Demo 11: Flipping Coins—the Law of Large Numbers

How the proportion of heads approaches 0.5 as sample size increases • How the number of heads does not approach half the sample size

This demonstration shows how the proportion of heads you flip approaches the "true" probability of heads. First, we'll look at a file where you flip the coins one at a time.

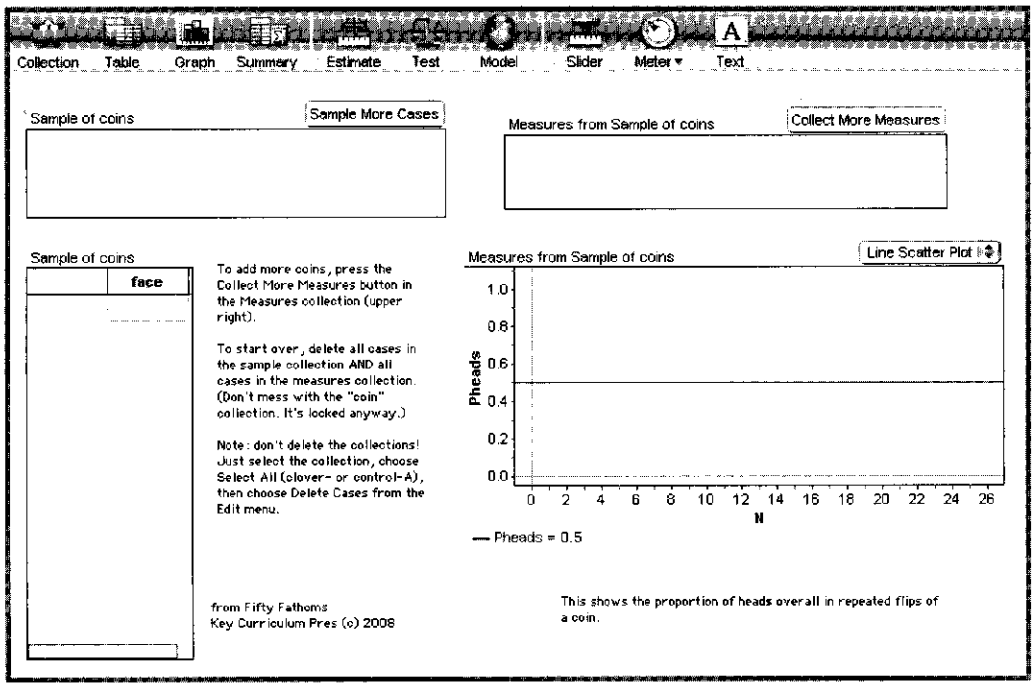

What To Do

▷ Open **Law of Large Numbers.ftm.** It will look like the illustration.

Here we have fairly complicated collections that start out empty. We will mostly watch the graph, which will show the proportion of heads for a number of consecutive flips. But we haven't flipped anything yet, so everything is empty.

▷ Click the **Collect More Measures** button in the measures (right-hand) collection. Fathom will flip a coin. You'll see the coin appear in the sample collection and the cumulative proportion appear in the graph.

▷ Keep doing this and see how the proportion of heads tends toward 0.5.

▷ To start over, either reopen the file or delete all of the cases in the **Sample of coins** collection and all of the cases in the **Measures from Sample of coins** collection. *Do not delete the collections*—just the cases within them! (Note: for reasons that are really esoteric, Fathom's Undo feature does not work if you want to take back collecting measures. So that particular path—ordinarily one of the best—will not serve.)

We have built up the graph conceptually, starting with a single coin. Now, we will use special formulas to get faster performance. In our new file, we will use a collection that does its cases all at once instead of one at a time.

▷ Open **Law of Large Numbers 2.ftm.** It will look something like the next illustration.

Note: To delete all of the cases in a collection, select an object of that collection (the collection, a case table, or even a graph); choose **Select All Cases** from the **Edit** menu; and then choose **Delete Cases** from the **Edit** menu.

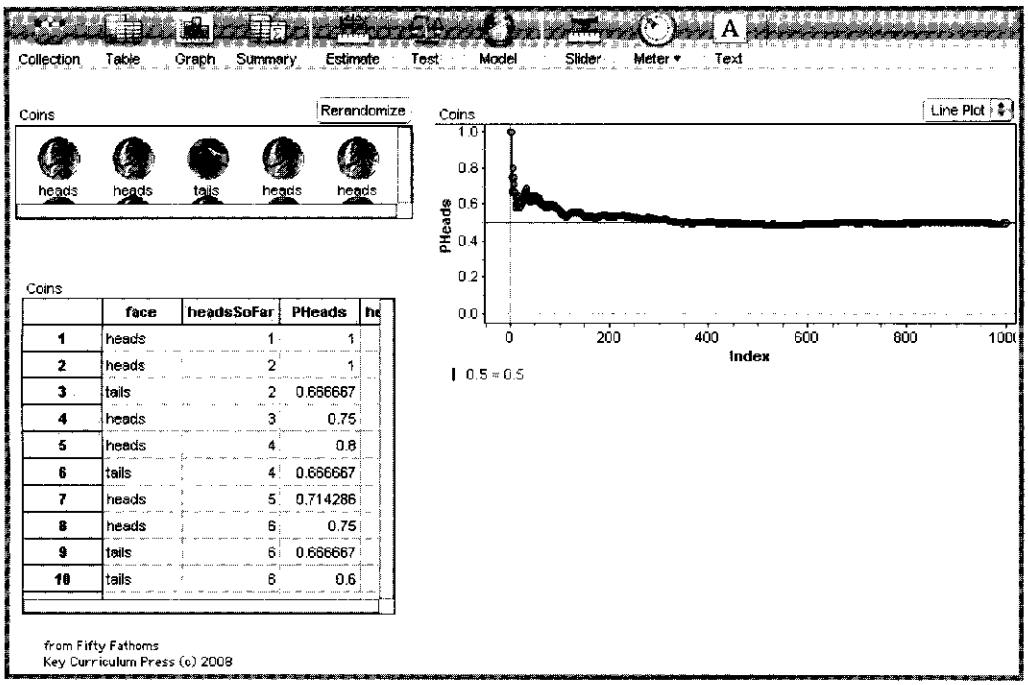

This illustration shows how the proportion of heads (**PHeads**) changed over the course of 1000 flips. In the upper left, you see the collection itself and the first five coins. For example, here, you can see how **PHeads** is 1 after one and two coins (they're both heads) and 0.667 after six coins. Below the collection, the case table shows three attributes for the first 10 cases: the **face** of the coin, the number of **headsSoFar**, and **PHeads**, the proportion of heads in the collection so far.

- ▷ Click **Rerandomize** repeatedly to make Fathom flip the coins and redraw the graph.
- ▷ On the graph, zoom in to the first 50 or so cases and rerandomize a lot. A good way to do this, besides dragging the big numbers off the edge, is to hold down **Option** (Mac) or **Control** (Windows) and click near the zero in the horizontal axis. The axis will expand; with so many points, it may be a little slow.
- ▷ Zoom in to the *last* 100 or so cases—be sure to expand the vertical scale—and rerandomize a lot.

Questions

1. About what is the range of values you typically get for the proportion at 30 cases?
2. What is the range after 1000 cases?
3. Why is the graph so much flatter at the end than it is at the beginning? **Sol**

Extension

- ▷ Restore the scale of the graph by rechoosing **Line Plot** from the menu in the corner of the graph.
- ▷ Make a new graph by dragging one off the shelf. Make it about the same size and shape as the existing graph, if you can.
- ▷ Double-click the collection (the box with the coins in it) to open its inspector. Note: Do your double-clicking down where the coins are—you can even double-click a coin—rather than up between the name and the **Rerandomize** button. That zone "between the ears" is not really part of the collection.
- ▷ Click on the **Cases** tab to make the **Cases** panel appear.
- ▷ Drag the last attribute, **headsExcess**, to the *vertical* axis of the new graph. It will make a dot plot.
- ▷ Close the inspector.
- ▷ Change the graph to a line plot by choosing **Line Plot** from the pop-up menu in the graph itself. Your screen will look something like the illustration below.

This new attribute, **headsExcess**, is how many extra heads you have flipped so far, as compared with exactly one-half of the coins. So, for example, if the first coin is tails, you have −0.5 excess heads: You're "expecting" one-half of a head, and you have none. So your "excess" is −0.5. You can expand the case table to see this attribute as well—look at it until you understand what it does.

- ▷ As before, repeatedly rerandomize, zooming in to various parts of the graph. You'll discover that you need vertical bounds running from about −40 to +40 to get all the points in reliably. Compare this graph with the one for **PHeads.**

More Questions

4 What range of values do you get for **headsExcess** after 1000 cases?

5 What is the range after 30 cases?

6 Why is this graph no flatter at the end than it is at the beginning (in contrast to the **PHeads** graph)?