Standard Error of the Mean

PSCI 1801 · Class Handout

Dr. Marc Trussler

This handout is designed as an in-class supplement to the textbook. The material contained on this handout does not represent the full set of information you need to know for the exams. Exclusively studying off of this handout will not be sufficient for the exams.

The Three Distributions

Three distributions get confused constantly. Keep them separate.

1. **The population** — the true data-generating process. Can be any shape.
2. **The sample** — a set of n observations drawn from the population. A rough approximation.
3. **The sampling distribution of the mean** — the distribution of sample means we would get if we repeated the sampling process. Always normal (by CLT), centered on μ, with SD $= \sigma/\sqrt{n}$.

Statistics has nothing *useful to say about how close the sample distribution is to the population distribution. The whole game is that we can predict how close the* sample mean *is to the population mean.*

But We Don't Know the Population

The CLT gives us $\bar{X}_n \sim N(\mu, \sigma^2/n)$*. But* σ^2 *is a population quantity — in real life we don't know it!*

The fix: use the sample variance S^2 *as a stand-in for the population variance* σ^2*.*

$$S^2 = \frac{1}{n-1}\sum_{i=1}^{n}(X_i - \bar{X})^2$$

The divisor is $n-1$, not n. This adjustment makes S^2 an unbiased estimator of σ^2. R's `var()` and `sd()` do this automatically.

The estimated standard error

Once we plug in S for σ, we get the estimated standard error:

$$\widehat{SE}(\bar{X}) = \frac{S}{\sqrt{n}}$$

Turns out this substitution works well — the sampling distributions built with S^2 are almost indistinguishable from ones built with the true σ^2.

Confidence Intervals

A confidence interval is a zone around our sample mean that is likely to contain the true population mean.

Construction: center the sampling distribution on the sample mean *(not the population mean, which we don't know), then find the range that captures the desired probability mass.*

Confidence level $= (1-\alpha)\cdot 100\%$. So $\alpha = 0.05$ gives a 95% CI, $\alpha = 0.10$ gives a 90% CI, etc.

$$CI(\alpha) = \left[\bar{X}_n - z_{\alpha/2}\cdot SE,\ \bar{X}_n + z_{\alpha/2}\cdot SE\right]$$

Where $z_{\alpha/2}$ is the critical value from the standard normal. Recall from Sampling that the standard error is the standard deviation of the sampling distribution, $\sigma/\sqrt{n}$ — here we use the sample-based version $s/\sqrt{n}$ because we don't know σ.

Common critical values:

Confidence	α	$z_{\alpha/2}$
90%	0.10	1.645
95%	0.05	1.960

Confidence	α	$z_{\alpha/2}$
99%	0.01	2.576

Worked example

Sample: $n = 100$ *women's heights,* $\bar{X} = 63.92$*,* $S^2 = 7.88$*. Compute a 70% CI:*

$$SE = \sqrt{7.88/100} = 0.28$$

$$\alpha = 0.30, \quad z_{\alpha/2} = 1.04$$

$$CI = [63.92 - 1.04 \cdot 0.28,\ 63.92 + 1.04 \cdot 0.28] = [63.62, 64.21]$$

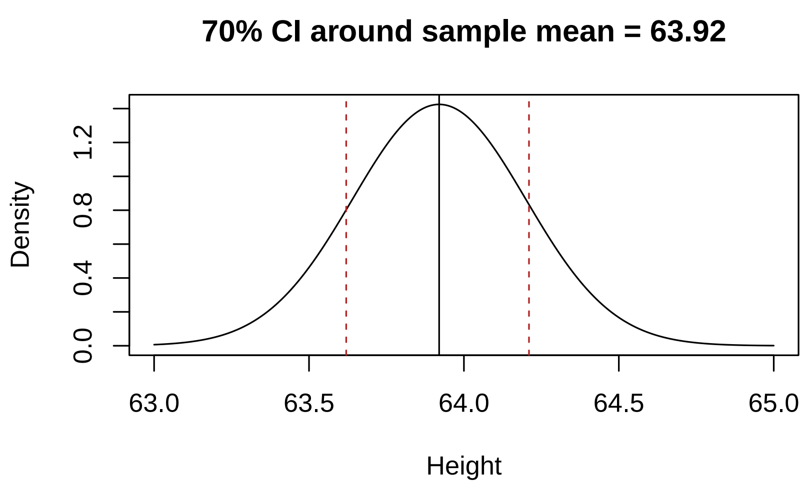

Interpretation: 70% of similarly constructed confidence intervals contain the true population parameter.

Pedantic language note: *A given CI either contains the truth or it doesn't — there's no probability attached to it once we've calculated it. The formally correct statement is "70% of similarly constructed CIs contain the truth." In casual usage, saying "we're 70% confident the truth is in this range" is common but technically wrong.*

Simulation check

If we repeatedly sample and construct 70% CIs, about 70% of them should contain the truth ($\mu = 64$):

```
[1] 0.6977
```

0.70. Confirmed.

The T Distribution

One complication with substituting s^2 for σ^2: at small n, the sampling distribution isn't quite normal.

Test it at $n = 10$. We use $z_{0.15} = 1.04$ (the standard-normal 70% cutoff) and simulate 10,000 CIs:

```
[1] 0.682
```

0.68. Systematically wrong — only 68% coverage instead of 70%.

The reason: once we sub s^2 for σ^2, the pivot is not standard normal but rather Student's t *with $n - 1$ degrees of freedom:*

$$\frac{\bar{X}_n - E[X]}{s/\sqrt{n}} \sim t_{n-1}$$

The t distribution looks a lot like the standard normal but has fatter tails *at small n. It converges to the normal as n grows (see figure — dashed line is standard normal, solid line is t_9).*

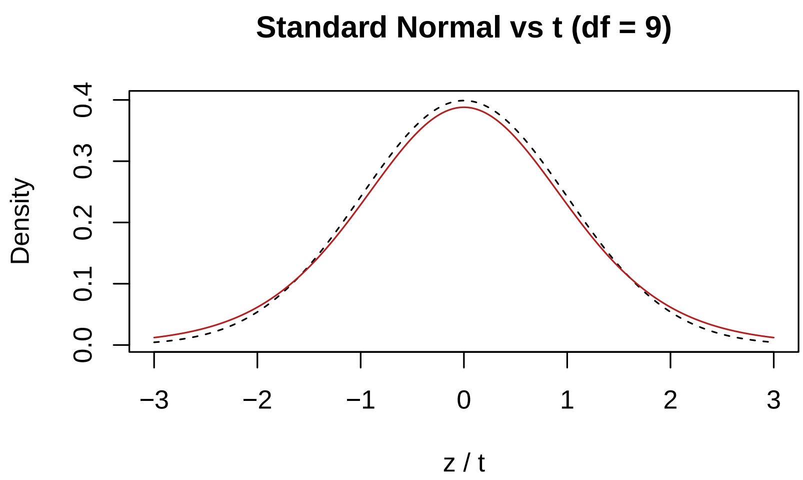

Use `qt(p, df)` *for critical values from the t distribution instead of* `qnorm(p)` *from the standard normal. Correct 70% CI at $n = 10$:*

```
[1] -1.036433
```

```
[1] -1.099716
```

Rebuilt with $t_{9,\,0.15} = 1.10$ instead of $z_{0.15} = 1.04$, coverage lands where it should.

Rule: *Whenever you use s^2 in place of σ^2 (which is always, in real research), the CI critical value comes from the t distribution with $n - 1$ degrees of freedom. For $n > 30$ the difference from the normal is negligible; for smaller n it matters.*

Coming Next

Confidence intervals answer "given my sample, where does the truth plausibly lie?" Hypothesis testing asks the mirror-image question: "given a specific claim about the truth, is my sample consistent with it?" Same machinery, flipped framing. That's next chapter.