n-1,\alpha/2}}, \quad \frac{(n-1)s^2}{\chi^2_{n-1,1-\alpha/2}} \right)$$

A level $100(1-\alpha)\%$ confidence interval for the standard deviation σ is

$$\left(\sqrt{\frac{(n-1)s^2}{\chi^2_{n-1,\alpha/2}}}, \quad \sqrt{\frac{(n-1)s^2}{\chi^2_{n-1,1-\alpha/2}}} \right)$$

Section 5.9: Prediction Intervals and Tolerance Levels

- A confidence interval for a parameter such as a population mean is an interval that is likely to contain the true value of the parameter.
- In contrast, prediction and tolerance intervals are concerned with the population itself, and with values that may be sampled in the future.
- The methods presented here are valid only when the population is known to be normal.

Prediction Intervals

- A **prediction interval** is an interval that is likely to contain the value of an item sampled from a population at a future time.

Prediction Intervals

- Let $X_1, \ldots, X_n$ be a sample from a *normal* population. Let *Y* be another item to be sampled from this population, whose value has not been observed.
- A 100(1 – α)% prediction interval for *Y* is

$$\bar{X} \pm t_{n-1,\alpha/2}\, s\sqrt{1+\frac{1}{n}}$$

- The probability is 1 – α that the value of *Y* will be contained in this interval.

Example 14

- A sample of 10 concrete blocks manufactured by a certain process had a mean compressive strength of 1312 MPa, with sample standard deviation 25 MPa. Find a 95% prediction interval for the strength of a block whose strength has yet to be measured.

Comparison of Prediction Intervals and Confidence Intervals

- The only difference between prediction intervals and confidence intervals from a computational standpoint is the ‘1’ that is under the square root for the prediction interval.
- This ‘1’ makes a prediction interval wider than its corresponding confidence interval, reflecting the fact that there is always uncertainty in the value of an item to be sampled from a population.

One-Sided Prediction Intervals

- One-sided prediction intervals are analogous to one-sided confidence intervals.
- Let $X_1, \cdots, X_n$ be a sample from a *normal* population. Let *Y* be another item to be sampled from this population, whose value has not been observed.

One-Sided Prediction Intervals

- A $100(1-\alpha)\%$ upper prediction bound for Y is

$$\bar{X} + t_{n-1,\alpha} s \sqrt{1 + \frac{1}{n}}$$

- A $100(1-\alpha)\%$ lower prediction bound for Y is

$$\bar{X} - t_{n-1,\alpha} s \sqrt{1 + \frac{1}{n}}$$

Tolerance Intervals for a Normal Population

- A **tolerance interval** is an interval that is likely to contain a specified proportion of the population.
- The method we will use requires that the population be normal.

Tolerance Intervals

- Let $X_1, \ldots, X_n$ be a sample from a *normal* population.
- A tolerance interval for containing

  100(1 – γ)% of the population with confidence 100(1 – α)% is $\bar{X} \pm k_{n,\alpha,\gamma} s$

- We use Table A.4 to find the value of $k_{n,\alpha,\gamma}$.

Example 15

- The lengths of bolts manufactured by a certain process are known to be normally distributed. In a sample of 30 bolts, the average length was 10.25 cm, with a standard deviation of 0.20 cm. Find a tolerance interval that includes 90% of the lengths of the bolts with 95% confidence.

Section 5.10: Using Simulation to Construct Confidence Intervals

If $X_1,\ldots, X_n$ are independent random variables with known standard deviations $\sigma_1,\ldots,\sigma_n$ and $U = U(X_1,\ldots, X_n)$ is a function of $X_1,\ldots, X_n$, then it will often (not always) be the case that U is approximately normally distributed and that its standard deviation can be estimated.

A simulation can be used to determine if U has an approximately normally distribution or not and to estimate its standard deviation.

Simulation

For the simulation, we generate say 1000 *X*'s from the normal distribution that specified the *X*'s. We transform these *X*'s into *U*'s. We can then check if the *U*'s are normally distributed by looking at a normal probability plot.

Summary

- We learned about large and small CIs for means.
- We also looked at CIs for proportions.
- We discussed large and small CIs for differences in means.
- We explored CIs for differences in proportions.