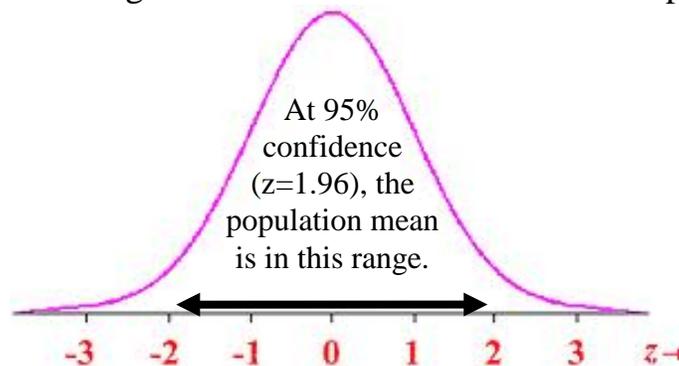


LECTURE 21: UNDERSTANDING HYPOTHESIS TESTING

- I. The Nature of Hypothesis Testing
- So far I've given you a lot of information on the parts of hypothesis testing. Today we assemble it all together.
 - Let's begin on familiar ground: the Central Limit Theorem. We know that if you take a bunch of samples, the means of those samples form a normal distribution. In practice, this implies that for any sample we pull, we don't know where on the normal distribution we fall. It could be in the middle, it could be unusually small, unusually large, etc.
 - But we do know where the *most* likely location is: the middle. In other words, our sample mean—assuming the sample isn't biased—is the best estimate of the population.
 - Note the connection between this and the confidence interval: the margin of error revolves around the sample mean.



- A large margin of error and/or a high confidence level means the margin of error is probably in the interval you calculated. In other words, your sample is about spot on.
- But let's go a different direction. The world is in a constant state of change and there are many good reasons that a sample taken from unusual circumstances, such as a subgroup of a population at a different point in time, will not capture the population mean.
 - This is a problem if you're trying to figure out what's going on under normal circumstances but what if you're trying to figure out if something is unusual? That's what hypothesis testing is for.
 - Recall what the null hypothesis is: the usual result. The alternative hypothesis is the unusual result.

- d. If the sample mean is close what we normally see, it stands to reason that we got a different mean because of random chance. But if the sample mean is far from what we normally see, then there's probably something unusual about the sample.
- e. The question becomes, "What's close? What's far away?" That depends on standard deviation and sample size.
 - i. The larger the standard deviation, the more likely your sample mean is different by chance.
 - ii. The larger the sample size, the more likely your sample mean is different because of some interesting reason.
- f. When the calculations are complete, we get a result which, after taking the absolute value, compare to a critical z-score.
 - i. If our result is less than the critical value, we *fail to reject the null hypothesis*. It's not unusual enough to be noteworthy.
 - ii. If our result is greater than the critical value, we *reject the null hypothesis*. This result is probably not due to chance.
- g. It can still be due to chance—95% confidence isn't 100% confidence—but probably not. The phrase "statistically significant" comes from this idea that passing the threshold means the results are interesting.
 - i. It's a controversial mindset. The process states if your result is just below the critical score, nothing interesting is going on. But just above the score is suddenly interesting. The lack of a gray area is frustrating.
 - ii. That said there is still some gray area: samples significant at the 95% level may not be significant at the 99% or 99.9% level.

II. Known σ

- a. If you know the population standard deviation:

$$z_{\bar{x}} = \frac{\bar{x} - \mu_{H_0}}{\sigma/\sqrt{n}}$$

- i. Where $z_{\bar{x}}$ is the z-test statistics;
 - ii. \bar{x} is the sample mean;
 - iii. μ_{H_0} is the mean of the sample distribution, which is assumed to be true for the null hypothesis;
 - iv. σ is the population standard deviation; and
 - v. n is the sample size.
- b. Example: You want to know if your new sports drink improves athletic performance over what people normally do. Suppose the average athlete can run a mile in exactly 15 minutes with a standard

deviation of exactly 2 minutes. You give 100 athletes your sports drink and time their run. Your sample mean is exactly 14 minutes.

- i. Your null hypothesis is that it does nothing to reduce the time it takes to run a mile: $\bar{x} \geq \mu$
- ii. Your alternative hypothesis is that it reduces the time it takes to run a mile: $\bar{x} < \mu$.
- iii. Here's what your equation should look like:

$$z_{\bar{x}} = \frac{\bar{x} - \mu_{H_0}}{\sigma/\sqrt{n}} = \frac{14 - 15}{2/\sqrt{100}} = \frac{-1}{2/10} = \frac{-1}{0.2} = -5.00$$

- iv. Note that because this is a one-tailed test, our critical z-scores are 1.645 (95%); 2.330 (99%); and 3.090 (99.9%). Because the absolute value is greater than 3.090, we have evidence—strong evidence—that our sports drink improves performance.
- c. You might be tempted to say that we've "proved" our sports drink make a positive difference. We haven't "proved" anything. Because there's always a chance of luck, statisticians state that we have evidence for something. "Proof" is not an option.
- d. The example also highlights the difference between statistical significance and practical significance.
 - i. Yes, our sports drink really does improve performance (statistical significance).
 - ii. But is a minute really that big of a deal? If the answer is no, then it's not really interesting and people might not be willing to spend the money on it (practical significance).

III. Unknown σ

- a. You don't know the population standard deviation:

$$t_{\bar{x}} = \frac{\bar{x} - \mu_{H_0}}{s/\sqrt{n}}$$

- i. Where $t_{\bar{x}}$ is the z-test statistics; and
- ii. s is the standard deviation of the sample.

IV. Proportion

- a. You are using a proportion:

$$z = \frac{p - \pi}{\sqrt{\frac{\pi(1 - \pi)}{n}}}$$

- i. Where π is the estimate of the population's proportion; and
- ii. p is the sample estimate.