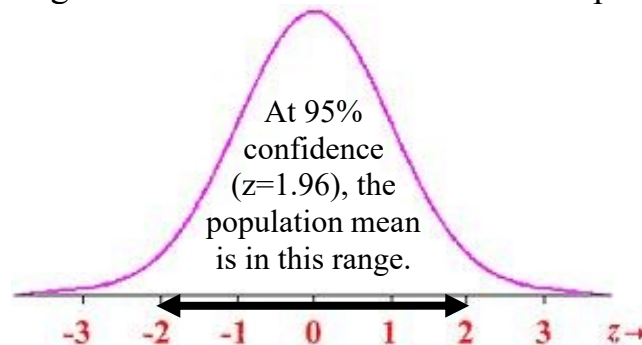


LECTURE 17: HYPOTHESIS TESTING I

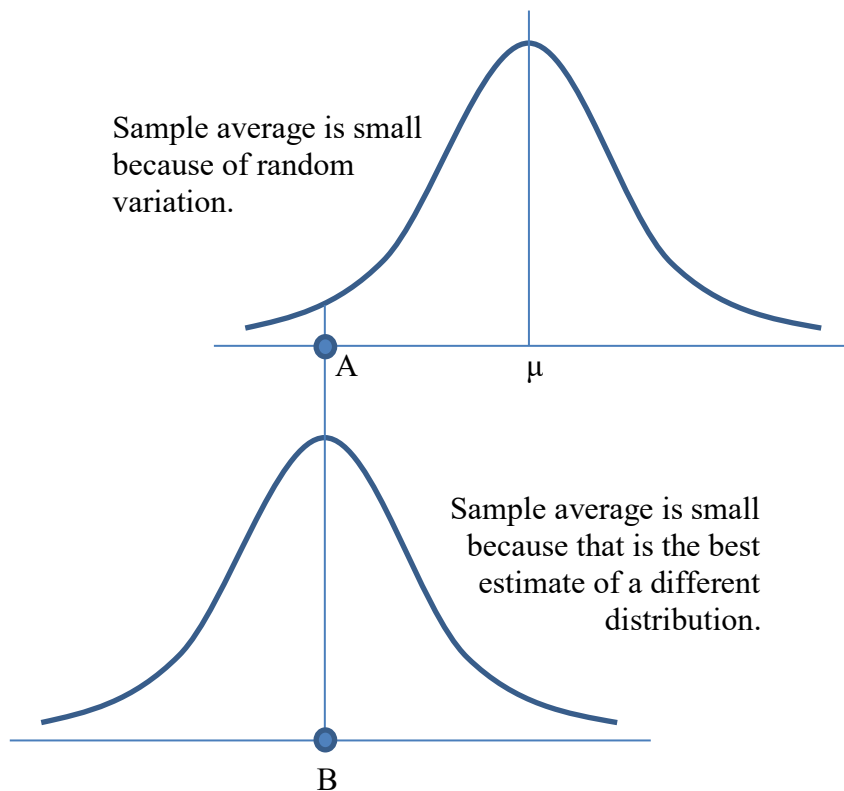
I. The Nature of Hypothesis Testing

- a. 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.
 - i. 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.
 - ii. Note the connection between this and the confidence interval: the margin of error revolves around the sample mean.



- iii. 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.
- b. 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.
 - i. 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.
 - ii. Recall what the null hypothesis is: the usual result. The alternative hypothesis is the unusual result.
- c. If the sample mean is close what we normally see, it stands to reason that we got a different mean because of random chance.

- d. But if the sample mean is far from what we normally see, then there's probably something unusual about the sample. In other words, the sample means is *from a different distribution* than the population mean.
- e. Graphically, hypothesis testing is about answering this question: Is the sample mean from A (unusual due to randomness) or B (unusual due to a genuine difference in populations).



- f. If it's A or if it's B depends on how far away the sample mean is from the population mean, the standard deviation and the sample size.
- i. The bigger the difference between the sample mean and the population mean, more likely your sample mean is different because of some genuine difference.
 - ii. The smaller the standard deviation, the more likely your sample mean is different because of some genuine difference.
 - iii. The larger the sample size, the more likely your sample mean is different because of some genuine difference.
- g. 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. It is "not statistically significant."
 - ii. If our result is greater than the critical value, we *reject the null hypothesis*. This result is probably not due to chance. It is "statistically significant."
- h. 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 it's likely we're estimating a different distribution rather than getting a sample from a strange part of the original (population) distribution.
 - 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:

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

- 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.
 - vi. Note that we take the absolute value; this only for purposes of comparing to critical values.
- 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 fifteen minutes with a standard deviation of exactly two minutes. You give 16 athletes your sports drink and time their run. Your sample mean is exactly 13.75 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:

$$z_{\bar{x}} = \left| \frac{\bar{x} - \mu_{H_0}}{\sigma/\sqrt{n}} \right| = \left| \frac{13.75 - 15}{2/\sqrt{16}} \right| = \left| \frac{-1.25}{2/4} \right| = \left| \frac{-1.25}{0.5} \right| = 2.50$$

- iv. Note that because this is a one-tailed test, our critical z-scores are 1.645 (95%); 2.326 (99%); and 3.090 (99.9%). Because the absolute value of -2.5 is greater than 2.326, we have evidence—strong evidence—that our sports drink improves performance. It is statistically significant at the 99% level. There's a 1% chance (technically, less than a 1% chance) that our difference is merely a coincidence.
 - v. However, the calculated value is not greater than 3.090, our threshold for 99.9% confidence. There is greater than a 0.1% chance that our difference is merely a coincidence.
 - 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. "Prove" 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 the one minute difference 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. Connection to Confidence Intervals
- a. To better understand what we're doing, note there's a connection between confidence intervals and hypothesis testing.
 - b. Consider a confidence interval at 95% confidence. That means there's a 95% chance that the true population mean is in that interval. That also means that there's a 5% chance the population means is above or below that interval.
 - c. Suppose the true population mean is outside the interval. That would be surprising (as it's unlikely; 5% is a small chance) and would suggest that maybe that true population mean isn't really the true population mean, or that the sample mean and the population mean are actually from two different distributions.
 - i. Suppose we're estimating the costs of wind vs solar energy. Solar seems more expensive—the difference averages 5 cents per

kilowatt-hour (meaning solar is somewhat cheaper)—but there's a confidence interval which ranges from -1 cent (wind is slightly cheaper) to 11 cents (solar is *much* cheaper).

- ii. Note this range includes our null hypothesis: zero, or that there's no cost difference between them. In other words, there's less than a 95% chance that the null hypothesis is included in our range and the 5-cent difference is not statistically significant.

d. Here's some algebra to show you the connection:

$$z_{\bar{x}} = \left| \frac{\bar{x} - \mu_{H_0}}{\sigma/\sqrt{n}} \right| = \frac{\bar{x} - \mu_{H_0}}{\sigma/\sqrt{n}} \text{ or } -\frac{\bar{x} - \mu_{H_0}}{\sigma/\sqrt{n}}$$

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

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

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

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

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

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

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

e. Look, it's a confidence interval!