Chapter 4
Confidence Interval of the Mean

Section 4.4:
Confidence Level, Sample Size and Margin of Error

School of Business Administration
Portland State University

- Confidence Level, Sample Size and Margin of Error
  - Confidence Level
  - Decreasing the Margin of Error
  - Choose the Needed Sample Size
  - An Application

4.4a
Confidence Level

How Likely that the Interval Contains μ?

Confidence Level and t-cutoff

Many possible confidence levels

- 0.95 is only one of many possible confidence levels
- To change the confidence level, change the t-cutoff value
- Each t-cutoff value corresponds to a different range of sampling variation, such as for 90% and 99%
- Consider first the corresponding z-cutoff values, which set the baseline for the slightly larger t-cutoffs

90% Range of Variation: 0.95 Quantile, $z_{.05} = 1.64$

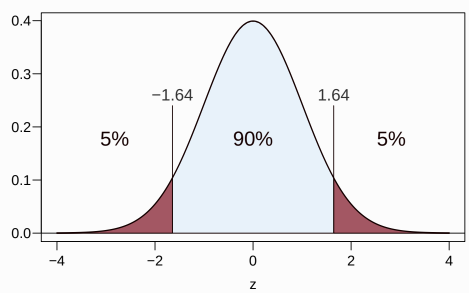

Figure: Normal Distribution: 90% Range of Variation, within 1.64 standard errors of the population mean, μ

- 90% of the values of any normal distribution, including for m, fall within 1.64 standard deviations of the mean, μ

99% Range of Variation: 0.995 Quantile, $z_{.005} = 2.58$

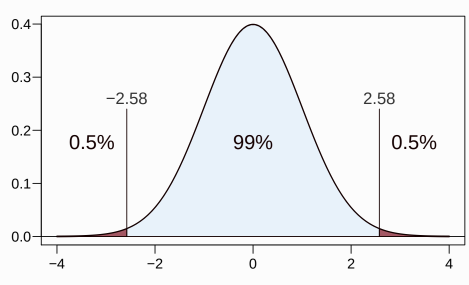

Figure: Normal Distribution: 99% Range of Variation, within 2.58 standard errors of the population mean, μ

- 99% of the values of any normal distribution, including for m, fall within 2.58 standard deviations of the mean, μ

Confidence Level and t-cutoff

To change the confidence level, change the t-cutoff value

- In practice, use the corresponding t-cutoff to set the confidence level
- Each t-cutoff is a little larger than the corresponding z-cutoff
- 90% Confidence Interval About a Sample Mean

  $m \pm t_{.05}(s_m), \quad t_{.05} > 1.65$
- 95% Confidence Interval About a Sample Mean

  $m \pm t_{.025}(s_m), \quad t_{.025} > 1.96$
- 99% Confidence Interval About a Sample Mean

  $m \pm t_{.005}(s_m), \quad t_{.005} > 2.57$

The Ideal Confidence Interval

Ideal not usually obtainable

- **Key Concept**: The ideal confidence interval has two desirable goals
  - is narrow, that is, a low margin of error, E
  - with a high level of confidence of containing μ
- For example, when estimating mean income, would like to have μ within a range of $\pm$\$1 units at 99% confidence
- Unfortunately, there is a trade off between having two mutually incompatible, distinct goals:

  high confidence vs. a narrow interval
- Recalculating a confidence interval to obtain either more confidence or a narrower interval results in a *less desirable* value of the other characteristic

Trade off: Confidence Level vs. Margin of Error

Desire: High confidence the CI contains unknown μ

- Action: To increase probability of containing the unknown value of μ, move confidence level from 95% to 99%
- Trade off: Larger confidence level results in wider interval, and so a larger margin of error

Desire: Low margin of error, i.e., precision

- Action: To decrease size of confidence interval, decrease confidence level from 95% to 90%
- Trade off: Narrower interval yields a lower probability that the interval actually contains μ

Information and Confidence Level

Trade off means confidence level is to some extent arbitrary

- Each choice of confidence level provides the same amount of information, balancing margin of error against confidence level
- So usually choose 95%
  - 95% is a high level of confidence
  - 95% is a "nice" number, divisible by 5, that still gets most of the values

Illustration: Different Confidence Levels

Most common are 95%, and, to a lessor extent, 90% and 99%

- Consider the previous application that analyzed the 95% confidence interval for the average ship time of a supplier
- Data Summary. $n = 15$, $m = 8.087$ days, $s = 1.587$ days

Level	$t_{.025}$	LB	UB	Width
.90	1.761	7.36	8.81	1.44
.95	2.145	7.21	8.97	1.76
.99	2.977	6.87	9.31	2.44

- The narrowest width, of 1.44 days, comes with the least confidence, .90
- Moving the confidence level up to .99 increases the confidence interval's width all the way to 2.44 days

4.4b
Decreasing the Margin of Error
A More Precise Estimate

Desired Margin of Error

Need more information

- The analysis of the same data with different confidence levels just re-packages the existing information
- **Key Concept**: To obtain a lower margin of error for a 95% confidence interval, increase sample size, n
- To illustrate, consider three simulations of 50 random samples, of size 10, 100 and then 1000, and observe the resulting confidence intervals, computed for each sample

> To obtain many confidence intervals of the mean from repeated sampling of simulated data, use the `lessR` function `sim.CImean`, introduced in Section 4.3
>
> Set the boundaries of the y-axis with the option `ylim.bound`, which specifies the upper and lower boundaries of the y-axis

Fifty 95% Confidence Intervals, $n = 10$ for Each Sample

> Set `ns` for 50 samples, and `n` for each sample to have 10 elements

```
> sim.CImean(ns=50, n=10, ylim.bound=1.5)
```

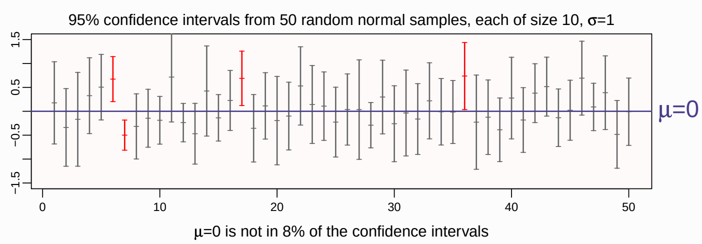

- At such a small sample size, $n = 10$, the 95% confidence intervals tend to be wide and of widely divergent sizes

Fifty 95% Confidence Intervals, $n = 100$ for Each Sample

```
> sim.CImean(ns=50, n=100, ylim.bound=1.5)
```

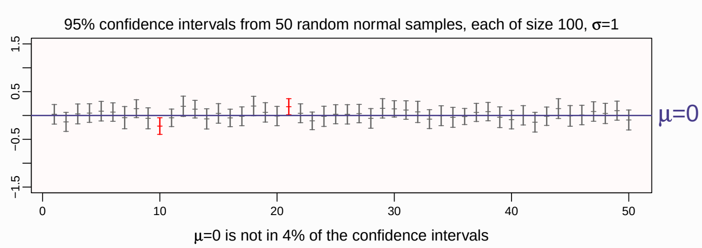

- With a sample size of 100, the 95% confidence intervals are much smaller than for a sample size of 10
- Their sizes still differ, but are much less divergent

Fifty 95% Confidence Intervals, $n = 1000$ for Each Sample

```
> sim.CImean(ns=50, n=1000, ylim.bound=1.5)
```

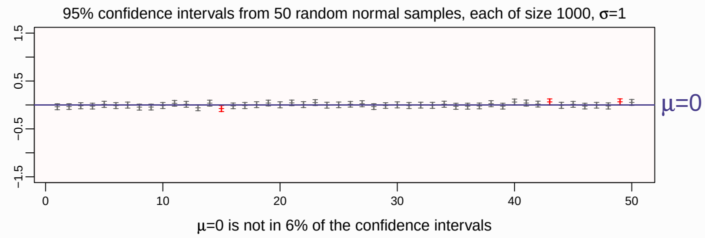

- For $n = 1000$, the 95% confidence intervals are very small
- Still, about 5% of the intervals do not contain μ=0, but all m's are relatively close to μ

Fifty 99% Confidence Intervals, $n = 1000$ for Each Sample

Change the confidence level from the default value of 0.95 with the option `cl`

```
> sim.CImean(ns=50, n=1000, ylim.bound=1.5, cl=.99)
```

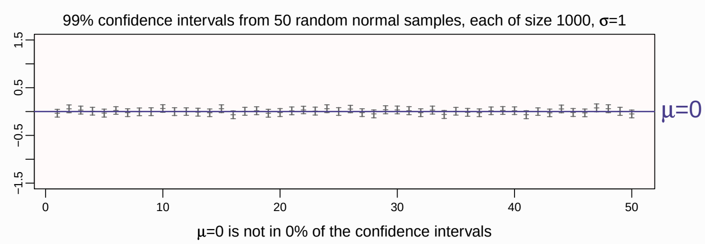

- With $n = 1000$, leverage the small margin of error
- **Key Concept**: With a large n, move to a 99% level of confidence and still obtain a small margin of error

Margin of Error as a function of n and Confidence Level

- This 3-D graph shows the dependence of margin of error, E, on n and confidence level
- Use z, so known σ is just a constant scale factor, set $\sigma = 1$,

$$E = [z\text{-cut}]\left[\frac{\sigma}{\sqrt{(n)}}\right]$$

- The largest error occurs at small n and large confidence level
- Huge decrease in E from the peak

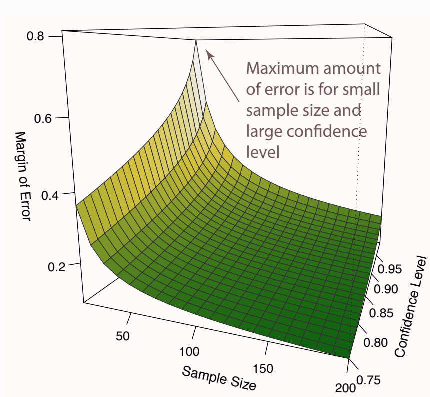

More is Always Better

Qualification . . . Diminishing returns on sample size

- The larger the sample, the more information provided
- However, the relationship between sample size and obtained information is nonlinear
- Huge gains result from increasing a very small sample size to a moderate sample size
- Unfortunately, the marginal increase in information diminishes rapidly as sample size increases beyond the very small sample sizes

4.4c
Choose the Needed Sample Size

Desired Margin of Error

Distinguish between what you got and what you want

- What you have: E, the margin of error *obtained* with the acquired sample of size of n
- What you want: $E_{desired}$, the **margin of error *desired***, which requires an increase of the current n to the larger value n_{needed}
- Question: How large a sample is needed to obtain the smaller margin of error that is desired?
- As always in life and statistics, there are no guarantees
- The goal is to have a very high probability that when the new confidence interval is calculated from the larger sample, the desired margin of error will be obtained
- **Key Concept**: Calculate the needed larger sample size, n_{needed}, so that the new 95% confidence interval has a .90 probability of obtaining the desired margin of error, $E_{desired}$

Overview of Sample Size Procedure

Move from initial sample to final, larger sample

- Specify the desired margin of error (precision), $E_{desired}$
- Obtain initial data sample
- Calculate margin of error, E, then the confidence interval
- If $E_{desired} < E$, calculate n_{needed}
- Gather new data for larger sample
- Re-calculate the margin of error, E, and the confidence interval

Two-step Procedure to Calculate the Needed Sample Size

Two-step procedure for obtaining needed sample size

- 1. From the initial sample, get the standard deviation, s, and then calculate preliminary estimate of sample size, n_s :

$$n_s = \left[\frac{(1.96)(s)}{E_{desired}}\right]^2$$

- 2. Calculate actual estimate of sample size: s may underestimate σ, so revise n_s upward for a given probability[1] of obtaining $E_{desired}$ with a 95% level of confidence
  .70 probability: $n_{needed} = 1.054n_s + 4.532$
  .90 probability: $n_{needed} = 1.132n_s + 7.368$ ▷ Most often used
  .99 probability: $n_{needed} = 1.242n_s + 10.889$

[1] These coefficients come from analysis of a paper by Kupper and Hafner in the *American Statistician*, 43(2):101-105, 1989

Two-step Procedure to Calculate the Needed Sample Size

Two-step procedure for obtaining needed sample size

- The adjustment equation is a mathematical adjustment to explicitly account for the variability of the standard deviation estimate, s, across repeated samples
- In practice, the standard error of the sample standard deviation is quite large in small samples
- So by chance the one obtained sample standard deviation, s, may be a very large underestimate of the true population standard deviation, σ
- The result would be a value of estimated sample size that would be too small to obtain the desired margin of error with any reasonable probability
- Without this adjustment the actual probability of getting the desired margin of error is quite a bit less than .90

4.4d
An Application

Reconsidering Ship Times

Previous CI of 7.21 to 8.97 too large

- Previous business application constructed a confidence interval about the mean delivery time of the 15 last shipments from a single company
- Unfortunately, management deemed the precision of estimation as not sufficiently precise, achieving an obtained margin of error, E, of only 0.879 days
- Instead, management specifies a maximum desired margin of error of only one-half day
- Managerial Question: How many shipments need be sampled to reach a .90 probability of obtaining a margin of error of half a day or less at the 95% confidence level?

Excel Worksheet: Needed Sample Size for 95% Confidence

0.9 probability of getting desired margin of error of 0.5 days

	Description	Name	Value	Formula
INPUT: DESCRIPTIVE STATISTICS	count of data	n	15	*COUNT(data)*
	mean of data	mean	8.087	*AVERAGE(data)*
	standard dev of data	stdev	1.587	*STDEV(data)*

	Description	Name	Value	Formula
SAMPLE SIZE	desired precision	Edesire	0.5	
	z cutoff for 95% CI	zcut	1.960	NORMINV(0.975,0,1)
	initial sample size	ns	38.70	*((zcut*stdev)/Edesire)^2*
	needed sample size	needed	51.18	*1.132*ns+7.368*
	sample size rounded up	size	**52**	*CEILING(needed,1)*

Analysis of Sample Size

Step 1: Initial sample size

- From previous analysis of initial sample of $n = 15$ shipments, $s = 1.587$, with $E_{desired}$ specified as 0.5

$$n_s = \left[\frac{(z_{.025})(s)}{E_{desired}}\right]^2 = \left[\frac{(1.96)(1.587)}{0.5}\right]^2 = 38.69$$

Step 2: Upward adjusted actual sample size

- $n_{needed} = 1.132 n_s + 7.368$
  $= 1.132(38.69) + 7.368$
  $= 51.17$
- Now round up, 51.17 to: $n_{needed} = 52$

Conclusion

Results

- Now collect information on 52 shipments instead of 15, so 37 more shipments needed
- There is a .90 probability that when the revised 95% confidence interval is calculated over all 52 shipments, the resulting margin of error will be 0.5 or less

Potential Problems

- 52 shipments are not available
- Even if available, the older shipments are from so long ago the underlying process has changed, rendering them invalid for estimating the *current* population mean shipping time, μ
- Sometimes statistical analysis simply is not definitive

Index Subtract 2 from each listed value to get the Slide

- The End