Demo 17: What Is Standard Error, Really?

The connection between standard error and the sampling distribution of the mean •
How the sample size connects standard deviation and standard error

The difference between standard deviation and standard error is important, and it has everything to do with another concept, *sampling distributions.*

This demo gives you examples of all three of these close together so that you can tease them apart. Here's the main idea: If you take samples of size *n* from a population, and take the mean of each sample, you can construct a sampling distribution—the distribution of those means. That distribution will be roughly normal (see Demo 27, "The Central Limit Theorem," for restrictions). Its mean will be the mean of the population, and its standard deviation will be the standard deviation of the population *divided by the square root of n.*

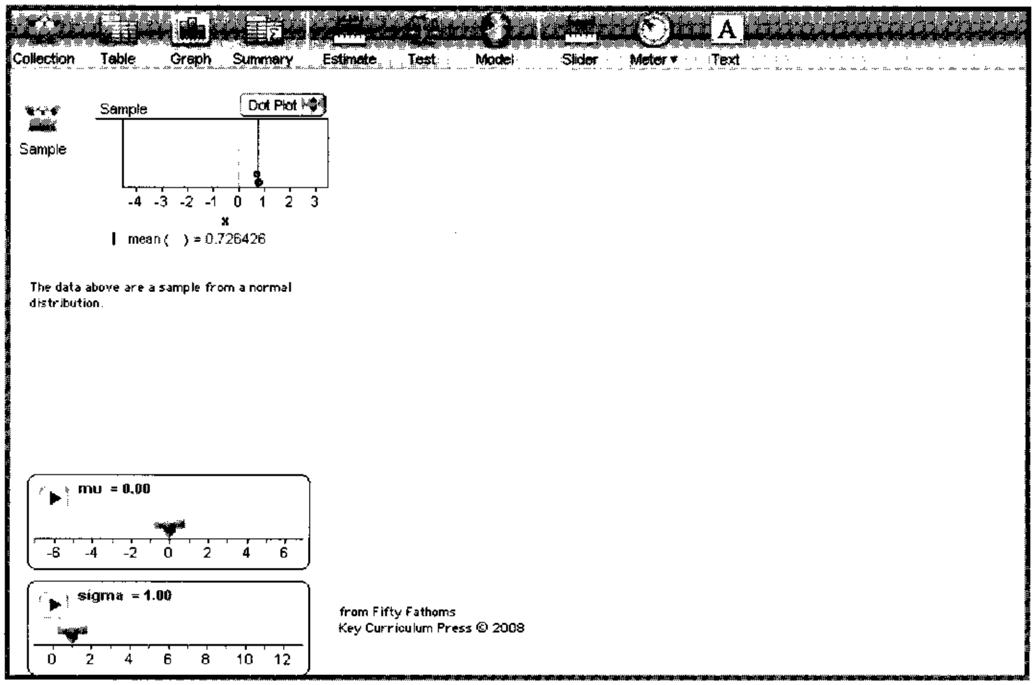

What To Do

▹ Open **What Is SE.ftm.** It will look something like the illustration.

This window shows a collection called **Sample**—the box in the upper left—with three cases. These are drawn from a normal distribution with mean **mu** and standard deviation **sigma,** controlled by the sliders. The little graph shows the values (**x**) of these three points.

▹ **Rerandomize** the sample using the keyboard shortcut: **⌘+Y** on the Mac, **Control+Y** in Windows.

Collecting the Means

It would be great to collect those means.

▹ Choose **Show Hidden Objects** from the **Object** menu.

The newly revealed objects—the collection called **Measures from Sample,** the summary table to its right, and the large graph—will show the results of our resamplings of the **Sample** collection. The summary table will show the numerical mean and standard deviation of those resampled means. That is, **Measures from Sample** is a *sampling distribution.*

▹ Press the **Collect More Measures** button in the measures collection. The **xbar** graph fills with data; each point is a different mean from those three **data** values.

- ▷ Note (in the summary table) that the standard deviation of this distribution of means is probably about 0.60.

Now we are going to investigate the spread of that sampling distribution as we change the number of cases in the **Sample** collection.

- ▷ Click once on the **Sample** collection to select it.
- ▷ Choose **New Cases** from the **Collection** menu. Add one case. You should see the new point appear in the graph.

Note: The right-hand graph *does not change,* because we have not collected the measures. We have to do that ourselves:

- ▷ Again, press the **Collect More Measures** button in the measures collection. The graph updates. Note that the standard deviation of this distribution has changed. Repeat a few times. You should get a standard deviation near 0.50.
- ▷ Now that we see what this is a distribution of, let's change the display. In the pop-up menu in the corner of the graph, choose **Histogram.** The graph will look something like the one in the illustration.

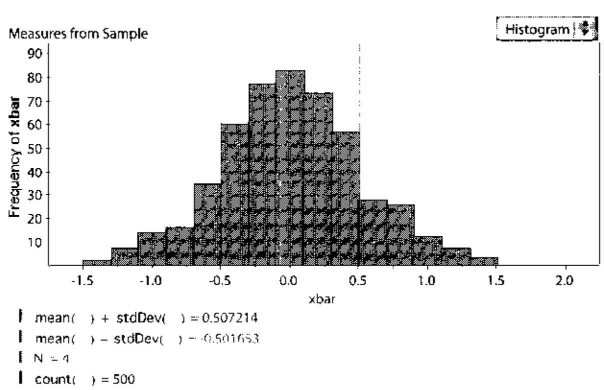

- ▷ Add 21 cases to **Sample** for a total of 25. **Collect More Measures** again. See how the standard deviation decreases.
- ▷ Add 75 cases for a total of 100. Again, **Collect More Measures.**

With $n = 4$ we get a standard deviation of about ½, with 25 we get about ⅕, and with 100 about 1⁄10. One conjecture is that n is the square of the denominator; in other words, the standard deviation of the sampling distribution is the population standard deviation (**sigma**) divided by the square root of n. Let's see:

- ▷ Click once on the summary table to select it.
- ▷ Choose **Add Formula** from the **Summary** menu. The formula editor opens.

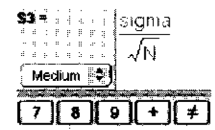

- ▷ Enter **sigma**/$\sqrt{N}$ as shown. (**N** is defined in the measures collection. Use the button for square root.) Press **OK** to close the editor.
- ▷ The newly computed quantity appears in the table. Compare it to the standard deviation above it.
- ▷ See if this near-equality holds true in other cases: Change the number of cases, the population mean **mu**, and the population spread **sigma.**

This quantity—the standard deviation divided by $\sqrt{N}$—is called the *standard error,* and the whole point is that it's the standard deviation of the sampling distribution. That may sound useless, but let's put it another way: The SD is a measure of how far a given data value is likely to be from the mean of the sample. The SE is a measure of how far that mean of the sample is likely to be from the mean of the *population.* That's the SE: the SD of the sampling distribution.

Note: Usually, when we compute standard error, we do so having only the results of a single sample; we divide the *sample* SD by $\sqrt{N}$ to get the SE. Here, we know the population SD, **sigma,** which we never do in real life.

Questions

1 The sampling distribution of the mean looks more or less normal. Is it? How would you know?

2 What standard deviation do you expect in the sampling distribution if there are 16 cases in each sample and the original standard deviation is 1.0? Try it to see if you're right. **Sol**

Challenges

3 With the sampling distribution as a histogram, make Fathom draw the relevant normal curve so that you can compare. (You may want to use a density scale on the graph—in the **Graph** menu, look under **Scale**; you can see what it looks like in the next demo.) **Sol**

4 Suppose we collected 1000 measures instead of only 500. Explain what difference that would make in this demo, and why.

Extensions

- Assess the normality of the sampling distribution. Change the histogram to a normal quantile plot (use the pop-up menu in the corner of the graph) and see if the data are in a straight line.
- Suppose the original data were not drawn from a normal distribution but, for example, from a uniform distribution. Would the sampling distribution of the mean still be normal? Change the formula for **x** to **random( )** and see what changes.
- Suppose the data were binary. Use **randomPick(−1, 1)** for **x** to see what happens.