USING FATHOM TO PROMOTE INTERACTIVE EXPLORATIONS OF STATISTICAL CONCEPTS ®

Robin H. Lock
St. Lawrence University
USA

Although the use of a statistical computer package has become an integral part of modern statistics courses, the primary goal of traditional software has been to do statistics rather than to learn statistics. Fathom: Dynamic Statistics Software is one of several newly developed packages that focus a greater emphasis on providing an atmosphere in which students can investigate statistical concepts. To facilitate learning, a key premise is that all aspects of an analysis are linked so that students can see how changes in one area are reflected in another. Fathom's developers have made a special effort to produce an intuitive interface that allows students to "drag & drop" to construct analyses from basic building blocks. It also provides a convenient environment for instructors to develop effective demonstrations. We give examples of these sorts of dynamic illustrations and discuss how Fathom can be used to encourage student explorations.

INTRODUCTION

Few would argue with the proposition that statistical computer software has played a significant role in the recent evolution of statistics teaching. As computers become faster, more powerful and more readily available in classroom settings, software to do more sophisticated statistical analyses becomes more accessible to students, even at introductory levels. Hand-held calculators can now do most of the computations and data analysis required for an introductory course. But just doing more complicated procedures more efficiently may not be the best way to utilize computing power in statistics education. Shouldn't we also work to develop statistical software that leverages computing power to help students learn important principles more effectively? Many of our students come to campus with vast experience in video games that simulate real world activities in environments where the users gets to control the action with a mouse or joystick. While we have no illusions that linear regression or hypothesis testing will ever compete with a flight simulator or graphic hand-to-hand combat, we see merit in letting students use statistical software that allows them to manipulate a statistical environment and observe the effects of their own actions.

Fathom: Dynamic Statistics Software, developed by KCP Technologies, is the product of a recent effort to develop new software that is designed specifically to facilitate the teaching of statistical concepts. In what follows we describe several examples of how we have used the software to modify our own teaching. Of course one faces an inherent drawback in trying to highlight the importance of dynamic interactions within the constraints of these static printed pages. We hope that the reader can combine the written descriptions with the graphic snapshots in order to get some sense for what a student would be experiencing when using the software.

SLIDING THE POINT HOME

One of is the convenience with which one may attach a "slider" to almost any numerical quantity in an analysis. As a basic example, suppose that we have a plot of a normal density curve and would like to see how the curve changes as we vary the mean and standard deviation parameters (see Figure 1). Drag the slider icon from Fathom's toolbar to any blank spot on the screen and replace the generic name (V1) with whatever you'd like to call the mean, then create a second slider for the standard deviation. Use those names in the formula for the normal density and that formula (and its plot) will change as one drags the slider back and forth with the mouse. One can adjust the scales for each slider and even click on the triangle to create an animation where the slider steps through successive values of the parameter. The non-dynamic (textbook) alternative is to show static plots of several normal curves with parameters labeled and hope that students can see the connection between the shapes and positions of the curves and those parameters. Allowing students to control those parameters themselves (or even watch a live instructor's demo in class) helps present the relationship much more vividly.

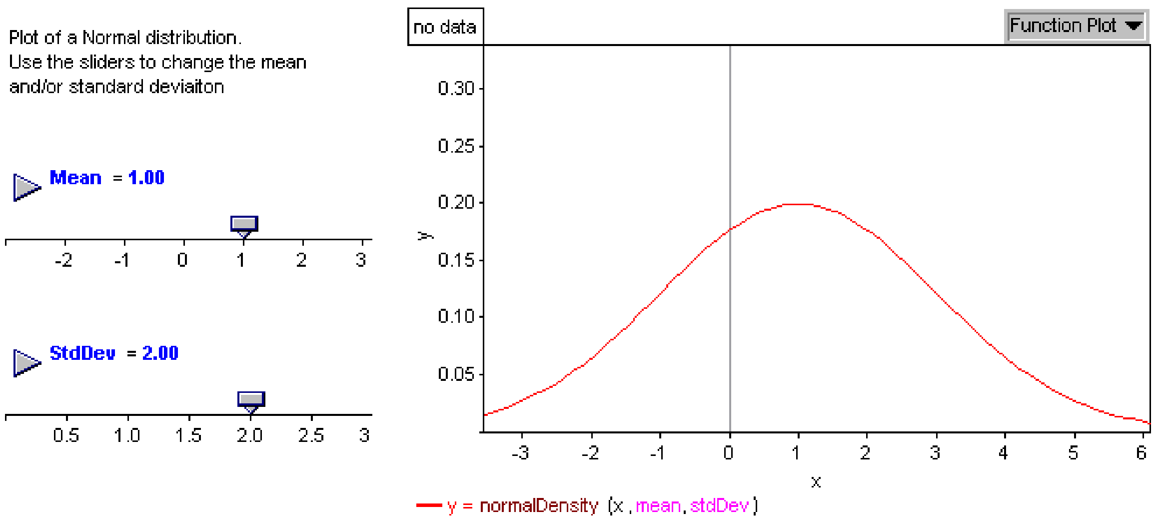

Figure 1. An Illustration of One of Fathom's Exploratory Features.

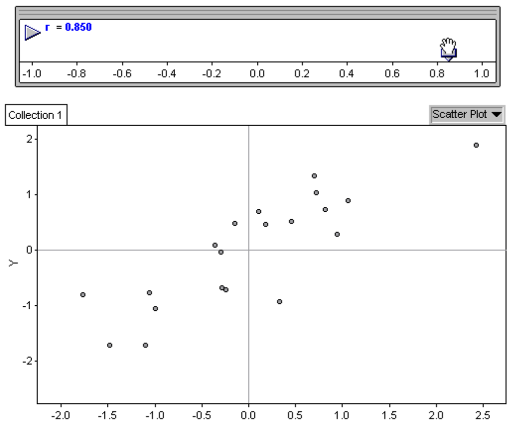

How is the association in a scatterplot connected to the sample correlation? An instructor can prepare a Fathom document that ties the value of the correlation in a slider to the points in a scatterplot so that the points move as a student drags the slider to produce a scatterplot with any desired correlation.

Students, even in a mathematical statistics course, often have difficulty with the concept of power when they first encounter it in studying hypothesis testing. In the past we have attempted to illustrate this idea by drawing a pair of normal curves, representing the sampling distribution of the mean under a null and alternative hypothesis (as shown in Figure 2), and asking students to imagine what would happen as the means moved farther apart or the significance level was decreased or the sample size increased. Wouldn't it be better if they could actually move those parameters themselves and study the effects?

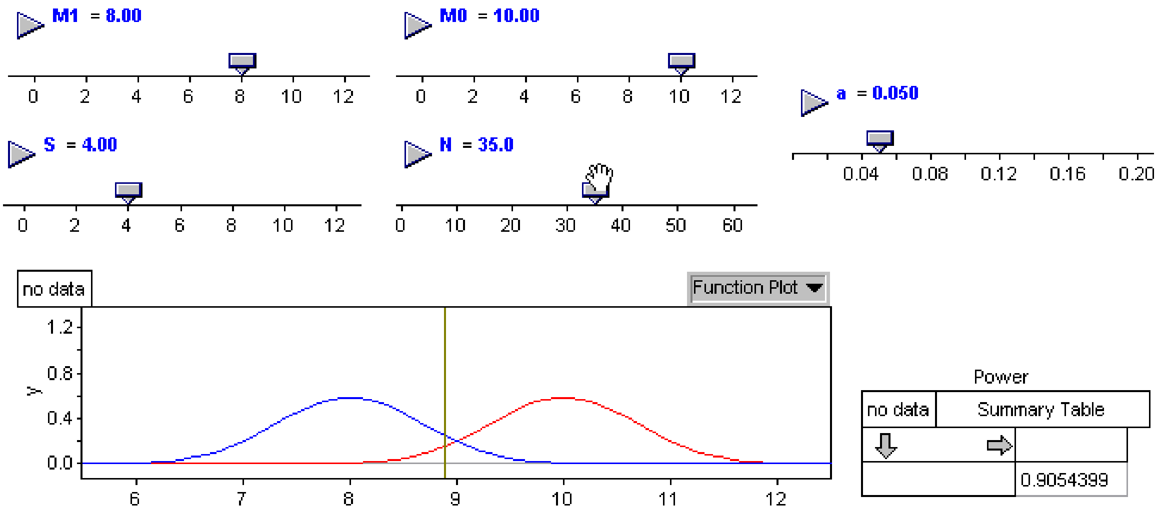

Figure 2. An Example of Fathom Document Representing the Sampling Distribution of the Mean under a Null and Alternative Hypothesis.

The Fathom document illustrated in Figure 2 starts with a test of H_o:μ=10 vs. a lower tail alternative at a 5% significance level. The original sample size is set at 35, the standard deviation is assumed to be 4.0 and the second normal curve represents an alternative hypothesis that the mean is 8. The vertical line represents the boundary for the rejection region as a function of the sample mean. Finally, the summary table contains the result of a calculation for the power of this test against this specific alternative. Students can drag any of the sliders to see how the sampling distributions and the power change. If a student wanted to know how big a sample they would need to obtain 95% power when the alternative was μ=9, a couple of adjustments on the sliders should give a quick answer. While one can also find web-based applets or dedicated software to do this sort of illustration, Fathom provides a convenient platform for creating one in just a few minutes with minimal programming effort and allows the instructor to customize the demonstration to fit a particular instructional agenda.

DRAG YOUR DATA

An effect similar to a slider can be accomplished within any Fathom plot that shows individual data points. One may click on any point in a plot to identify it in the case table (or any other plots that contain information for that data case) and then drag the point to a different location on the screen to change its value. Any calculation that depends on that data value will be updated dynamically as the point is moved.

For example, suppose that we wanted to demonstrate the relative resistance of the median to the presence of outliers when compared to the mean. Choose a reasonably small set of data (perhaps just ten values so the effect is noticeable), put the data in a Fathom collection and generate a dotplot. Use Fathom's "Plot Value" menu item to show both the mean and the median as separate colored lines on the dotplot. Finally, choose any data point and drag it back and forth to see how the mean and median react. A quick drag on the x-axis may be needed to change the scale to allow the point to move farther away from the rest of the data. When students drag such a point themselves, the concept of resistance becomes immediately obvious.

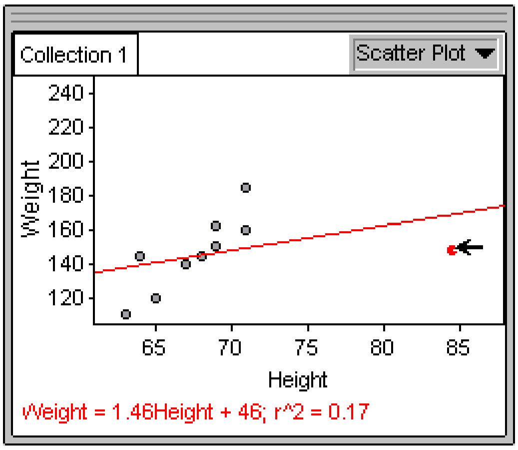

A related demonstration works well to illustrate the idea of an influential point in a regression setting. We'll start with a bivariate data situation - trying to predict people's weight based on their height. Again, a small data set helps to see the desired effect more clearly, so we'll use about ten cases. Drag the height variable to the horizontal axis of a blank plot to create a dotplot and then drag the weight variable to the vertical axis to add the second dimension and get a scatterplot. A Fathom menu item will add the least squares line to this plot. Grab any data value with the mouse and put that individual on super growth hormones to increase both the height and the weight to be much larger than the rest of the data - following a path along the regression line. The equation of the line shouldn't change much, although the displayed r-squared value should increase as you introduce more variability in the weights that the line explains well. Next put the person on a crash diet (by dragging the point straight down), decreasing the weight while leaving the height fixed at a high value and watch how the least squares line reacts - even to the point of turning an original positive association into a negative slope by just moving that single point. Note how the line reacts strongly to the movements of this point when the x-value is extreme, but not nearly so strongly when a point with an average height is moved up and down.

Before using Fathom, we would try to illustrate this point with more static software by running the original regression, plotting the line, then moving to a data table to adjust the influential point, re-running the regression, etc. This process took a great deal longer and students had trouble keeping track of what effect each of the changes had on the least squares line. The dynamic, real-time linkage in Fathom produces a much more compelling demonstration.

COLLECTIONS FROM COLLECTIONS (FROM COLLECTIONS...)

Fathom's basic data objects are called "collections". We can easily create new collections from existing ones by subsetting, sorting, stacking, scrambling, or sampling. We can also collect the values of any statistics (known in Fathom as "measures") into a new collection that records the results as the original parent collection changes. This allows students to simulate the sampling distribution of any statistic of their choice as well as investigate statistical techniques such as bootstrapping and permutation tests with relative ease.

For example, suppose that one wanted to produce a computer simulation of the "random rectangles" activity (as described in Scheaffer, Watkins, Gnanadesikan & Witmer, 1996). This activity has students select a sample of five rectangles from a sheet showing a "population" of 100 rectangles and use the average area of those five rectangles to estimate the average area in the population. When choosing samples "by eye", students tend to avoid the smaller rectangles and produce a distribution of means that is biased above the true mean. When they use a random number table to determine the samples the results will almost always cluster around the true mean. Having done their own samples by hand and viewed the results for the entire class on a plot, students are quite ready to acknowledge the superiority of the simple random sampling scheme. They are also in a great position to start exploring ideas about the sampling distribution of the mean, but often do not have enough data points to see that distribution well or to experiment with how the distribution may change as the sample size changes.

To set up this simulation in Fathom we first start with a collection that contains the areas of the 100 rectangles in the original population. Students can do a summary table for the population to find the "true" mean and standard deviation of the areas and produce a plot to see that the shape is fairly skewed. A selection from Fathom's menus creates a second collection with a sample from this first collection. Students may adjust the sampling size and determine whether or not the sampling should be done with replacement. Once the appropriate sample has been selected, hitting Ctrl-Y replaces the existing sample with a new one. If we had defined a statistic (like the sample mean) for that sample, it will be updated every time a new sample is drawn. But what we'd really like to do in order to mimic the class activity is to keep track of all those sample means as different samples are generated. Fathom's "Collect Measures" option is designed to do just that. This automatically creates a third collection that contains the sample means from repeated samples taken from the original population of rectangles. In just a moment, we can see what the results would look like if we had done the activity with a class of 500 students.

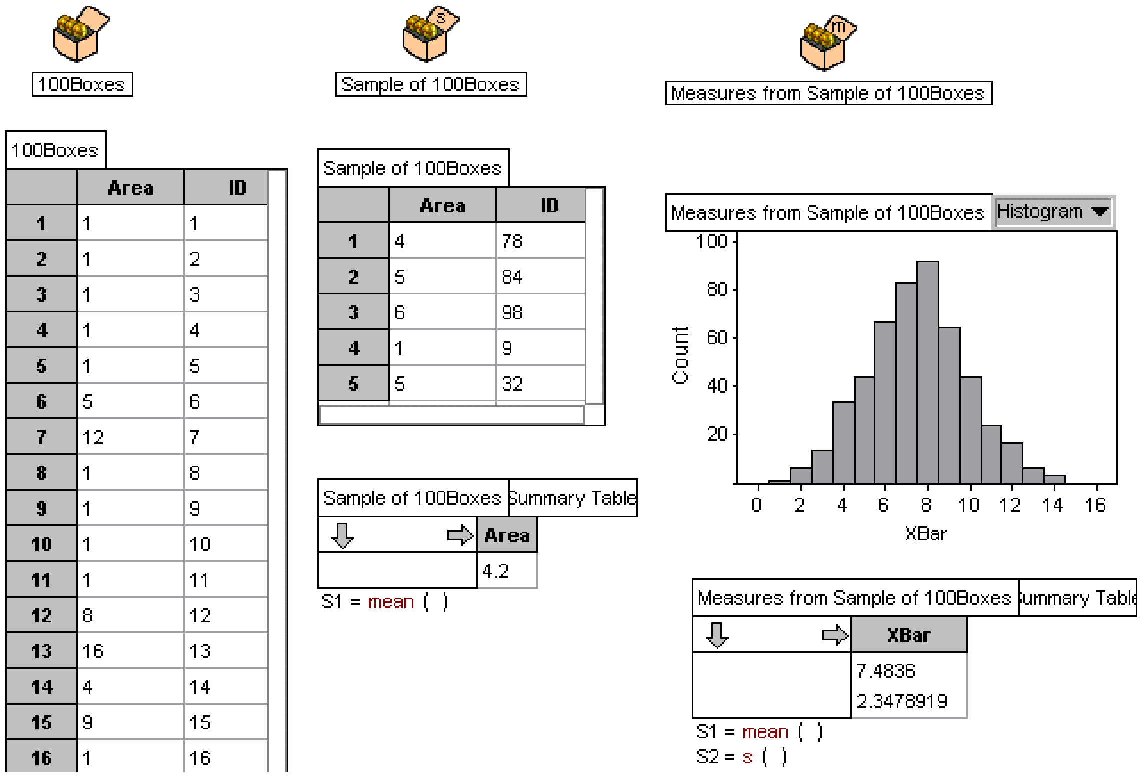

	Area	ID
1	1	1
2	1	2
3	1	3
4	1	4
5	1	5
6	5	6
7	12	7
8	1	8
9	1	9
10	1	10
11	1	11
12	8	12
13	16	13
14	4	14
15	9	15
16	1	16

	Area	ID
1	4	78
2	5	84
3	6	98
4	1	9
5	5	32

Figure 3. Illustration of the End of Fathom Simulation.

The end result of the Fathom simulation (as shown in Figure 3) may look fairly complicated, but, surprisingly, students are able to follow it quite well, especially if they have had some physical experience with the activity before seeing the computer simulation and if they see the simulation being built from the component pieces. Furthermore, students can freely modify any of the relevant parameters. What happens if we adjust the sample size in the middle collection to 20 or 80? What if we change the original population by making it more heavily skewed or adding a couple of really large outliers? What if we wanted to use the sample median or midrange or average of the quartiles as our measure of location? All of these sorts of questions can be addressed readily with a few clicks of the mouse.

As a final example, let us examine the distribution of the sample slope from a simple linear model. We start with a collection of midterm exam scores (out of a possible 50 points) that are assigned (at random) to students in the class. We assume that their final exam scores will be generated according to a linear model `Final=20+1.5*Midterm +Error` where the error terms are normally distributed with mean zero and standard deviation 10. After they have a chance to simulate a set of final exam scores by hand, with each student generating their own random error term, we turn to Fathom to automate the process. Entering the formula `round(20+1.5*Midterm+randomNormal(0,10))` for a new Fathom variable quickly generates a new set of final exam scores that we can plot and add a regression line. With a single keystroke we can re-run the simulation and watch how the plot and fitted line change. We again use Fathom's "Collect Measures" feature to save the sample slopes and intercepts for many iterations in a new collection, display their distribution, and see how that distribution changes as we vary parameters of the model. Although it slows the simulation down, we like to leave Fathom's animation "on" for some of these simulations since students are interested in watching how the plots vary from sample to sample (see Figure 4).

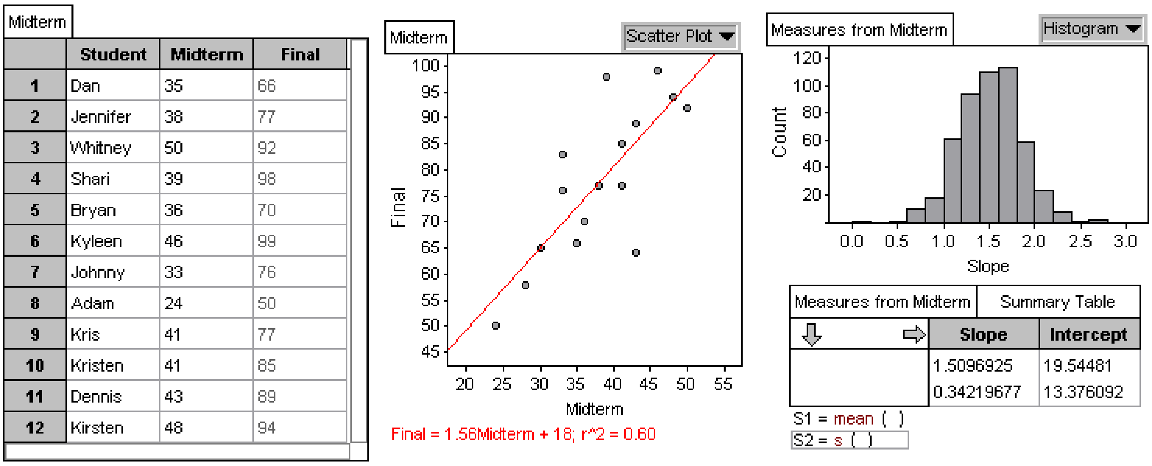

Figure 4. An Illustration of Simple Linear Model.

SUMMARY

The examples given in this paper have illustrated some of the dynamic interactions that computer technology allows our statistics students to experience. While such demonstrations can be created with various software packages (e.g., as Java applets), the Fathom statistical software provides an especially convenient and powerful environment for implementing these ideas.

RESOURCES

More information about the Fathom software can be found at the Key Curriculum Press website at http://www.keypress.com/fathom. Additional examples of guided student activities using Fathom for an introductory statistics course can be found in Rossman, Chance and Lock (2001).

REFERENCES

Rossman, A.J., Chance, B.C., & Lock, R.H. (2001). *Workshop statistics - discovery with data and Fathom*. Key College Publishing.

Scheaffer, R.L., Watkins, A., Gnanadesikan, M., & Witmer, J.A. (1996). *Activity-based statistics - student guide*. New York: Springer-Verlag.