Chapter 19
Confidence intervals:
The basics

Types of Statistical Inference

- ***Confidence intervals for estimating the value of a population parameter***
- ***Tests of significance assesses the evidence for a claim about a population.***
  - ✓ ***Both types of inferences are based on the sampling distributions of statistics***
  - ✓ ***Both report probabilities that state "what would happen if we used the inference method many times"***

When you use statistical inference, you are acting as if the data are a random sample or come from a randomized experiment.

Objectives

Confidence intervals: the basics

- **Estimating with confidence**
- **Confidence intervals for the proportion or mean**
- **How confidence intervals behave**
- **Choosing the sample size**

Estimating with confidence

Although the sample mean, $\overline{x}$, is a unique number for any particular sample, if you pick a different sample, you will probably get a different sample mean.

In fact, you could get many different values for the sample mean, and virtually none of them would actually equal the true population mean, μ.

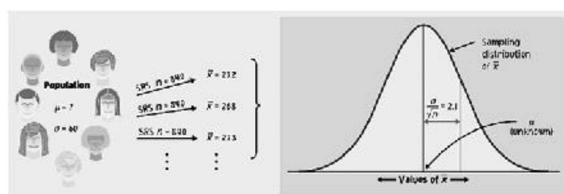

Statistical Inference

Statistical inference provides methods for drawing conclusions about a population from sample data.

What does ___ % confidence really mean?

"In repeated samples of the same size, the confidence created will catch the true value/parameter (p) _____ of the time."

In repeated samples of size 30, the conf. interval created will catch the true winning percent of the GAME 95% of the time.

We are 95% confident that the true % of chickens infected with Salmonella is between 11.95% and 18.05%.

What does ___ % confidence really mean?

Whenever we create a confidence interval, we write a sentence interpretation:

"Based on our sample, we are 95% confident that the <u>true % (or proportion)</u> of <u>(content)</u> is between <u>a</u> and <u>b</u> %.

Example: President Obama 45% approval rating, MOE of 3%, 95% confidence:

We are 95% confident that the true % approval rating of President Obama is between 42% and 48%.

68-95-99.7 Rule

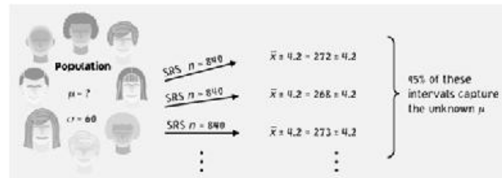

But the sample distribution is narrower than the population distribution, by a factor of √*n*.

Thus, the estimates $\bar{x}$ gained from our samples are always relatively close to the population parameter *μ*.

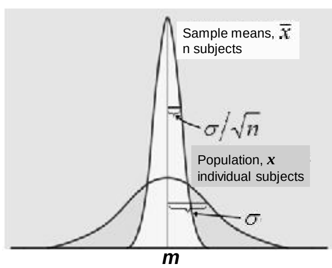

If the population is normally distributed $N(\mu,\sigma)$, so will the sampling distribution $N(\mu,\sigma/\sqrt{n})$.

Confidence interval

A level **C confidence interval** for a parameter has two parts:

- ✓ An interval calculated from the data, usually of the form

estimate ±margin of error

- ✓ A **confidence level C**, which gives the probability that the interval will capture the true parameter value in repeated samples, or the success rate for the method.

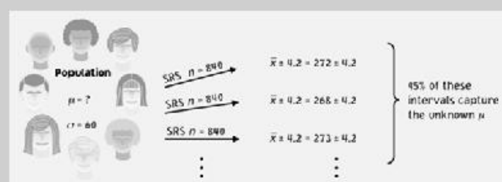

68-95-99.7 Rule

- ✓ In 95% of all samples, the mean score $\bar{x}$ for the sample will be within two standard deviations of the population mean score *m*. So the mean $\bar{x}$ of 500 SAT Math scores will be within 9 points of *m* in 95% of all samples.
- ✓ To say that $\bar{x} \pm 9$ is a 95% confidence interval for the population mean is to say that *in repeated trials*, 95% *of these intervals capture m*.
- ✓ ***We are** 95% **confident that the unknown mean SAT Math score for all California high school seniors lies between 452 and 470.***

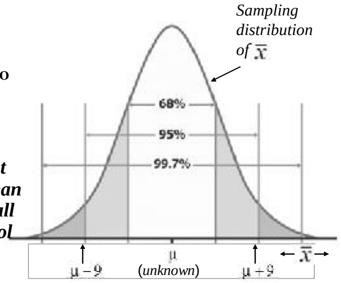

95% of all sample means will be within roughly 2 standard deviations (2**s*/√*n*) of the population parameter *m*.

Because distances are symmetrical, this implies that **the population parameter *m* must be within roughly 2 standard deviations from the sample average $\bar{x}$, in 95% of all samples.**

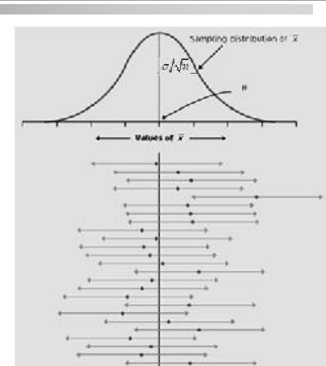

Red dot: mean value of individual sample

This reasoning is the essence of statistical inference.

The weight of single eggs of the brown variety is normally distributed *N*(65 g,5 g). Think of a carton of 12 brown eggs as an SRS of size 12.

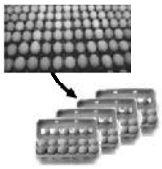

- What is the distribution of the sample means $\bar{x}$?
  Normal (mean m, standard deviation s/√*n*) = *N*(65 g,1.44 g).
- Find the middle 95% of the sample means distribution.
  Roughly ± 2 standard deviations from the mean, or 65g ± 2.88g.

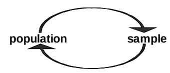

You buy a carton of 12 white eggs instead. The box weighs 770 g. The average egg weight from that SRS is thus $\bar{x}$ = 64.2 g.

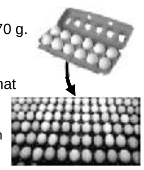

- Knowing that the standard deviation of egg weight is 5 g, what can you infer about the mean *μ* of the white egg population?

  There is a 95% chance that the population mean *μ* is roughly within ± 2s/√*n* of $\bar{x}$, or 64.2 g ± 2.88 g.

The important z* values

- Find the **z*** for a 90% C.I, 95% C.I. and for a 99% C.I.
- Summarize your results in a simple table

Confidence Level	Z*
90%	1.645
95%	1.960
99%	2.576

- **N(0, 1)**
- Use **invNorm(*p*, 0, 1)**

Confidence Interval for a Population Mean

Conditions for constructing a confidence interval for *m*

The construction of a confidence interval for a population m is appropriate when

- ***When the data come from an SRS from the population of interest, and***
- ***The sampling distribution of x-bar is approximately normal***

How do we find specific *z** values?

We can use a table of *z* values (Table A). For a particular confidence level *C*, the appropriate ***z**** value is just above it.

z*	0.674	0.841	1.036	1.282	1.645	1.960	2.054	2.326	2.576	2.807	3.091	3.291
Confidence level C	50%	60%	70%	80%	90%	95%	96%	98%	99%	99.5%	99.8%	99.9%

Ex. For a 98% confidence level, z*=2.326

We can use software. In **Excel:**

=NORMINV(probability,mean,standard_dev)
gives *z* for a given cumulative probability.

Since we want the middle *C* probability, the probability we require is (1 - C)/2

Example: For a 98% confidence level, = NORMINV (.01,0,1) = −2.32635 (= neg. *z**)

Constructing a level C confidence interval

Catch the central probability C under a normal curve

- ***Go out z* standard deviations on either side of the mean.***

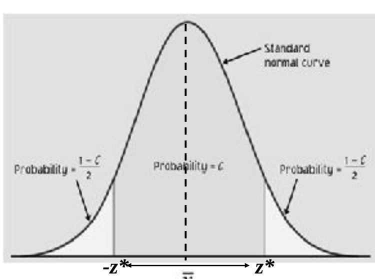

Interpreting a confidence interval for a mean

A **confidence interval** can be expressed as:

- $\bar{x} \pm z^*$ estimate

  z^* estimate is called the **margin of error**

A **confidence level *C*** (in %) indicates the success rate of the method that produces the interval.

It represents the area under the normal curve within ± ***z**** of the center of the curve.

- Two endpoints of an interval: ***m*** possibly within

  ($\bar{x} - z^*$ estimate) to ($\bar{x} + z^*$ estimate)

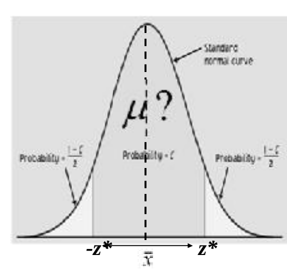

Confidence interval

The **confidence interval** is a range of values with an associated probability or **confidence level *C*.** The probability quantifies the chance that the interval contains the true population parameter.

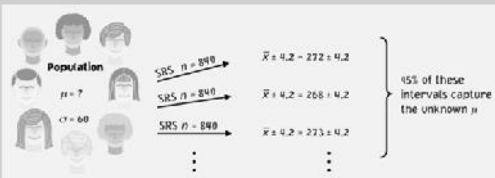

$\bar{x} \pm 4.2$ is a 95% confidence interval for the population parameter m

This equation says that in 95% of the cases, the actual value of mwill be within 4.2 units of the value of $\bar{x}$.

A **confidence interval** can be expressed as:

- Mean ± *m*
  m is called the **margin of error**
  mwithin $\bar{x} \pm m$
  Example: 120 ± 6
- Two endpoints of an interval
  mwithin ($\bar{x}$ - *m*) to ($\bar{x}$ - *m*)
  ex. 114 to 126

A **confidence level *C*** (in %) indicates the probability that the *μ* falls within the interval.

It represents the area under the normal curve within ± *m* of the center of the curve.

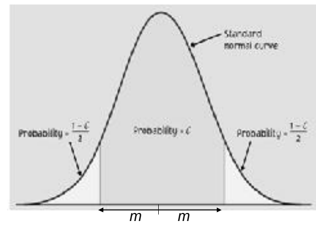

Implications

We <u>don't need</u> to take a lot of random samples to "rebuild" the sampling distribution and find m at its center.

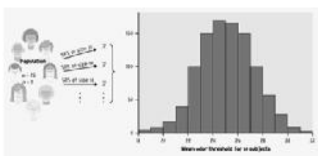

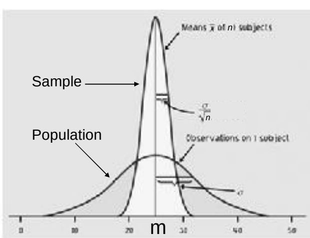

All we need is <u>one SRS</u> of size *n* and relying on the properties of the sample means distribution to infer the population mean m

Review: standardizing the normal curve using *z*

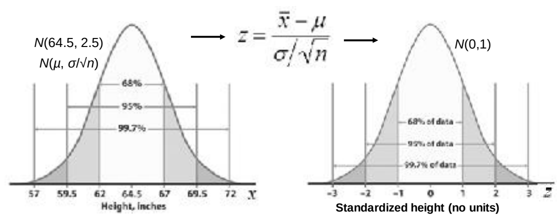

Here, we work with the sampling distribution, and *s*/√*n* is its standard deviation (spread).

Remember that *s* is the standard deviation of the original population.

Reworded

With 95% confidence, we can say that *μ* should be within roughly 2 standard deviations (2**s*/√*n*) from our sample mean $\bar{x}$ bar.

- In 95% of all possible samples of this size *n*, *μ* will indeed fall in our confidence interval.
- In only 5% of samples would $\bar{x}$ be farther from μ.

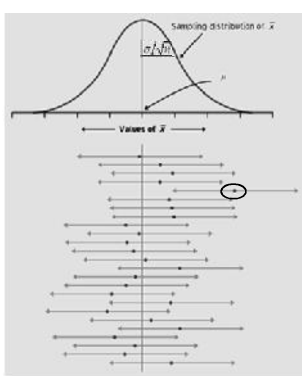

Varying confidence levels

Confidence intervals contain the population mean *m* in *C*% of samples.
Different areas under the curve give different confidence levels *C*.

Practical use of *z*: *z**

- *z** is related to the chosen confidence level *C*.
- *C* is the area under the standard normal curve between −*z** and *z**.

The confidence interval is thus:

$$\bar{x} \pm z^* \sigma/\sqrt{n}$$

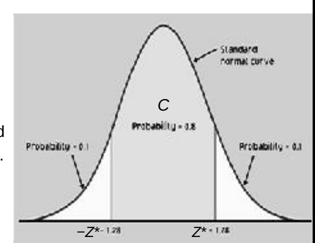

Example: For an 80% confidence level *C*, 80% of the normal curve's area is contained in the interval.

Link between confidence level and margin of error

The confidence level *C* determines the value of *z** (in Table *C*).

The margin of error also depends on *z**.

$$m = z^* \sigma / \sqrt{n}$$

Higher confidence ***C*** implies a larger margin of error ***m*** (thus less precision in our estimates).

A lower confidence level ***C*** produces a smaller margin of error ***m*** (thus better precision in our estimates).

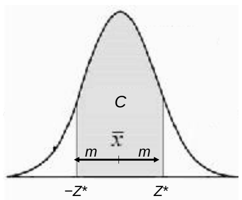

Sample size and experimental design

You may need a certain margin of error (e.g., drug trial, manufacturing specs). In many cases, the population variability (*s*) is fixed, but we can choose the number of measurements (*n*).

So plan ahead what sample size to use to achieve that margin of error.

$$m = z^* \frac{\sigma}{\sqrt{n}} \quad \Leftrightarrow \quad n = \left(\frac{z^* \sigma}{m}\right)^2$$

Remember, though, that sample size is not always stretchable at will. There are typically costs and constraints associated with large samples. The best approach is to use the smallest sample size that can give you useful results.

Different confidence intervals for the same set of measurements

Density of bacteria in solution:

Measurement equipment has standard deviation $s = 1*10^6$ bacteria/ml fluid.

3 measurements: 24, 29, and $31*10^6$ bacteria/ml fluid

Mean: $\bar{x} = 28*10^6$ bacteria/ml. Find the 96% and 70% CI.

- 96% confidence interval for the true density, $z^* = 2.054$, and write

$$\bar{x} \pm z^* \frac{\sigma}{\sqrt{n}} = 28 \pm 2.054(1/\sqrt{3})$$

$$= 28 \pm 1.19*10^6 \text{ bacteria/ml}$$

- 70% confidence interval for the true density, $z^* = 1.036$, and write

$$\bar{x} \pm z^* \frac{\sigma}{\sqrt{n}} = 28 \pm 1.036(1/\sqrt{3})$$

$$= 28 \pm 0.60*10^6 \text{ bacteria/ml}$$

z*	0.674	0.841	1.036	1.282	1.645	1.960	2.054	2.326	2.576	2.807	3.091	3.291
	50%	60%	70%	80%	90%	95%	96%	98%	99%	99.5%	99.8%	99.9%
	Confidence level C											

What sample size for a given margin of error?

Density of bacteria in solution:

Measurement equipment has standard deviation $\sigma = 1*10^6$ bacteria/ml fluid.

How many measurements should you make to obtain a margin of error of at most $0.5*10^6$ bacteria/ml with a confidence level of 90%?

For a 90% confidence interval, $z^* = 1.645$.

$$n = \left(\frac{z^* \sigma}{m}\right)^2 \Rightarrow n = \left(\frac{1.645*1}{0.5}\right)^2 = 3.29^2 = 10.8241$$

Using only 10 measurements will not be enough to ensure that *m* is no more than 0.5*106. Therefore, we need at least 11 measurements.

z*	0.674	0.841	1.036	1.282	1.645	1.960	2.054	2.326	2.576	2.807	3.091	3.291
	50%	60%	70%	80%	90%	95%	96%	98%	99%	99.5%	99.8%	99.9%
	Confidence level C											

Impact of sample size

The spread in the sampling distribution of the mean is a function of the number of individuals per sample.

- The larger the sample size, the smaller the standard deviation (spread) of the sample mean distribution.
- But the spread only decreases at a rate equal to $\sqrt{n}$.