

The Rice Virtual Lab in Statistics

DAVID M. LANE
Rice University, Houston, Texas

The Rice Virtual Laboratory in Statistics is an integrated combination of an electronic textbook, simulations/demonstrations, and case studies. The electronic textbook covers basic concepts in statistics and data analysis and contains links to data analysis tools, instructional demonstrations/simulations, and other on-line texts. The simulations and demonstrations help make abstract concepts concrete and allow students to investigate various aspects of statistical tests and distributions. Case studies demonstrate the real-world applicability of statistical methods. Materials are available on the World-Wide Web.

Teaching statistics is a challenging enterprise for several reasons: Students often find the material in statistics courses abstract and difficult, often do not see the relevance of statistics to their field and, as a result, are less interested in the course than they otherwise would be, and typically lack the background to understand a mathematical treatment of statistics. Perhaps the most difficult challenge is ensuring that students acquire a solid understanding of statistical concepts and how to apply them, rather than merely acquiring a set of rules and procedures to be applied by rote and without appreciation of their implications. Without this understanding, students will have difficulty knowing which technique to apply.

The Rice Virtual Lab in Statistics consists of (1) an electronic statistics textbook, (2) interactive simulations and demonstrations, and (3) case studies. Its goal is to promote hands-on, student-centered learning of statistics and, as a result, to increase students' understanding of concepts in statistics and their motivation for learning statistics. The emphasis is on topics that have important applications in the social sciences, with special attention to techniques pertaining to the analysis of data from experiments.

HYPERSTAT ONLINE: A STATISTICS TEXTBOOK

HyperStat Online is an introductory level statistics book. As can be seen from the table of contents shown in Table 1, HyperStat Online covers basic concepts in statistics and data analysis and emphasizes methods that can be used to analyze data from psychology experiments.

The table of contents for each chapter contains links to the chapter sections and related materials. The contents page for the chapter "Describing Univariate Data" is shown in Figure 1. The left-hand panel of this page contains links to related materials on the World-Wide Web. There

are three types of related materials: analysis tools, instructional demos, and explanatory text. To be included, analysis tools and instructional demos must be available at no cost and must run under the Macintosh, Windows, and UNIX operating systems. Many of the newer tools and demos are written in Java or JavaScript. Others use CGI scripts or are written in XLISP-STAT. Items listed under "Text" are links to portions of other on-line texts that relate to the current chapter. Because of space limitations, Figure 1 shows only 2 of the 13 on-line texts linked to by this chapter.

In addition to links to the chapter sections, each chapter's contents page contains links to a search facility and a glossary. The search facility is written in JavaScript and is based on an index of statistical terms. It does not do a full-text search. The glossary contains 140 entries and includes links to two other glossaries of statistical terms.

Each page contains a relatively small amount of information so that it can load quickly, even with a low-bandwidth connection. Therefore, as in a traditional book, each chapter contains many pages rather than one long scrolling page. Although numerous graphs are included, they are small and do not appreciably increase the time it takes to load a page.

HyperStat Online makes extensive use of hypertext, containing approximately 2,000 links among related concepts. To aid navigation, each page has links to the next and previous pages in the chapter (when there are such pages), the contents page for the chapter, and the main contents page.

SIMULATIONS AND DEMONSTRATIONS

A variety of statistical concepts and principles is illustrated by simulations and demonstrations. Each simulation includes instructions and exercises to help students understand the concept(s) illustrated.

The simulations and demonstrations are Java applets and, therefore, can be used by anyone with a Java-compatible browser. When a browser encounters a Java applet, it downloads the applet and runs it locally. User interface elements, such as buttons, follow the conventions

Partial support for this work was provided by the National Science Foundation's Division of Undergraduate Education through Grant DUE 9751307. Correspondence regarding this paper should be addressed to D. M. Lane, Department of Psychology MS-25, Rice University, 6100 Main Street, Houston, TX 77005 (e-mail: lane@rice.edu).

Table 1
Table of Contents of HyperStat Online

Introduction
Describing Univariate Data
Describing Bivariate Data
Introduction to Probability
Normal Distribution
Sampling Distributions
Point Estimation
Confidence Intervals
The Logic of Hypothesis Testing
Testing Hypotheses with Standard Errors
Power
Introduction to Between-Subjects ANOVA
Factorial Between-Subjects ANOVA
Within-Subjects ANOVA
Prediction
Chi Square
Distribution-Free Tests
Measuring Effect Size

of the local operating system. The figures included here show the way the applets look on a Macintosh; they look slightly different under Windows and UNIX.

Mean and Median

This applet can be used to illustrate several characteristics of the mean and median, including (1) that the mean deviation from the mean is zero, (2) that the sum of squared deviations from the mean is less than the sum of squared deviations from the median (unless the mean and median are equal), and (3) how skew affects the relative size of the mean and median.

Figure 2, the screen when the applet first opens, shows a histogram of nine numbers, their mean, their median, and the skew. The bottom portion of the display shows each of the nine numbers, their differences from the mean, their squared differences from the mean, their differences from the median, and their squared differences from the median. Students can change the distribution with the mouse. When the distribution changes, all the numerical values are updated automatically.

Sampling Distributions

As can be seen in Figure 3, the opening screen of this applet contains a histogram of a normal distribution with a mean of 16 and a standard deviation of 5. Below the histogram are empty histograms represented simply by the x - and y -axes. The best way to begin interacting with the applet is to click the "Animated Sample" button. This causes five scores to be sampled from the distribution one at a time and to float down to form a histogram of the sampled scores. The mean of the sampled scores is computed and floats down to begin a third histogram. The rectangle shown in the lowest histogram of Figure 3 is in motion on its way down to the x -axis. By repeating this

process several times, the student can see concretely that the lowest histogram is the distribution of means computed from samples from the first histogram. Once the student understands this, he or she can approximate the sampling distribution of the mean by clicking on the "5 Samples" button, which adds five sample means to the lowest distribution, and/or either the "500 Samples" or the "2000 Samples" buttons. Figure 4 shows the distribution after 4,000 samples.

It is instructive for students to examine the effect of sample size on the sampling distribution. The third and fourth histograms both allow students to specify the statistic and the sample size. Figure 5 shows a comparison of the sampling distribution of the mean for $N = 5$ with one for $N = 10$ on the basis of 4,000 samples. Notice that the histograms are scaled identically to facilitate the comparison. Students see concretely that the spread of the sampling distribution depends on the sample size. They can also compare the empirical values of the standard deviation with the theoretically expected values. In this case, the expected values for $N = 5$ and $N = 10$ are 2.24 and 1.58, respectively; the empirical estimates are 2.22 and 1.57.

This applet can estimate the sampling distribution of the following statistics: mean, median, standard deviation, variance, mean absolute deviation, and range. Comparing the sampling distributions of two statistics can help students understand the concept of efficiency. Figure 6 compares the sampling distribution of the mean with the sampling distribution of the median, with $N = 5$ and 4,000 samples.

Students can choose the parent population to be normal, uniform, and skewed, or they can customize the distribution, painting it with the mouse. This makes it easy to demonstrate the central limit theorem and concepts such as robustness. For example, Figure 7 shows estimates of the sampling distribution of the mean and the sampling distribution of the median for $N = 10$ for numbers sampled from a uniform distribution. This demonstration shows that the standard error of the mean is smaller than the standard error of the median and that the sampling distribution of the mean is closer to a normal distribution than is the sampling distribution of the median.

Normal Approximation to the Binomial Distribution

This applet illustrates how the normal distribution can be used to approximate the binomial distribution. As is shown in Figure 8, students enter N (the number of trials) and π (the probability of success on a given trial). Binomial probabilities for each possible number of successes are shown by vertical lines, and the normal approximation is shown as a smooth curve superimposed on these binomial probabilities. Students can ask to see binomial and normal approximations of ranges of outcomes. Figure 8 shows that the probability of from six to eight successes out of eight events is .1445 and the normal approximation is .1444. The graph shows that the normal approxima-

Other Sites

Analysis Tools

Many descriptive statistics
and stem and leaf displays
by David Lane

WebStat

by Webster West

Statiscope

by Mikael Bonnier

Interactive histogram

by Webster West

Computing studio

by John Behrens

Mean, median, and sd

by Online Educational
Services

Instructional demos

Histogram explorer
by Eric Scheide

Text

Means, medians, and modes
by Jay Hill

A single variable, descriptive

by Jan de Leeuw

Contents

Main Contents | Search | Glossary

1. Central Tendency

1. Mean
2. Median
3. Mode
4. Trimean
5. Trimmed Mean
6. Summary

2. Spread

1. Range
2. Semi-Interquartile Range
3. Variance
4. Standard Deviation
5. Summary

3. Shape

1. Skew
2. Kurtosis

4. Graphs

1. Frequency Polygons
2. Histograms
3. Stem and Leaf Displays
4. Box Plots



Figure 1. The contents page for chapter 2 of HyperStat Online.

tion is computed as the area greater than or equal to 5.5. A discussion of this *correction for continuity* accompanies the demonstration.

Regression by Eye

The student's task is to examine a scatterplot and draw in an estimate of the regression line. After the line is drawn, the MS_e for the line is displayed, and the student has an opportunity to draw another line. The student can see the best-fitting line by clicking on the "Draw regression line" button. When the "Show minimum MS_e " box is checked, the mean square error associated with the regression line is displayed. Students can compare the mean squared error(s) from their regression line(s) with the one for the

best-fitting line. Figure 9 shows an estimated regression line that has an MS_e of 82 and the best-fitting line that has an MS_e of 69.

To help students learn what scatterplots showing data with various levels of correlation look like, students are asked to estimate the value of Pearson's correlation. After making their choice, the actual value can be shown by clicking on the "Show r " button.

Components of r

Neither Pearson's correlation coefficient nor any other single number can fully represent the relationship between two variables. This applet is designed to help students understand three factors that together determine the value

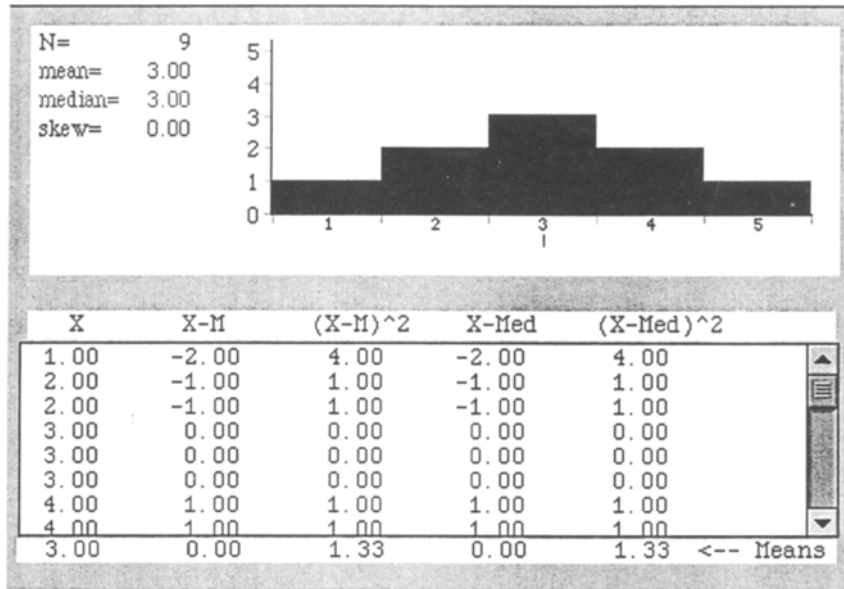



Figure 2. The applet demonstrating characteristics of the mean and median.

of r : the slope, the standard error of the estimate, and the standard deviation of x .

As can be seen in Figure 10, data based on these three input variables are generated and plotted. Students can see concretely that the standard error of the estimate de-


termines the spread around the regression line and, therefore, affects the unexplained variance without affecting the explained variance. They can also see that increasing the slope increases the variance explained without affecting the variance unexplained.

mean= 16.00
 median= 16.00
 sd= 5.00
 skew= 0.00
 kurtosis= 0.00




Clear lower 3
 Normal

N= 5
 mean= 18.42



Animated Sample
 5 Samples
 500 Samples
 2000 Samples



Mean
 N=5
 Fit normal

Figure 3. An example of animated sampling. The rectangle shown in the lowest histogram is in motion on its way down to the x-axis.




Figure 4. An approximation of the sampling distribution of the mean, based on 4,000 samples.

The applet allows the three input parameters to change automatically in small steps. By “animating” the standard error of the estimate, the scatterplot shows the spread around the regression line gradually decreasing until all the points are on the line and then gradually increasing until the spread is quite large.

A Small Effect Size Can Make a Large Difference

Most statistics courses now stress that it is important to consider the practical as well as the statistical significance of an effect. This applet illustrates that an effect that is very small by most conventional assessments can have a large practical effect. Figure 11 shows one such situation. The distribution of the population on the left (which is in blue in the simulation) has a mean score of 50 on some hypothetical employment test; the distribution to the right (which is red in the simulation) has a mean score of 52. The standard deviation of each distribution is 10. If an applicant needs a score of 70 or higher to be qualified, then, assuming the populations are the same size, there will be more than 1.5 qualified applicants from the population shown in red for every qualified applicant from the distribution shown in blue. This large practical effect occurs despite the fact that the difference between populations explains only 1% of the variance in test scores.

Students can explore the effects of the difference between population means and the cutoff score on the rel-

ative number of qualified applicants from the two distributions. They can discover that assessing the practical importance of an effect is tricky and cannot be done without considering the context in which it occurs.

Restricted Range

This applet illustrates the effect of restricting the range of x on Pearson’s r , the slope, and the standard error of the estimate. Figure 12 shows made-up data for which the relationship between x and y is linear and the variance around the regression line is the same for all values of x (homoscedasticity). The upper scatterplot shows the complete dataset. The lower scatterplot shows just the data that fall between the vertical lines in the upper scatterplot. Students can move the vertical lines with the mouse to select the portion of the distribution they wish to examine.

For these data, the slope and the standard error of the estimate are the same in the lower and upper scatterplots. Naturally, Pearson’s r is smaller in the restricted dataset. For the data in Figure 12, Pearson’s r is .60 for the complete dataset but only .39 in the restricted dataset. Figure 12 also shows that $r(c)$, the correlation corrected for range restriction, exactly equals r from the full dataset.

Students can explore the effects of range restriction in “real” datasets for which the assumptions of linearity and homoscedasticity do not necessarily hold. They can examine its effects on the slope, standard error of the es-

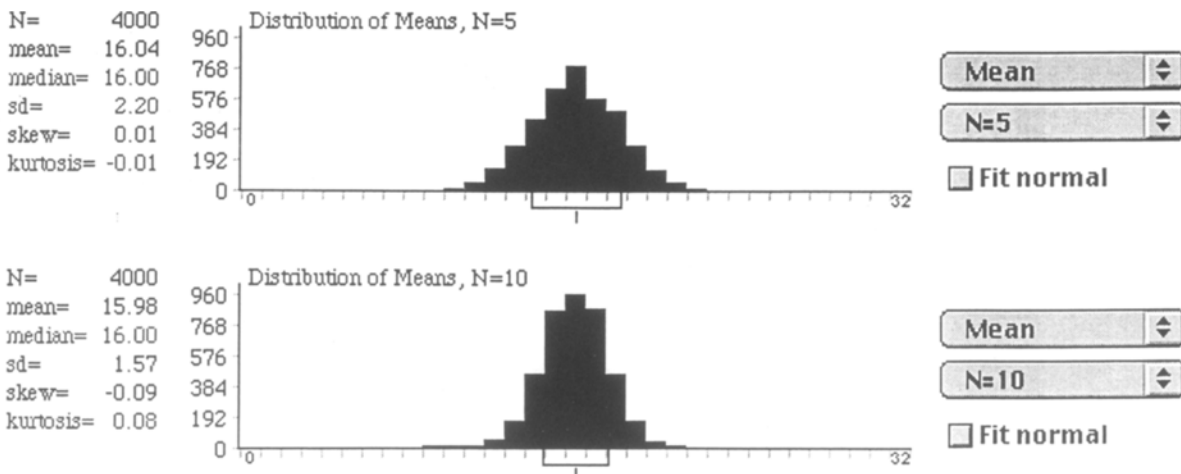


Figure 5. A comparison of approximations of the sampling distribution of the mean for $N = 5$ and $N = 10$, based on 4,000 samples.

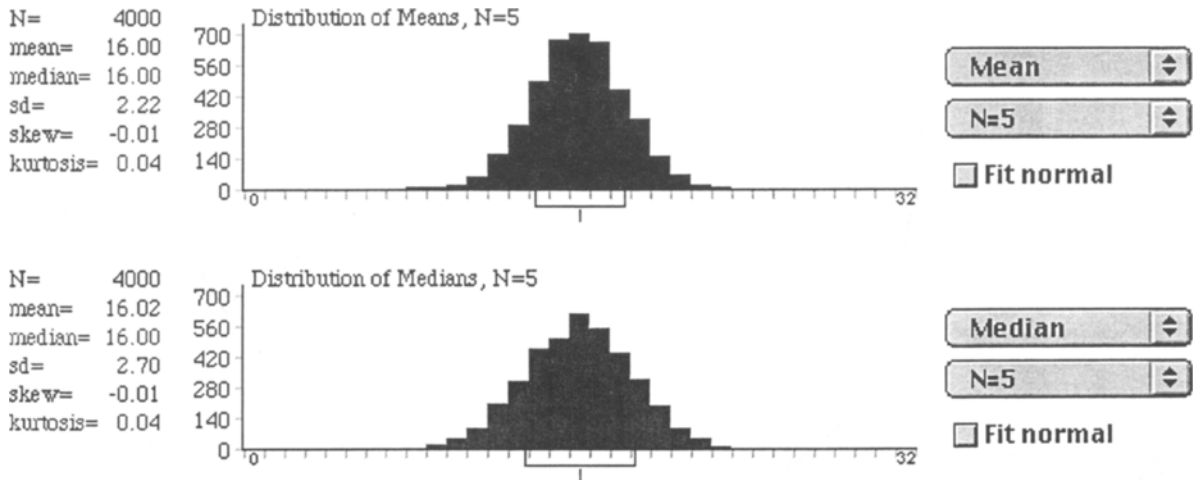


Figure 6. A comparison of approximations of the sampling distribution of the mean and sampling distribution of the median for $N = 5$, based on 4,000 samples.

timate and Pearson's r and see for themselves how well the correction for range restriction works in practice.

Repeated Measures

The applet shown in Figure 13 demonstrates the similarities and differences between independent and related t tests. The student begins by specifying the sample size, the mean for Population A, the mean for Population B, the population value of the correlation between measures (ρ), the standard deviation (which is the same for each population), and whether the design is between- or within-subjects. If the design is between-subjects, ρ is automatically set to 0. The two populations are distributed normally. When the "Sample" button is pressed, one simulated experiment is performed, and the data are displayed. A tally is kept of the number of significant and

nonsignificant outcomes. The tally is reset when any parameters are changed. By pressing the appropriate button, students can perform 500 or 1,000 simulated experiments.

This simulation can be used to illustrate the power difference between independent and related t tests, the role of ρ in determining power, and how ρ affects the error term. The graph of individual data points can be used to show that when ρ is high, the lines tend to be close to parallel. This can lead into a discussion of why the treatments \times subjects interaction is the error term in within-subjects designs.

CASE STUDIES

Case studies can serve several valuable functions in the teaching of statistics. First, case studies have important

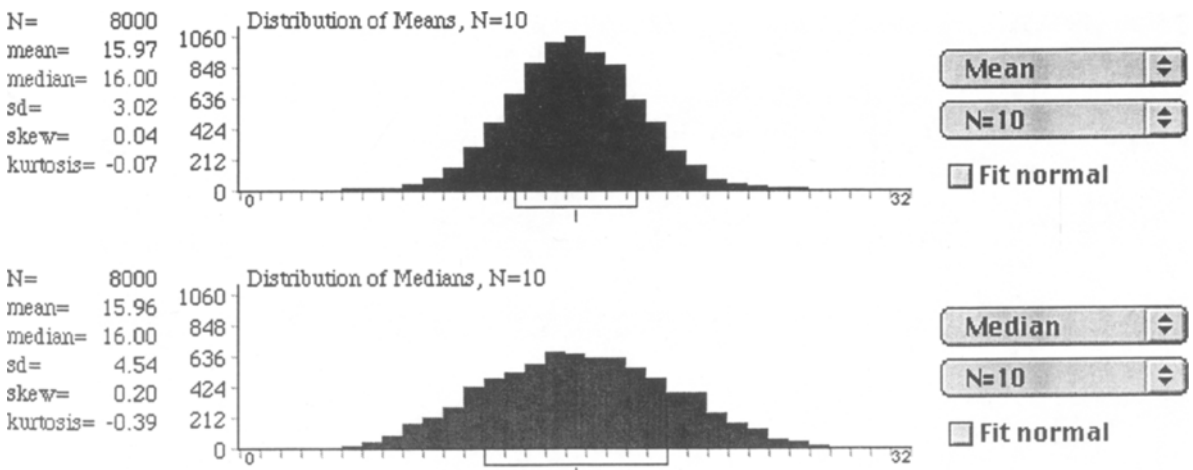


Figure 7. A comparison of approximations of the sampling distribution of the mean and sampling distribution of the median for $N = 10$, on the basis of 8,000 samples. The parent population has a uniform distribution.




Figure 8. The normal approximation of the probability of six or more successes out of eight trials with a .5 probability of success on each trial.

motivational effects, because (1) case studies make clear the applicability to and importance of statistics for real-world problems, (2) students have more interest in the outcome of statistical analyses when they know that it is a real problem being investigated, and (3) students can become actively involved with the data from case studies by exploring hypotheses that interest them. In addition to their motivational benefits, case studies have the important benefit of making issues in statistical analysis more concrete. For example, it is much more instructive for students to see how a logarithmic transformation affects the outcome of a statistical analysis of real data than it is to

read an abstract description of the effects of a logarithmic transformation.

Although several case studies are under development, only the case study “The Effect of Smiles on Leniency” is available as of this writing. The raw data from an article by LaFrance and Hecht (1995) are analyzed with popular statistical packages to illustrate statistical methods, including quantile plots, boxplots, linear contrasts, Dunnett’s test, and the Bonferroni correction. Instructions for doing the analyses and explanations of the output are given. Raw data are available for students to do their own analyses.




Figure 9. The line with the steeper slope was drawn by eye. It has an MS_e of 82. The other line is the best-fitting line and has an MS_e of 69.

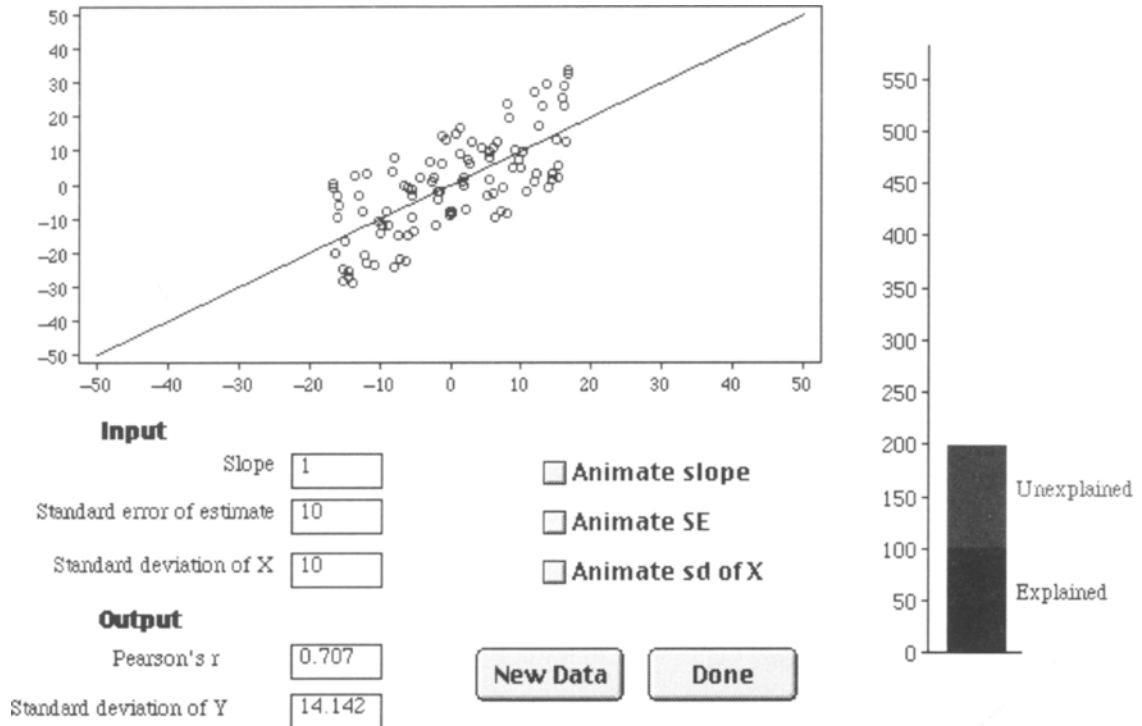


Figure 10. Students can manipulate the input parameters and observe the effects on the graph and on the output parameters.

This study produced the following three results: (1) The mean of the three smiling condition means was significantly greater than the control mean; (2) only one of the three smiling conditions was significantly different from the neutral control; and (3) there were no significant differences among the three smiling conditions. A simulation was developed to help interpret this set of findings. Students specify the population means of the four conditions

used in the experiment and the within-group standard deviation; simulated experiments are then performed to determine the number of times the statistical results matched the actual results described above.

Figure 14 shows 100 simulations of the situation in which the three smiling condition population means are equal and greater than the control condition mean. For these 100 simulated experiments, 88 showed a significant

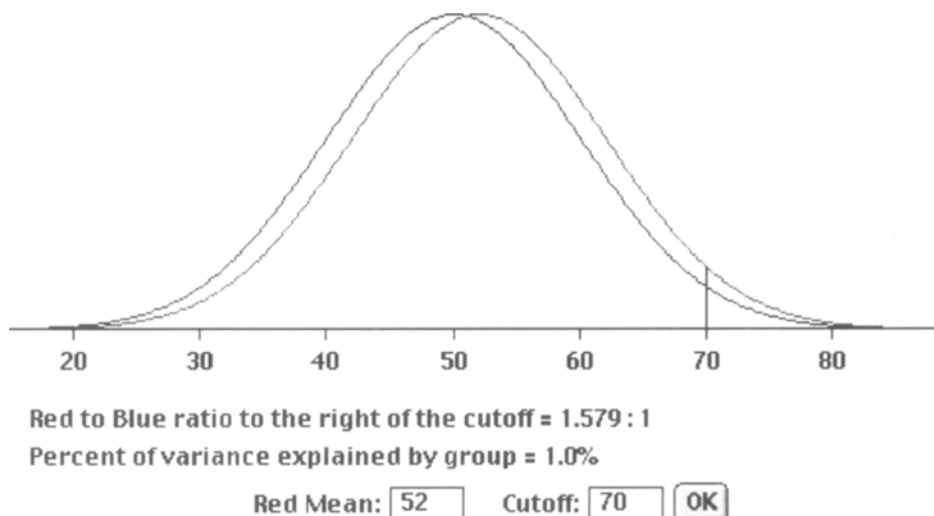


Figure 11. The applet demonstrating that a small effect can make a large difference.

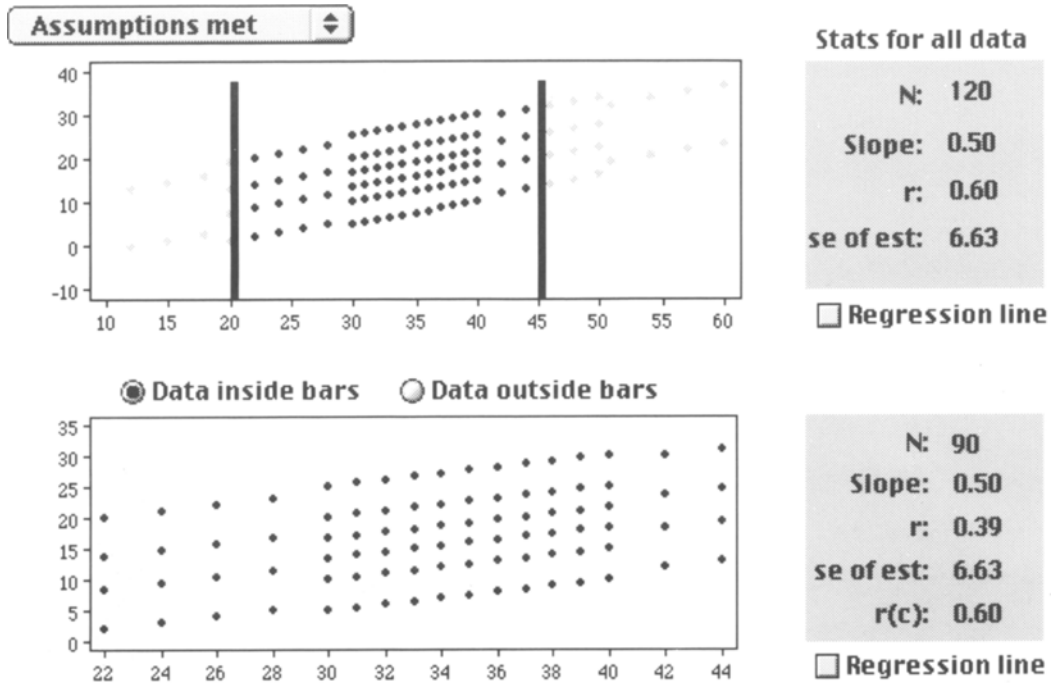


Figure 12. The effects of range restriction. Students move the vertical bars in the upper plot and observe the effects on the lower plot.

difference between the mean of the smiling conditions and the control. Of these 88 experiments, 21 found exactly one experimental mean differing significantly from the control. Finally, 20 of these 21 experiments found no significant differences among the three smile conditions. This

shows that the statistical outcome in the experiment is not particularly unlikely if the three smile population means are equal. Students explore how likely the three experimental outcomes would be if only one of the three population smile means were different from the control con-

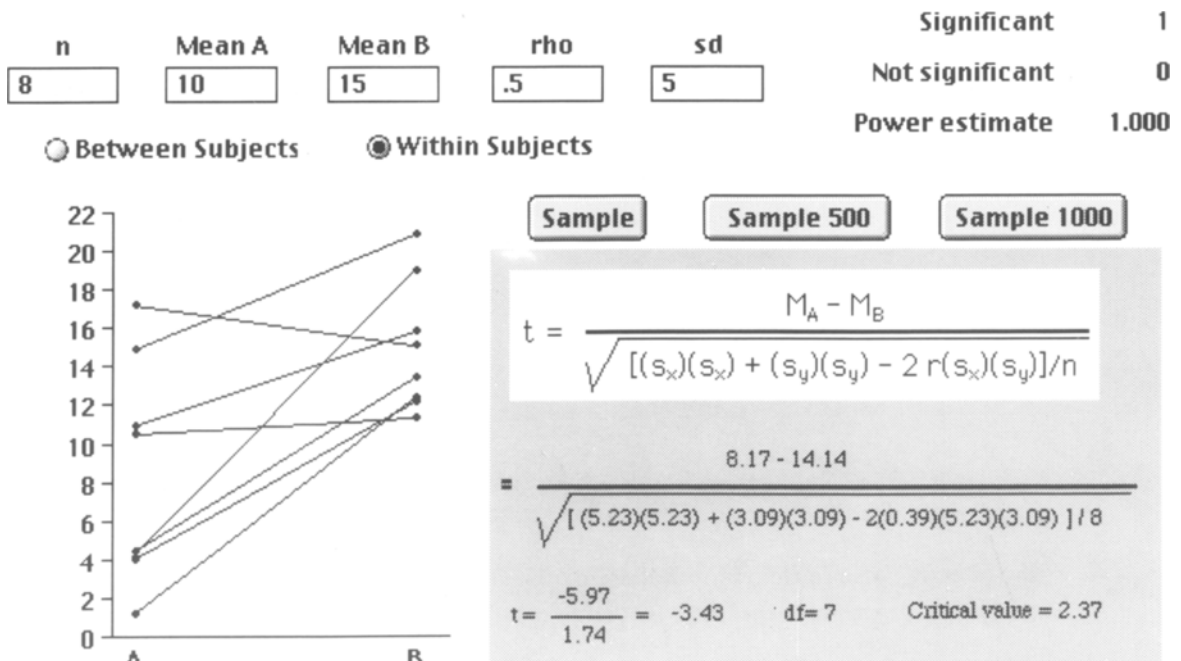


Figure 13. This applet shows the similarities and differences between the independent and correlated t tests.

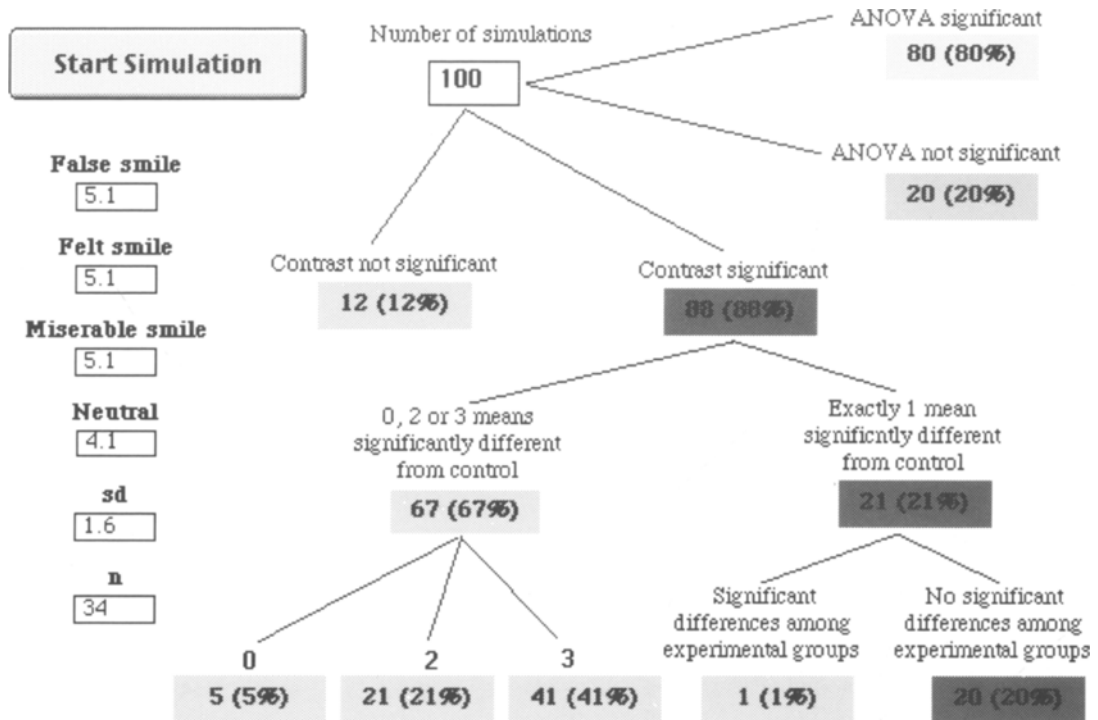


Figure 14. Students specify the population parameters on the left and run simulated experiments.

dition. It turns out that it is very unlikely for all three experimental findings to occur in that situation.

INTEGRATION OF HYPERSTAT, SIMULATIONS/DEMONSTRATIONS, AND CASE STUDIES

The ability to provide links on the World-Wide Web greatly facilitates the integration of three components of the Rice Virtual Lab for Statistics. First, the simulations and demonstrations have links to the explanatory material in HyperStat Online. For example, the sampling distribution simulation provides links to the following sections of HyperStat Online: sampling distribution, normal distribution, histogram, sampling distribution of the mean, standard deviation, mean, median, range, standard error, bias, efficiency, and central limit theorem.

Second, each chapter in HyperStat contains links to relevant simulations and demonstrations. These include simulations and demonstrations in the Rice Virtual Lab for Statistics and those developed by others.

Finally, the case studies and simulations are designed to be tightly integrated. The case study on smiles and le-

niency includes a simulation designed to help interpret the pattern of results. In a case study under development, students will be able to explore the effects of a violation of the assumption of normality apparent in the experimental data by simulating experiments in which the data are sampled from a distribution resembling the distribution found in the experiment.

AVAILABILITY

The URL for the Rice Virtual Lab in Statistics is <http://www.ruf.rice.edu/~lane/rvls.html> and can be freely accessed. Most of the simulations require a browser that supports Java version 1.1.

REFERENCES

LaFrance, M., & Hecht, M. A. (1995). Why smiles generate leniency. *Personality & Social Psychology Bulletin*, 21, 207-214.

(Manuscript received November 4, 1997; revision accepted for publication May 28, 1998.)