Chapter 5:

Point and Interval Estimation for a Single Sample

Introduction

- When data is collected, it is often with the purpose of estimating some characteristic of the population from which they came.
- The sample mean and sample proportion are examples of **point estimates** because they are single numbers, or points.
- More useful are **interval estimates**, also called **confidence intervals**.

Section 5.1: Point Estimation

- A numerical summary of a sample is called a **statistic**.
- A numerical summary of a population is called a **parameter**.

✓ Sample statistics are often used to estimate parameters.

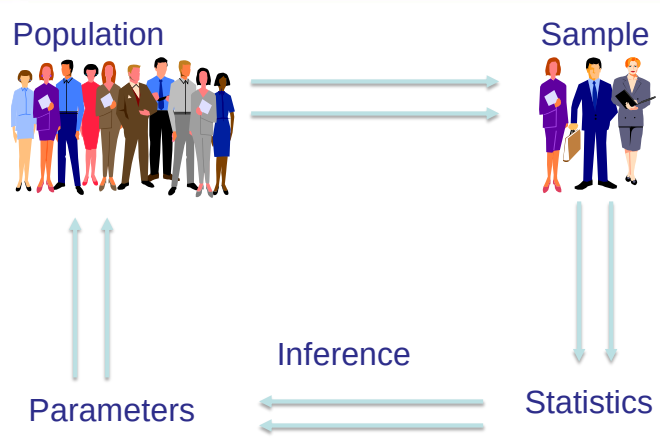

Summary of Process of Estimation

- We collect data for the purpose of estimating some numerical characteristic of the population from which they come.
- A quantity calculated from the data is called a statistic, and a statistic that is used to estimate an unknown constant, or parameter, is called a **point estimator**. Once the data has been collected, we call it a **point estimate**.

Questions of Interest

- Given a point estimator, how do we determine how good it is?
- What methods can be used to for construct good point estimators?

Notation: θ is used to denote an unknown parameter, and $\hat{\theta}$ to denote an estimator of θ.

Measuring the Goodness of an Estimator

- The accuracy of an estimator is measured by its bias, and the precision is measured by its standard deviation, or uncertainty.
- The bias is $Bias = \mu_{\hat{\theta}} - \theta$ (5.1)
- An estimator with a bias of 0 is said to be unbiased.
- To measure the overall goodness of an estimator, we used the **mean squared error** (MSE) which combines both bias and uncertainty.

Mean Squared Error

Let θ be a parameter, and $\hat{\theta}$ an estimator of θ. The mean squared error (MSE) of $\hat{\theta}$ is

$$MSE_{\hat{\theta}} = (\mu_{\hat{\theta}} - \theta)^2 + \sigma_{\hat{\theta}}{}^2 \quad (5.2)$$

An equivalent expression for the MSE is

$$MSE_{\hat{\theta}} = \mu_{(\hat{\theta}-\theta)^2} \quad (5.3)$$

Bias and Variance

Small Bias and Variance

Large Bias, Small Variance

Small Bias, Large Variance

Large Bias, Large Variance

Example 5.1 / 5.2

Let $X \sim Bin(n,p)$ where p is unknown. Find the bias and MSE of the estimator $\hat{p} = X/n$.

$$\mu_{\hat{p}} = \mu_{X/n} = \frac{\mu_X}{n} = \frac{np}{n} = p$$

$$Bias = \mu_{\hat{p}} - p = p - p = 0$$

$$\sigma_{\hat{p}}{}^2 = \sigma_{X/n}{}^2 = \frac{\sigma_X{}^2}{n^2} = \frac{np(1-p)}{n^2} = \frac{p(1-p)}{n}$$

$$MSE_{\hat{p}} = (\mu_{\hat{p}} - p)^2 + \sigma_{\hat{p}}{}^2 = \frac{p(1-p)}{n}$$

Limitations of Point Estimates

- Point estimates are almost never exactly equal to the true values they are estimating.
- Sometimes they are off by a little, sometimes a lot.
- For a point estimate to be useful, it is necessary to describe just how far off the true value it is likely to be.
- We could report the MSE with the point estimate.
- Many people just use interval estimates, called **confidence intervals.**

Confidence Intervals

- Confidence intervals consist of a lower limit and an upper limit.
- It includes a level of confidence, which is a number that tells us just how likely it is that the true value is contained within the interval.

Section 5.2: Large-Sample Confidence Interval for a Population Mean

Example 2: An important measure of the performance of an automotive battery is its cold cranking amperage, which is the current, in amperes, that the battery can provide for 30 seconds at 0°F while maintaining a specified voltage. An engineer wants to estimate the mean cold cranking amperage for batteries of a certain design. He draws a simple random sample of 100 batteries, and finds that the sample mean amperage is 185.5 A and the sample standard deviation is 5.0 A.

Example 2 (cont.)

- The sample mean of 185.5 is a point estimate of the population mean, which will not be exactly equal to the population mean.
- The population mean will actually be somewhat larger or smaller than 185.5.
- To get an idea of how much larger or smaller, we construct a confidence interval around 185.5 that is likely to cover the population mean.

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 $X_1, \ldots, X_{100}$ be the 100 amperages of the sample batteries. The observed value of the sample mean is 185.5. Since $\bar{X}$ 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 $\bar{X}$. The middle 95% of the curve, extending a distance of $1.96\sigma_{\bar{X}}$ on either side of the population mean μ, is indicated. The following illustrates what happens if $\bar{X}$ lies within the middle 95% of the distribution:

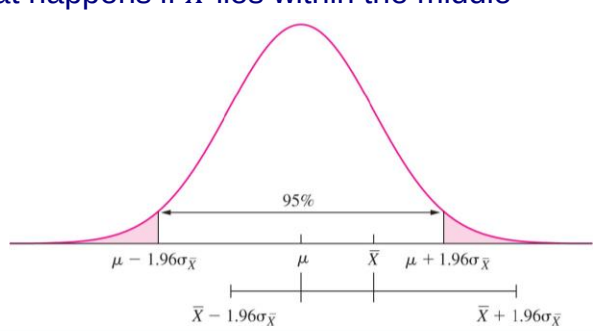

Illustration of Not Capturing True Mean

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.

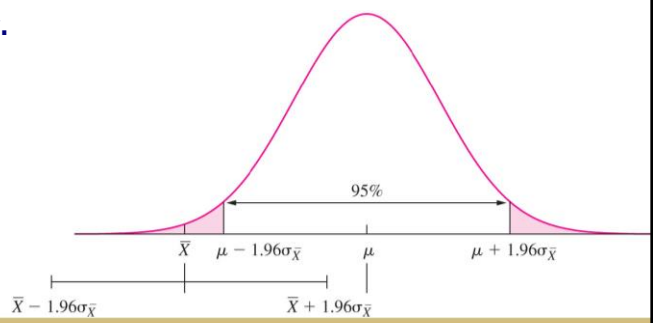

Computing a 95% Confidence Interval

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

So, a 95% CI for the mean is $185.5 \pm 1.96(\frac{5}{\sqrt{100}})$ or 185.5 ± 0.98 or (184.52, 186.48). 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 amperage lies between 184.52 and 186.48.

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 repeated these confidence intervals over and over, then 95% of the samples will have means in the middle 95% of the population. Then 95% of the confidence intervals will cover the population mean.

Confidence Level of a Confidence Interval

The confidence level is the proportion of all possible samples for which the confidence interval will cover the true value.

Pieces of CI

- Recall that the CI was 185.5 $\pm$ 0.98.
- 185.5 was the sample mean which is a point estimate for the population mean.
- We call the plus-or-minus number 0.98 the **margin of error**
- The margin of error is the product of 1.96 and $\sigma_{\bar{X}} = 0.5$.
- We refer to $\sigma_{\bar{X}}$ which is the standard deviation of $\bar{X}$, as the **standard error**.
- In general, the standard error is the standard deviation of the point estimator.
- The number 1.96 is called the **critical value** for the confidence interval. The reason that 1.96 is the critical value for a 95% CI is that 95% of the area under the normal curve is within – 1.96 and 1.96 standard errors of the population mean.

General Form of CI

Many confidence intervals follow the pattern just described. They have the form

point estimate ± margin of error

where

margin of error = (critical value)(standard error)

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 this reasoning to compute confidence intervals with various confidence levels.

Other CI 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 185.5 $\pm$ 1.0(0.5), or (185.0, 186.0).

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 $\bar{X}$ is approximately normal. Then a level 100(1 - α)% confidence interval for μ is

$$\bar{X} \pm z_{\alpha/2}\sigma_{\bar{X}} \qquad (5.4)$$

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

100(1 - α)% CI

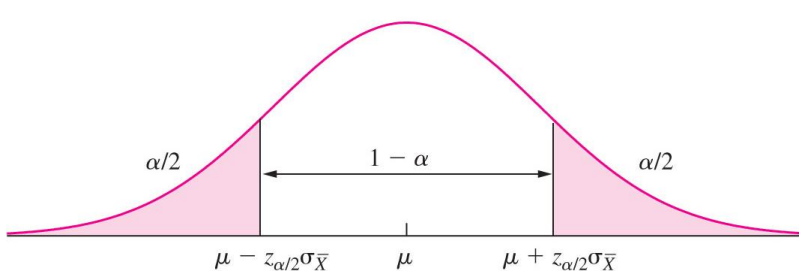

Specific Intervals for μ

- $\bar{X} \pm \frac{s}{\sqrt{n}}$ 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 5.4

In a sample of 50 microdrills drilling a low-carbon alloy steel, the average lifetime (expressed as the number of holes drilled before failure) was 12.68 with a standard deviation of 6.83. Find a 95% confidence interval for the mean lifetime of microdrills under these conditions.

Example 5.6

There is a sample of 50 micro-drills 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 CI's

- The confidence level of an interval measures the reliability of the method used to compute the interval.
- A level 100(1 - α)% confidence interval is one computed by a method that in the long run will succeed in covering the population mean a proportion 1 - α 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 CI, such as the one of amperage of batteries: (184.52, 186.48), 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.
- 184.52 and 186.48 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 that the population mean is in this interval.

Probability vs. Confidence

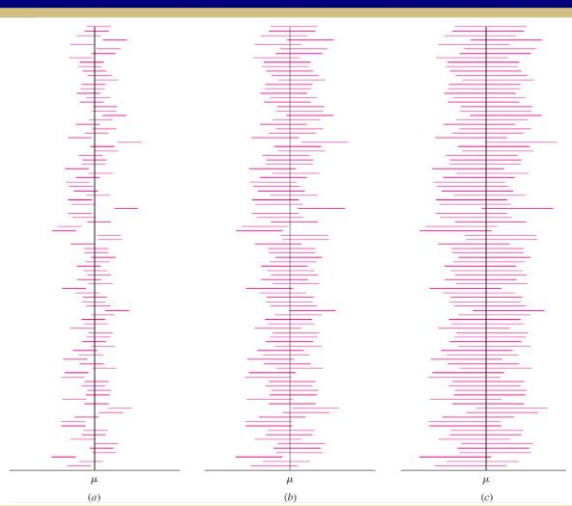

a) 68% CI
b) 95% CI
c) 99.7% CI

Example 5.7

A 90% confidence interval for the mean resistance (in Ω) of resistors is computed to be (1.43, 1.56). True or false: The probability is 90% that the mean resistance of this type of resistor is between 1.43 and 1.56.

Determining Sample Size

- In Example 5.4, a 95% CI was given by 12.68 ± 1.89.
- This interval specifies the mean to within ± 1.89. 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 ± 0.5.
- To do this, the sample size must be increased.
- The width of a CI is specified by $\pm z_{\alpha/2}\frac{\sigma}{\sqrt{n}}$
- If the desired with is ± *w* then $w = z_{\alpha/2}\frac{\sigma}{\sqrt{n}}$
- Solving this equation for *n* yields $n = {z_{\alpha/2}}^2\sigma^2/w^2$. (5.5)

Example 5.9

In the amperage example discussed earlier in this section, the sample standard deviation of amperages from 100 batteries was s = 5.0 A. How many batteries must be sampled to obtain a 99% CI of width ± 1.0 A.

One-Sided Confidence Intervals

- We are not always interested in CI's with both an upper and lower bound.
- For example, we may want a confidence interval on battery life. We are only interested in a lower bound on the battery life.
- With the same conditions as with the two-sided CI, the level 100(1-α)% lower confidence bound for μ is $\bar{X} - z_\alpha \sigma_{\bar{X}}$ (5.6)

  and the level 100(1-α)% upper confidence bound for μ is $\bar{X} + z_\alpha \sigma_{\bar{X}}$ (5.7)

Example 5.10

Find both a 95% lower confidence bound and a 99% upper confidence bound for the mean lifetime of the microdrills.

Section 5.3: 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 is drawn.
- We constructed a CI for the mean lifetime of a certain type of microdrill when drilling a low-carbon alloy steel.
- Now assume that a specification has been set that a drill should have a minimum lifetime of 10 holes drilled before failure.
- A sample of 144 microdrills is tested, and 120, or 83.3%, meet this specification.
- Let p represent the proportion of microdrills in the population that will meet the specification.
- We wish to find a 95% CI for p. To do this we need a point estimate, a critical value, and a standard error.

Traditional Approach

- To construct a point estimate for p, let X represent the number of drills in the sample that meet the specification.
- Then $X \sim Bin(n, p)$, where n = 144 is the sample size.
- The estimate for p is $\hat{p} = X/n$.
- In this example, X = 120, so $\hat{p}$ is 120/144 = 0.833
- Since the sample size is large, it follows from the Central Limit Theorem that $X \sim N(np, np(1-p))$.
- Since $\hat{p} = X/n$, it follows that

$$\hat{p} \sim N(p, (p(1-p)/n)$$

Standard Error of p

- The standard error of $\hat{p}$ is

$$\sigma_{\hat{p}} = \sqrt{\frac{p(1-p)}{n}}$$

- We can't use this value in the CI because it contains the unknown p.
- The traditional approach is to replace p with $\hat{p}$, obtaining $\sigma_{\hat{p}} = \sqrt{\frac{\hat{p}\,(1-\hat{p})}{n}}$

Traditional 95% CI for *p*

- Since $\hat{p}$ is approximately normal, the critical value for a 95% CI is 1.96 (obtained from the *z* table).
- So the traditional CI is $\hat{p} \pm 1.96\sqrt{\frac{\hat{p}\,(1-\hat{p})}{n}}$

Comments

- Recent research shows that a slight modification of *n* and the following estimate of *p* provide a better confidence interval.
- Define $\tilde{n} = n + 4$ and $\tilde{p} = \frac{X+2}{\tilde{n}}$

CI for *p*

Let X be the number of successes in n independent Bernoulli trials with success probability p, so that $X \sim Bin(n, p)$.

Then a 100(1 - α)% confidence interval for p is

$$\tilde{p} \pm z_{\alpha/2}\sqrt{\frac{\tilde{p}(1-\tilde{p})}{\tilde{n}}} \qquad (5.8)$$

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 5.11

Interpolation methods are used to estimate heights above sea level for locations where direct measurements are unavailable. In an article in *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 $\tilde{p} - z_{\alpha}\sqrt{\frac{\tilde{p}(1-\tilde{p})}{\tilde{n}}}$ (5.9)
- A level 100(1 - α)% upper confidence bound for p is $\tilde{p} + z_{\alpha}\sqrt{\frac{\tilde{p}(1-\tilde{p})}{\tilde{n}}}$ (5.10)
- The traditional CI (must have a least 10 successes and 10 failures in the sample) is $\hat{p} \pm z_{\alpha/2}\sqrt{\frac{\hat{p}\,(1-\hat{p})}{n}}$ (5.13)

Sample Size Formula

The sample size *n* needed to construct a level 100(1 – α)% CI of width ± *w* is

$$n = \frac{{z_{\alpha/2}}^2 \tilde{p}(1-\tilde{p})}{w^2} \quad (5.11)$$

if an estimate of $\tilde{p}$ is available, or

$$n = \frac{{z_{\alpha/2}}^2}{4w^2} \quad (5.12)$$

if no estimate of $\tilde{p}$ is available.

Example 5.12 (Cont. from 5.11)

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

Section 5.4: 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 general methods for finding CI's.
- If 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.

More on CI's

- What can we do if $\bar{X}$ is the mean of a *small* sample?
- If the sample size is small, *s* may not be close to σ, and $\bar{X}$ 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 CI's.
- 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 $(\bar{X} - \mu)/(\frac{s}{\sqrt{n}})$, 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

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

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 normal.
- Table A.3, called a ***t* table**, provides probabilities associated with the Student's *t* distribution.

Example 5.14

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

Example 5.17

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

Student's *t* CI

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

$$\bar{X} \pm t_{n-1,\alpha/2}\frac{s}{\sqrt{n}} \qquad (5.14)$$

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 is approximately normal. Samples such as these rarely contain outliers. So if a sample contains outliers, this CI should not be used.

Other CI's

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
  $\bar{X} + t_{n-1,\alpha} \frac{s}{\sqrt{n}}$ (5.15)
- Then a level 100(1 - α)% lower confidence bound for μ is
  $\bar{X} - t_{n-1,\alpha} \frac{s}{\sqrt{n}}$ (5.16)
- 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 5.18

Measurements of the nominal shear strength (in kN) for a sample of 15 prestressed concrete beams are

580 400 428 825 850 875 920 550

575 750 636 360 590 735 950

Assume this data follows a normal distribution, construct a 99% CI for the mean shear strength.

Example 5.19

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 180.0 kN. Find a 99% confidence interval for the mean shear strength.

Section 5.5: Prediction Intervals and Tolerance Intervals

- A confidence interval for a parameter is an interval that is likely to contain the true value of the parameter.
- Prediction and tolerance intervals are concerned with the population itself and with values that may be sampled from it in the future.
- These intervals are only useful when the shape of the population is known, here we assume the population is known to be normal.

Prediction Interval

- A prediction interval is an interval that is likely to contain the value of an item that will be sampled from the population at a future time.
- We "predict" that a value that is yet to be sampled from the population will fall within the predication interval.

100(1 – α)% Prediction Interval

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

$$\bar{X} \pm t_{n-1,\alpha/2} s \sqrt{1 + \frac{1}{n}} \qquad (5.19)$$

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

One sided intervals may also be constructed.

Example 5.21

A sample of 10 concrete blocks manufactured by a certain process has a mean compressive strength of 1312 MPa, with standard deviation of 25 MPa. Find a 95% prediction interval for the strength of a block that has not yet been measured.

Comparing CI and PI

- The formula for the PI is similar to the formula for the CI of a mean of normal population.
- The prediction interval has a small adjustment to the standard error with the additional + 1 under the square root.
- This reflects the random variation in the value of the sampled item that is to be predicted.
- Prediction intervals are sensitive to the assumption that the population is normal.
- If the shape of the population differs much from the normal curve, the prediction interval may be misleading.
- Large samples do not help, if the population is not normal then the prediction interval is invalid.

Tolerance Intervals

- A tolerance interval is an interval that is likely to contain a specified proportion of the population.
- First assume that we have a normal population whose mean μ and standard deviation σ are known.
- To find an interval that contains 90% of the population, we have $\mu \pm 1.645\sigma$.
- In general, the interval $\mu \pm z_{\gamma/2}\sigma$ will contain $100(1 - \gamma)\%$ of the population.
- In practice, we do not know μ or σ. Instead we use the sample mean and sample standard deviation.

Consequences

- Since we are estimating the mean and standard deviation from the sample,
  - We must make the interval wider than it would be if μ and σ were known.
  - We cannot be 100% confident that the interval actually contains the required proportion of the population.

Construction of Interval

- We must specify the proportion 100(1 – γ)% of the population that we wish the interval to contain.
- We must also specify the confidence 100(1 – α)% that the interval actually contains the specified proportion.
- It is then possible to find a number $k_{n,\alpha,\gamma}$ such that the interval

$$\bar{X} \pm k_{n,\alpha,\gamma} s \qquad (5.22)$$

will contain at least 100(1 – γ)% of the population with confidence 100(1 – α)%. Values of $k_{n,\alpha,\gamma}$ are presented in Table A.4.

Tolerance Interval Summary

Let $X_1,\ldots,X_n$ be a random sample from a *normal* population. A tolerance interval for containing at least 100(1 – γ)% of the population with confidence 100(1 – α)% is

$$\bar{X} \pm k_{n,\alpha,\gamma} s \qquad (5.22)$$

Of all the tolerance intervals that are computed by this method, 100(1 – α)% will actually contain at least 100(1 – γ)% of the population.

Example 5.22

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.

Summary

- Large sample CI's for means
- CI's for proportions
- Small sample CI's for means
- Prediction Intervals
- Tolerance Intervals