

The Sampling Distribution of the Mean and Hypothesis Testing

Let's first review what we know about sampling distributions of the mean:

1. The mean of the sampling distribution will be equal to μ .
2. The standard error (standard deviation of the sampling distribution) will be less than σ . In fact, the standard error equals:

$$\sigma_{\bar{x}} = \frac{\sigma}{\sqrt{n}}$$

3. As the sample size increases, the sampling distribution of the mean approaches a normal distribution, regardless of the shape of the parent population. You should recognize this as the Central Limit Theorem.

Okay, how is all this information useful? Suppose that someone proposes that a population has parameters $\mu = 100$ and $\sigma = 15$. (You should recognize that this is roughly true of the normal distribution of IQ scores. However, there might be several populations with such parameters.) You take a sample of $n=9$ from the population to see if the proposed parameters are reasonable. Your friend takes a sample of $n=25$ to test the same hypothesis.

Thus, you and your friend have identical hypotheses. That is:

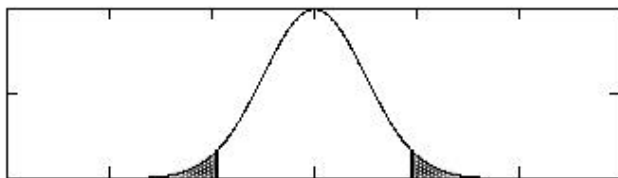
$$H_0: \mu = 100$$

$$H_1: \mu \neq 100$$

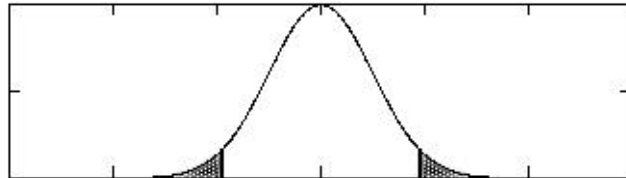
What's different between you and your friend is the sampling distribution from which your samples were drawn. Your sampling distribution has a standard error of 5. Your friend's sampling distribution has a standard error of 3. What are the implications of the differences in standard error?

To see what's different, calculate the scores that would determine the critical regions of the sampling distribution for both you and your friend. Fill in the appropriate values on the figures below. Which of you requires larger scores to be able to reject H_0 ? Why is that so? Which of you is more likely to detect departures from H_0 ?

Your Distribution



Your Friend's Distribution



So, here's the good news about the sampling distribution of the mean. With all that we know about it, we can place some faith in characteristics of the sample means that we draw in research. In other words, if our sample is reasonably large, we can expect that its mean is actually drawn from a sampling distribution of the mean that is normal (regardless of the population's shape) and whose standard error is small enough that our sample mean is likely to be near the μ of the population from which it was drawn. Note that all of this is true in the abstract — when we don't even have a clue about what μ or σ might be. Of course, in reality-land we'll never know what μ or σ actually are!