Beyond the Formula VIII
August 5-6, 2004

Using Fathom Dynamic Statistics Software
to explore
Sampling Distributions,
the Central Limit Theorem,
Confidence Intervals
and
the Robustness of the t-procedure

Presented by
Ruth Carver
Germantown Academy
rcarver@germantownacademy.org

Coin Ages Activity – Part I

Overview: We will simulate in Fathom what we did in the Coin Ages Activity.

1. We will take a sample of size n=5 (with replacement), calculate the mean age of the sample, plot the mean and repeat this process 100 times.
2. We will then repeat this process for samples of size n=10, n=25 and n=50. We will use a slider to control the sample size.

Fathom Instructions:

1. Open the file: Coin Ages 03.ftm.
2. Drag the age attribute to the graph in the upper right of your screen and drop it on the x-axis.

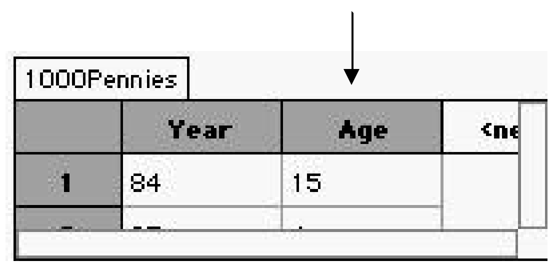

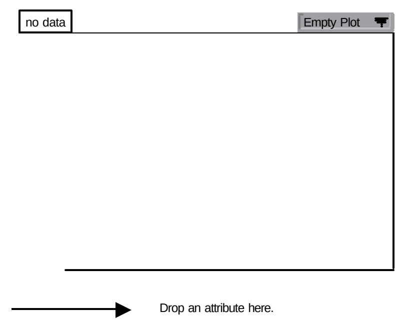

3. Change the graph to a histogram by clicking on the triangle in the upper right hand corner of the graph window.
4. The mean age of the pennies is 12.264 and the standard deviation is 9.913. Sketch a graph of the histogram below. (Put your mouse over a bar in the histogram until it changes to a hand. The status bar in the lower left hand corner of your screen will display the interval and the number of cases in the interval).

5. Would a Normal distribution be an appropriate model for the ages of coins in circulation graph? Justify your answer.

6. The sample size is set to n=5. Click on the button, *Sample More Cases* to take a sample of 5 pennies.

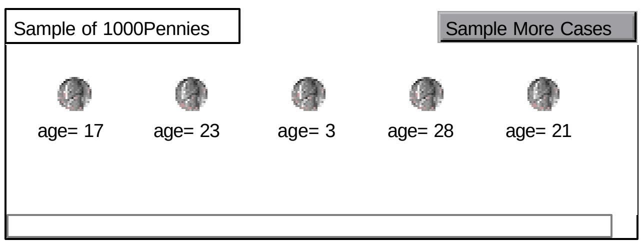

7. As we did in our hands-on activity, we want to calculate the mean of our sample of 5 pennies and graph this mean on a dotplot. We will then replace our sample and take another random sample of 5 pennies, calculate the mean, graph the value on our dotplot, and repeat the process until we have collected 100 means from samples of 5 pennies (taken with replacement) from our collection of pennies. The dotplot will represent the ***sampling distribution of sample means of size n=5 of penny ages in our collection***.

8. Once we click on the button to have *Fathom* carry out the above operations, several things will be happening at once, so it is important to take a moment **beforehand** to make sure that you understand what is happening. Below is a graphic after taking three samples of size n= 5 with replacement. Here are a few things to note as the animation is taking place:

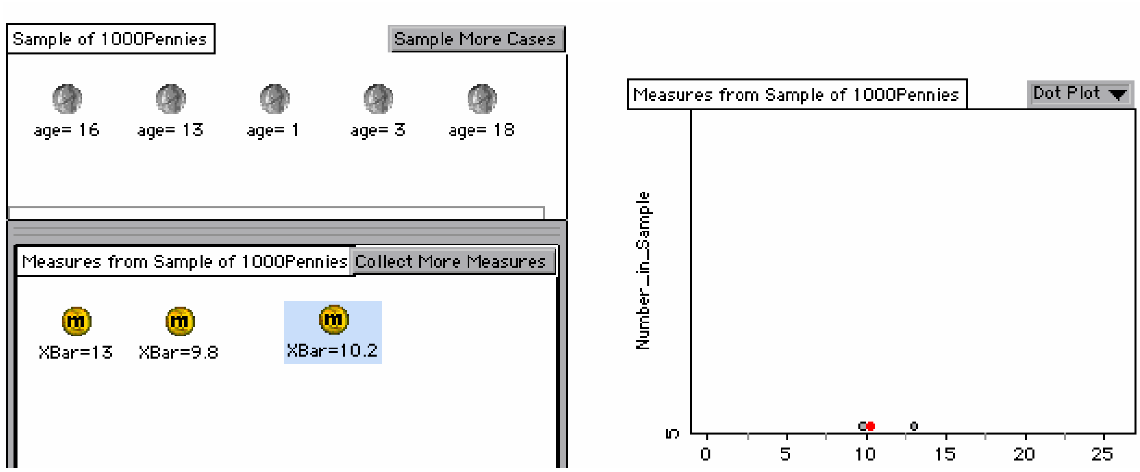

a. The value of each coin in the sample changes to reflect each new random sample of 5 pennies.
b. The mean (x-bar) of each new sample of 5 pennies is recorded in the measures collection below the sample.
c. Each new mean (x-bar) is plotted on the dotplot to the right of the sample.

9. Okay you are ready to go! Click on the Collect More Measures button.

Note: I left the animation feature on so that you could see the sampling distribution being created..

10. The summary table at the bottom of the screen gives the mean and standard deviation for the given sample size. Record the mean and standard deviation for sample means of size n=5 and also describe the shape of the distribution.

11. Now we will have *Fathom* take samples of size n=10, calculate the mean age of each sample and plot the mean. We will take 100 samples of size n=10 (with replacement).

It would probably be a good idea at this point to take the animation off to speed things up. To accomplish this you need to uncheck the animation feature in *both* ***the sample collection and the measures collection.***

Place your mouse anywhere inside the ***sample collection*** and double click to bring up the inspector. Click on the *Sample* tab on the far right and uncheck the box next to *Animation on*.
You **must** also repeat this process with the ***measures collection.*** After opening the inspector for the measures collection, click on the *Collect Measures* tab and uncheck *Animation on*. **Close both inspectors.**

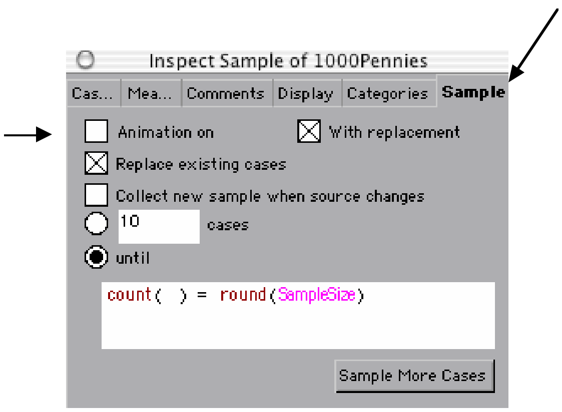

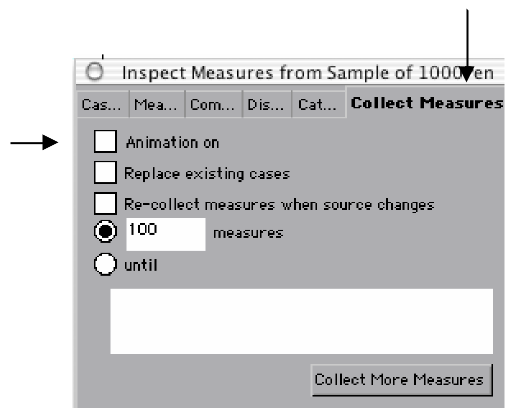

12. To change the sample size to n=10, double click on the number 5 on the **SampleSize slider** or drag the slider to 10 (or a number close to 10 since it is set to round the value). Make sure to <u>hit the return key</u> after changing the value on the slider.

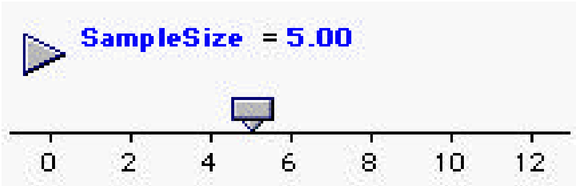

13. Click on the button, *Sample More Cases* to take a sample of 10 pennies.

14. Click on the *Collect More Measures* button again to have *Fathom* take 100 samples of size n=10 and to record the mean of each sample. Record the mean and standard deviation (from the summary table – bottom left) for sample means of size n=10 and also describe the shape of the distribution.

15. Change the SampleSize slider to 25 (don't forget to hit return). Click on the button, *Sample More Cases* to take a sample of 25 pennies then click on the *Collect More Measures* button. Record the mean and standard deviation for sample means of size n=25 and also describe the shape of the distribution.

16. Change the SampleSize slider to 50 and repeat the above steps. Record the mean and standard deviation for sample means of size n=50 and also describe the shape of the distribution.

17. Describe the similarities and differences of the four graphs (n=5, 10, 25 and 50) and the graph of the entire collection of pennies. Make sure to address center, shape and spread.

18. Look up *Central Limit Theorem (CLT)* in your text and write a few sentences about how this activity relates to the CLT.

19. To put a normal density curve on the graphs, first change the graph to a histogram. From the graph menu choose Plot Function. Click on the triangle to the left of *Functions*, then the triangle to the left of *Distributions*, then the triangle to the left of *Normal.* Choose ***normalDensity*** by double clicking on it and type the following in the parentheses: (x_bar, mean(x_bar), s(x_bar)). The expression should look like the window below.

In the **Graph** menu select scale and density to rescale the axes.

20. Fill in the information for the center, spread and shape of the distribution of sample means for the various sample sizes from your simulation in the table below.

Population Mean	Population std dev.	Population Shape
$m = 12.264$	$s = 9.613$	skewed right

Sample Size	**Mean ($\bar{x}$) of Sample Means**	**Std. Dev. (s) of Sample Means**	**Shape of Sample Means**	$\frac{s}{\sqrt{n}}$
1	**12.264**	**9.613**	**Skewed right**	$9.613/\sqrt{1} = 9.613$
5				$9.613/\sqrt{5} = 4.299$
10				$9.613/\sqrt{10} = 3.040$
25				$9.613/\sqrt{25} = 1.923$
50				$9.613/\sqrt{50} = 1.359$

21. The standard deviation of the sample means varies with the size of the sample according to the rule $\frac{s}{\sqrt{n}}$. Are the values in the table for the sample standard deviation calculated by the computer and those calculated using the formula *sample std dev* = $\frac{s}{\sqrt{n}}$ reasonably close enough to justify using it as an estimate?

Coin Ages Investigation – Part II

Exploration of the "Robustness of the t-procedure"

Recall that our population of coin ages for pennies in circulation was extremely skewed to the right with

$$m = 12.264 \quad and \quad s = 9.613$$

The t-procedure is said to be ***robust***. We would like to test this notion. How reliable are the results we obtain when the conditions for the t-procedure are violated? We will start by taking a sample of size n = 5, and using the mean of this sample, construct a 95% confidence interval for the true mean age of the population of pennies. We will repeat this 100 times so that we have 100 confidence intervals. Since we have the good fortune of actually knowing that $m = 12.264$, we can count the number of confidence intervals that "capture" the true mean age of our population of pennies and thus test how well our procedure works.

Fathom Overview and Instructions:

Open the document ***coin ages robust 03.ftm***

To begin the investigation, you will click on the *Collect More Measures* button, BUT before you do that, below is a brief explanation of what's going on.

The vertical line in the diagram represents the Population Mean, $m = 12.264$. If a particular confidence interval "captures" the population mean of 12.26, it is plotted with a green line; if it does not, it is plotted with a red line.

1. Click on the *Collect More Measures* button

to construct 95% confidence intervals for 100 samples of size n=5. How many of the intervals generated captured m, the true population mean? Repeat several times.

2. How does the definition of a 95% confidence interval

> *If you collect many samples of the same size from a population and construct a 95% confidence interval for each, about 95% of the intervals would cover/capture/contain the true population mean, m*

correspond to what you observed (the Summary Table shows the number and proportion of intervals that *captured* the true population mean)?

Investigate:

1. You can change the sample size by dragging on the sampleSize slider

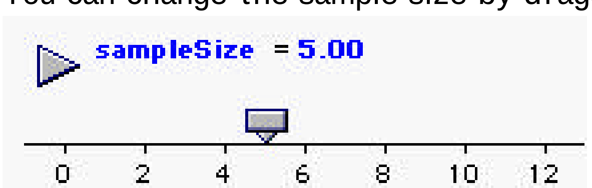

or double clicking on the value, typing a new value and hitting return. To see the effect of your change, click on the *Collect More Measures* button. You might want to do this several times for each investigation.

a) What effect did changing the sample size have on the number of intervals that captured the true population mean? Why?

b) What effect did changing the sample size have on the size of the intervals? Why?

2. In the *Estimate of Sample of 1000 Pennies* box, the confidence level (in blue) can be edited by clicking on it. This brings up the formula editor. Type in a new value and click OK. To see the effect of the change, click on the *Collect More Measures* button.

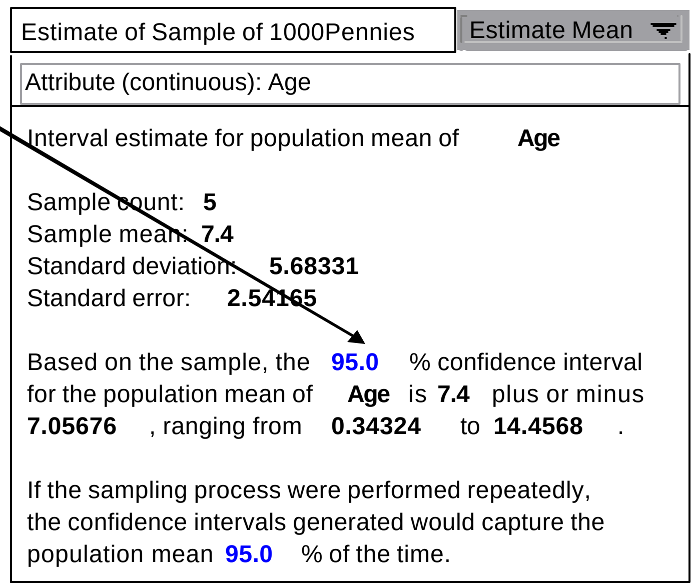

a) What effect did changing the confidence level have on the number of intervals that captured the true population mean? Why?

b) What effect did changing the confidence level have on the size of the intervals? Why?

After completing this investigation, check your understanding of confidence intervals by reviewing the information from the website

Tools for Teaching and Assessing Statistics
http://www.gen.umn.edu/research/stat_tools/

What students should understand about confidence intervals:

* A confidence interval for a population mean is an interval estimate of an unknown population parameter (the mean), based on a random sample from the population

* A confidence interval for a population mean is a set of plausible values of the parameter (μ) that could have generated the observed data as a likely outcome.

* A confidence interval for a population mean consists of a sample statistic ($\bar{x}$) plus or minus a measure of sampling error (which is error from random sampling), when we have approximate normality of the sampling distribution.

* The level of confidence tells the probability the method produced an interval that includes the unknown parameter.

* The probability relates to the method (data, interval), not to the parameter.

* An increase in sample size leads to a decreased interval width: large samples have narrower widths than small samples (all other things being equal).

* Higher confidence levels have wider intervals than lower confidence levels (all other things being equal).

* Narrow widths and high confidence levels are desirable, but these two things affect each other.

* If many random samples are independently sampled from a population and 95% confidence intervals constructed for each one, we would expect about 5% of the intervals to not include the population mean (the population parameter). This 95% refers to the process of taking repeated samples and constructing confidence intervals for each.

* Confidence intervals for a population mean should be based upon a t statistic when the population distribution is approximately normal (or at least not too skewed) and σ is unknown.

* A confidence interval suggests what parameter values are reasonable given the data and all values in the interval are equally plausible as values of μ that could have produced the observed sample mean.

* After you calculate one confidence interval, the parameter is either included or not, but you don't know.

* It is desirable to have a narrow width (for more precise estimates) with a high level of confidence. A narrow width alone is not sufficient (if it has a low level of confidence).

II. What students should be able to DO with this knowledge:

* Know how to make a confidence interval wider or narrower (what factors can be changed)

* Know how to compute a confidence interval for a mean given sample data

* Know how to interpret a confidence interval, make an appropriate inference (in context) and be able to make a correct probability statement as an interpretation of a confidence

III. Some common misconceptions students should NOT have:

* there is a 95% chance the confidence interval includes the sample mean

* there is a 95% chance the population mean will be between the two values (upper and lower limits)

* 95% of the data are included in the interval

* a wider interval means less confidence

* A narrower confidence interval is always better (regardless of confidence level)