

HOW JMP IS MOVING THE FENCES IN UNDERGRADUATE STATISTICS EDUCATION

Robert H. Carver, Stonehill College, Easton MA 02357, (508-565-1130), rcarver@stonehill.edu

ABSTRACT

As powerful software and large real-world data sets have entered the mainstream college classroom, and statistics has entered the K-12 classroom, university faculty have been slowly modifying "introductory stats" courses to emphasize statistical reasoning and discovery. JMP has clearly played a role in facilitating and adapting with these modifications. This talk reviews nine specific ways that JMP can be instrumental in redefining what we regard as fundamental and introductory topics, and argues that the structure and approach of JMP can shift traditional boundaries between introductory and more advanced topics. The software's native features – such as Graph Builder's visualizations, simulation, seamless integration of parametric and non-parametric tests, painless logistic regression, bootstrap CIs, sampling weights and intuitive data management tools – allow us to re-imagine the intro course, inspire undergrads to pursue further study in analytics, and provide value to analytically oriented firms. The presentation will include some classroom-tested demonstrations of "advanced" topics, couched in a way that introductory students can grasp them. Additionally audience members will be invited to share their insights about skills, concepts and orientations that they seek in entry-level employees fresh out of college.

INTRODUCTION

In most professional sports, the dimensions of the playing area are fixed by custom and convention. The lines defining what is “in” and what is “out” of bounds are the same regardless of arena or stadium. In baseball we have the somewhat remarkable situation that allows for variation in the precise angles and location of outfield fences, so that there are local differences in both the heights and distances of the fences. The consequence is that a long fly ball might be a routine out in one stadium, a home run in another, or a single that ricochets off an outfield wall. From time to time, ballpark owners may elect to modify their fences to tip the balance and adjust the number of homeruns hit in the park.

In this paper, I carry the metaphor of moving the fences into the realm of the collegiate statistics curriculum. The triangular relationship of computing technology, real-world statistical applications, and statistics education is a dynamic one that calls upon us frequently to re-evaluate the scope, sequence, and content of the introductory statistics course. Each new generation of hardware and each new iteration of statistical software offer alternatives to traditional methods of statistical practice and education [1], [10], [20]. Decisions about which application areas are essential vary by client discipline and over time as well. Whereas at one time, statistics courses were viewed as a narrow specialty, enrollments have grown steadily as universities have almost universally required courses to promote quantitative literacy and statistical reasoning. Recently, across a wide variety of public and private sector disciplines, “Big Data” and “analytics” are being applied to numerous purposes [7], [16].

Within the statistics education community, calls for reform and adjustment to changing realities are a long-standing tradition [10], [20]. In recent years there has been growing attention given to the importance of using real data and taking full advantage of technology, culminating in the 2005 GAISE report [1], [5]. More recently still we find arguments that, in response to the big data explosion, further changes are needed, especially a shift to visualization and to randomization-based inference [2], [6], [8], [9]. As on-line sources of real-world data continue to become more extensive and easier to access freely, textbooks and courses continue to replace artificial small data sets with larger real data files. Among the many “new” sources of reliable real raw data are large-sample surveys from governmental and non-for-profit organizations (see, for instance, [4], [21], [23]). The combination of user-friendly software, fast inexpensive computers, and widely-available datasets create numerous opportunities to rethink which topics “make the cut” for the introductory statistics course. Simultaneously, secondary school curricula now routinely include statistics and probability concepts and techniques. Some topics which were formerly in the canon are no longer pertinent, and others that once were considered too advanced for college sophomores may now have a place. Hence there is ample reason and opportunity to reevaluate yet again. In particular, in the following discussion I suggest nine ways that JMP can help us to adapt to changing realities and to seize opportunities to elevate statistics education. Each of the nine is present within the new edition of my book, recently published by SAS Press. [3]

9 WAYS THAT JMP ENABLES FENCE-MOVING

My home institution has a site license for JMP11 Pro, allowing anyone on campus to download the software, data tables, and scripts for use on a personal computer on or off campus. The software is not restricted to the classroom or lab, and runs in both the Windows and MacOS environments. The universal access and JMP's features have changed what I teach and how I teach in at least these nine fence-moving ways.

- Take a Look at This!
 - Engaging interactive visualizations via Graph Builder, 3-D scatterplots, and Bubble plots
 - Embedded and add-in simulations that illustrate important concepts
- Keeping it Real
 - Creation and use of sampling weights
 - Intuitive data management tools
- Accessible introduction to heretofore “advanced” techniques
 - Seamless integration of parametric and non-parametric tests
 - Bootstrap Confidence Intervals
 - Painless logistic regression
- Using Authentic Professional Tools
 - Outstanding support for undergraduates – menu-driven, ample documentation, help, on-screen tips, etc.
 - Removal of the DO NOT ENTER gates by eliminating or lowering barriers to entry (and providing good reasons to stay)

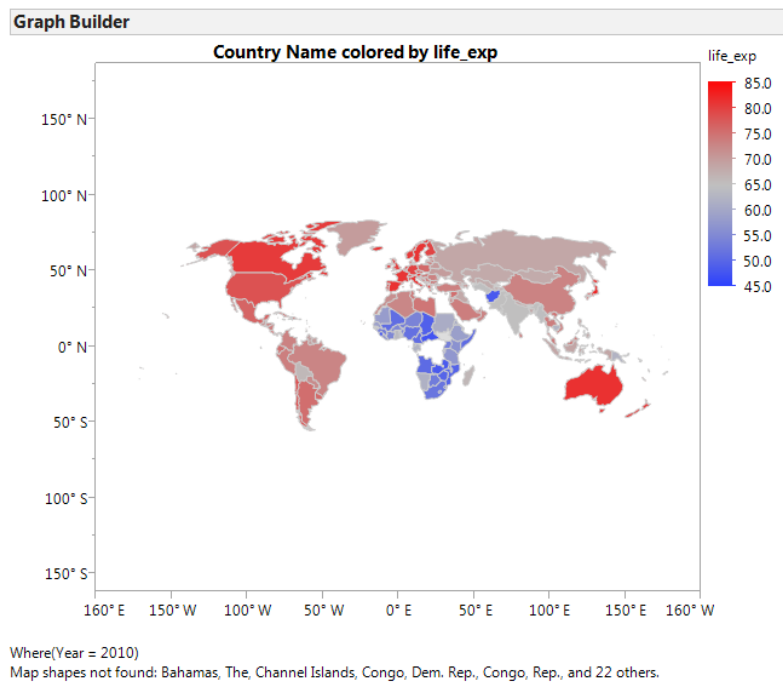
TAKE A LOOK AT THIS!

Engaging interactive visualizations via Graph Builder, 3-D scatterplots, and Bubble plots

Before graphs were called “visualizations” JMP was incorporating visual analysis with computational results, and linking open graphs from a single data table. For working statisticians, the benefits are clear. But for learners the power of JMP’s visual approach goes well beyond discovering patterns in data. Standard JMP graphs help learners build understanding of important statistical *concepts* and to experience the process and the thrill of statistical reasoning. Interactive and animated graphics go a step further by having a learner/user engage and control the software with almost no prior experience.

On the first day of a freshman business statistics class, I show Hans Rosling’s short video on the wealth and health of nations, featuring an animated bubble plot with national life expectancy on the Y-axis and per capita GDP on a log-scaled X-axis [19]. Students are immediately drawn in by the animation. They understand quickly what the display communicates, and react audibly as one small blue bubble plummets and then rebounds on the vertical life expectancy axis in the 1990s. When they discover that the distinctive action reflects the Rwandan genocide, they are silenced—and fully engaged in thinking about the data and its compelling story.

Soon thereafter, we explore the same data in JMP, using Graph Builder and the Bubble Plot tool. At that point, they are controlling the graph, wandering off into their own explorations and *not thinking* about the software, but rather thinking about substantive questions and examination of evidence. These are college freshmen.

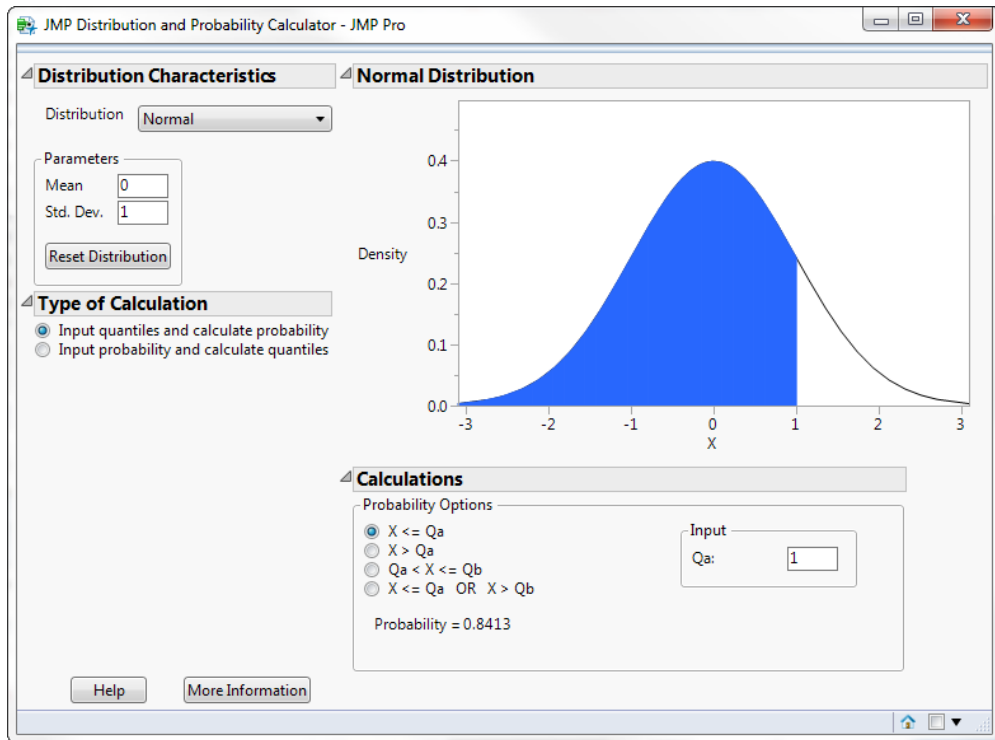




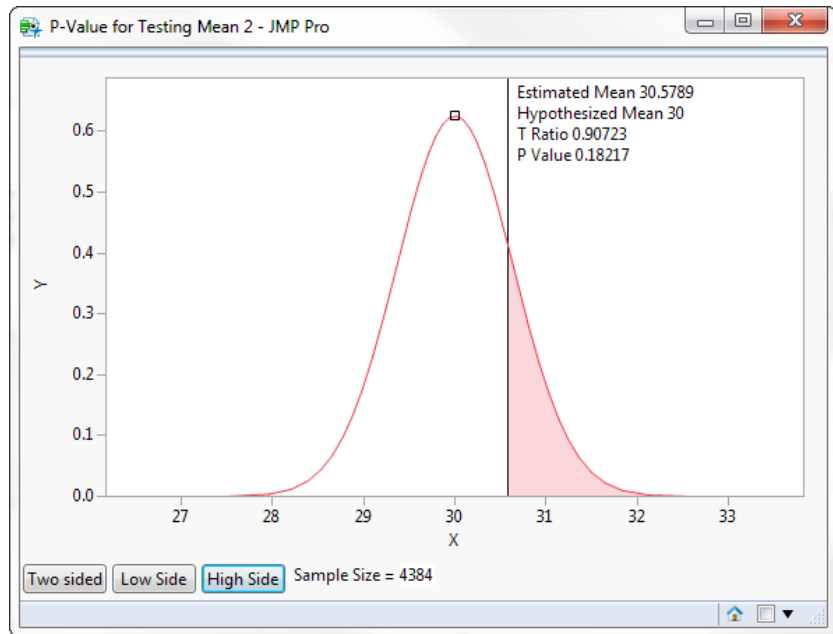
Embedded and add-in simulations that illustrate important concepts

Similar to the insight-developing properties of data visualizations, JMP also includes numerous simulations and animations that can clarify subtle foundational concepts within the introductory statistics course. There are many applets and standalone tools available online that accomplish the same goals, but as students are moving along the JMP learning curve, it is advantageous to have these tools incorporated within the single software environment, relying (often) on the very same data that students are already working with.

For example, the add-on distribution calculators—which are also interactive and animated—convey probability concepts very directly get to the heart of the conceptual understanding of probability and density functions, allowing students to develop facility with finding probabilities or quantiles. Along the way, they can expend their cognitive effort on understanding the underlying ideas rather than on doing mental acrobatics with a table and performing arithmetic operations to arrive at “the right answer.”



Similarly, for instance, the P-value and Power animations available in the Test Means commands also communicate very challenging concepts just when needed – within the context of a specific test using a particular set of data. Here again, an analyst may want to look at power because of the potential stakes in an analysis. A first-year statistics student probably doesn't choose to look at power, but has an instructor who invites her to understand that sometimes a statistical test may not see the signal amidst the noise. Both the professional and the learner can find value in these animations, but for quite different reasons.



KEEPING IT REAL

As noted above, there is widespread agreement on the utility of using real data in statistics courses. One might reasonably ask, though, precisely what we intend by the term “real” and just how much reality can introductory students tolerate before they are either overwhelmed or disillusioned? Much of the available real data is observational, and many large datasets are sparsely populated and/or peppered with errors which can easily become stumbling blocks for novices.

Introductory students are typically taught that one develops a research question and a research hypothesis, gathers a simple random sample of suitable data, performs the appropriate inferential technique and evaluates the hypothesis that prompted the effort. Instructors model this behavior by raising interesting questions, providing real (or realistic) and moderately clean data sets, and walking the students through an analysis through its statistical and contextual conclusions. Perhaps the instructor then sends students off to locate their own real data and conduct a start-to-finish study of their own. This is an approach that I have adopted.

In the name of constructing examples that “work out” or a desire to maintain focus on analytic techniques rather than the drudgeries of data preparation, it is quite reasonable for instructors to prepare real datasets in advance, purging them of some of their natural imperfections. At the same time, if we thoroughly sanitize raw data we compromise some of the reasons for including it in the course.

Excessive data cleansing is just one issue. Another thorny issue is one raised by Prof. Nick Horton and colleagues about the facts that, on the one hand, we hold up the simple random sample as the only authentic and reliable foundation for drawing generalizable inferences but on the other hand know that many students are entering fields where the SRS is a rare phenomenon [11].

Moreover, we also know that when governments, NGOs, and survey research firms conduct survey research, they routinely use multistage complex samplings techniques, and have done so for years [13],[14]. Though many first-level courses and introductory texts may devote space to probabilistic sampling methods in addition to the SRS, it is rare to find fulsome treatments of sampling weights and their use in *analysis* of such data.

Historically, it makes a great deal of sense to defer the topic of post-sampling adjustments to a subsequent course past the first course. Articles treating the topic (for example, [13], [14], [24]) tend to be well beyond the reach of undergraduates in a service course. Yet, just as we have found ways to introduce the *t*-test accessibly with software, current software offerings can make it quite reasonable to teach proper methods of analysis using sampling weights.

In a recent empirical study of learning outcomes associate with the use of real-life data, Neumann *et al.* [17] note that “research experiences should be authentic and grounded in context,” reporting that students in their study found the prospect of real-world application to both motivate their efforts and make concepts memorable. Many instructors have likely observed similar impacts.

Perhaps students just overlook the gap, but the practice of taking a real dataset and then analyzing it in an “unreal” way—that is, a way that practicing statisticians would *not* analyze it—potentially deflates some of the motivational charge that students get from knowing that they are actually and actively doing data analysis. For any students that notice the lapse, we set up an opportunity for cognitive dissonance.

Sampling Weights for Freshmen

There are surely reasons to delay the topic of sampling weights until a later course, perhaps a course for students in fields that rely heavily on survey research. The introductory course is often overloaded with topics. Complex sampling only makes sense subsequent to understanding simple random (and non-random) sampling, so in some respects it is a follow-on concept. Post-sampling adjustments to standard errors *are* complicated both in concept and in terms of the underlying mathematics, and hence probably do go beyond the scope of most introductory courses that de-emphasize formulae and theory. And, of course, we've never done it.

For this paper I have not conducted an extensive review of textbooks on the market, but one has the impression that it is quite common for authors to include a brief mention that surveys often employ a combination of stratified and cluster sampling. In one recent noteworthy newer text [18], after introducing sampling variation and sampling distributions, the authors go on to note that a stratified sample can provide more efficient estimates than the SRS:

“The mean and standard deviation from each stratum would be calculated separately and combined to form an estimate of the population parameters. The standard deviation of each subgroup should be smaller than the standard deviation of the population. When those standard deviations are combined to form an estimate of the population standard deviation, the resulting standard deviation will be smaller than if a simple random sample had been used.” [18]

The authors prudently stop at that point, without explaining precisely how the means and standard deviations of the strata are combined to achieve this result, but the point in this discussion is that the idea is delivered in three sentences. The idea may be a subtle one, and candidly may be lost on some students, but with care the idea can be expressed simply and briefly.

Somewhat unfortunately, this is typically where the treatment of complex sampling and the subsequent combination of subsample statistics ends. This would make sense given the somewhat difficult computational investment that it required—except for the fact that software gracefully and efficiently manages the task readily, much as it computes regression parameter estimates without first requiring students to compute sums of squares.

Rather than belabor the point further, consider two examples as an introduction to estimation based on complex sampling. Both examples use real data from the Web, and both analyses are performed using JMP 11. The first illustration uses a sampling frame from a known population at a given point in time, inspired by Diez' excellent on-line text [6]. The goal here is to introduce students to two ideas through example: 1) that in a stratified sample, observations from different strata represent a different number of population elements and 2) that weighting each observation proportion to the size of its stratum generates a more accurate estimate of a population parameter than simply treating all observations equally.

Once those concepts are established, the second example uses the weights supplied with the NHANES data. NHANES refers to a set of studies administered by the Centers for Disease Control and Prevention. The particular survey illustrated here was a national survey of more than 10,000 individuals, administered in 2005.

Both examples are introduced and accompanied by prose aimed at students. In other words, the corresponding explanations should guide a student through the intended thought process. One might distribute the passages as assigned reading or use them as the foundation for in-class presentation.

We start with a known population and investigate the outcomes of two different sampling methods. We begin with a sampling frame from the United States armed forces [6]. The data table contains selected attributes of 1,048,575 active-duty personnel from four branches of the U.S. military, excluding the Coast Guard, in April 2010. We'll treat this group as a population for the sake of this example. Naturally, in a real situation in which we have the luxury of having access to raw population we can just go ahead and calculate the population parameters. In this example, though, we are trying to get a feel for the accuracy and precision of different sampling methods, so we will start by computing a parameter and then subsequently drawing some samples to see how different methods compare.

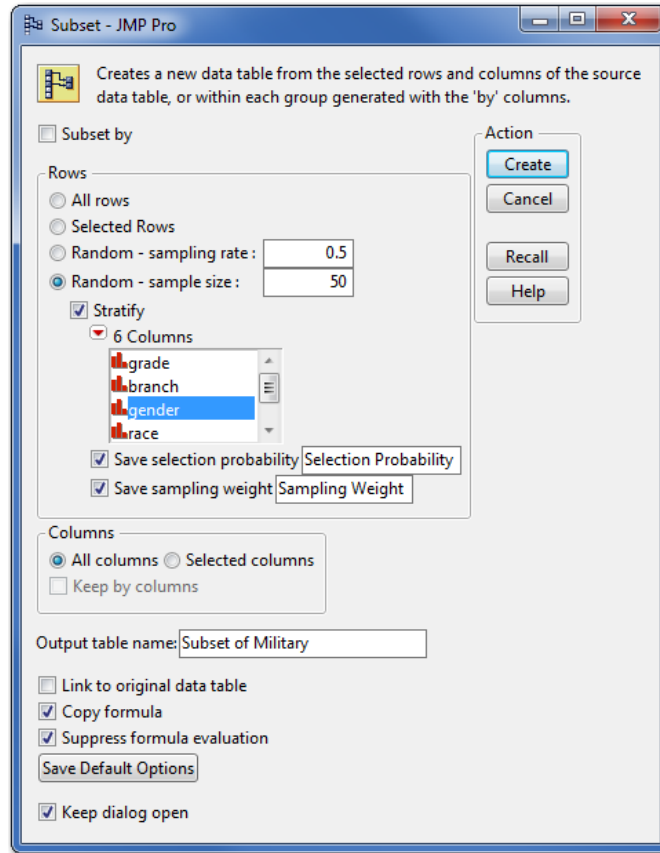
The variables in the dataset are these:

- Grade (enlisted, officer or warrant officer)
- Branch (Air Force, Army, Marine Corps or Navy)
- Gender (female, male)
- Race (7 categories)
- Hisp (Hispanic ethnicity: false, true)
- Rank (ordinal value from 1 to 12)

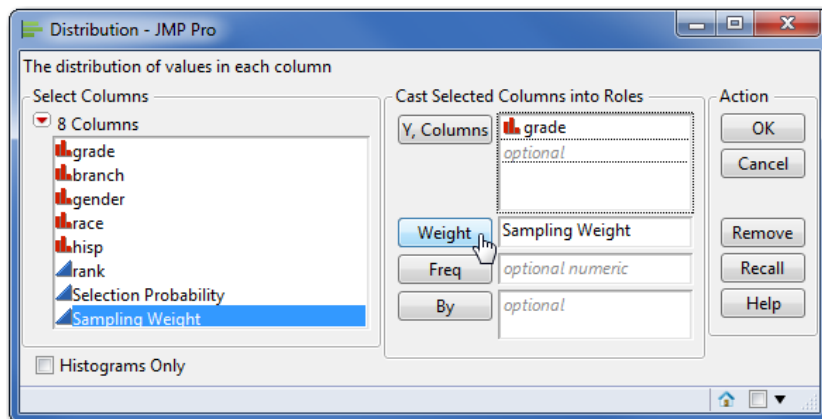
Within this population, approximately 78% of the individuals are enlisted personnel, that 6% are in the Air Force, and that 13% are female. We'll keep these population proportions in mind as we examine some samples. Recall that our goal here is to gain insight in the extent to which different methods of sampling produce samples that represent the population. In particular, we want to allow first-year students to develop some understanding of random sampling that is not *simple* random sampling.

Clustering and stratification both rely on random selection. To illustrate the general approach, we'll use JMP's Subset platform (command) to first select a stratified sample, compute sampling weights, and then develop sample proportions using and not using the weights. So as not to get lost in the bulrushes of numerical detail, I demonstrate the approach using gender to stratify, and indicate that if (for example) we choose a total sample of 100 personnel, 50 women and 50 men, we recognize that a sample with an equal number of men and women will initially misrepresent the entire military. Simply put, each man in the sample will represent a larger number of men than a corresponding woman in the sample.

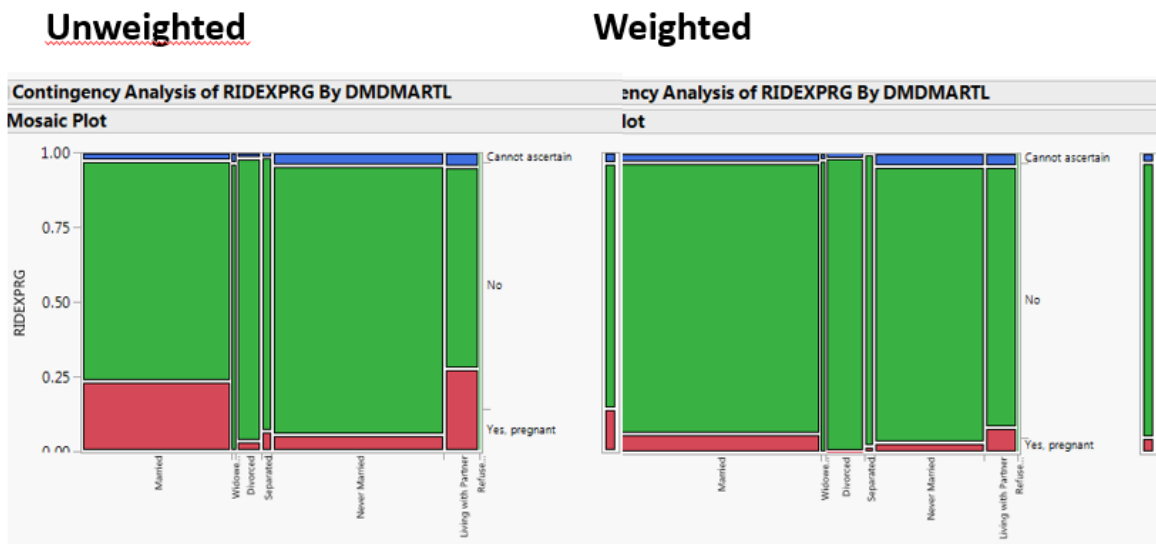
The JMP dialog that will automatically select the stratified sample is shown below. Once the student understands that the goal is to randomly choose 50 women and 50 men, executing the command is straightforward.



At this point, students are positioned to estimate the proportion of (say) enlisted personnel in our stratified sample of 100 individuals. We can anticipate that the disproportionately large number of women in the sample will bias the estimate of enlisted individuals. Indeed, the next class activity would be to have all students generate a stratified sample as just shown, and compute the naïve, unweighted sample proportions. Then, to more properly use the stratified sample, JMP’s descriptive analysis command provides an input box for post-stratification weights. An introductory student can easily and confidently choose the pre-calculated sampling weight when instructed to do so.



The prior example is artificial but potentially instructive. Far more useful and powerful lessons can come from analysis of authentic survey data published by a respected research organization like, for example, the CDC. As noted earlier, the CDC’s NHANES program conducts survey research on, among other things, lifestyle, dietary, and health issues [4]. Investigators combine interviews with closed ended survey questions, and some of the data (e.g. blood pressure, weight) is actually measured. Given the interests of late adolescents, NHANES data is a rich source. In this example, we look at pregnancy status of women in the survey, crossed with marital status. NHANES is decidedly *not* a simple random sample. Among the many data columns is a column of sampling weights. At this course level, it is not necessary to fully explain the CDC’s sample design methodology; I would suggest that it is a huge step forward to have introductory students understand that complex sampling has some advantages over SRS, and that once a complex sample is gathered there are adjustments necessary in analysis. The figure below presents a visual demonstration of the difference that weighting makes:



Intuitive Data Management Tools

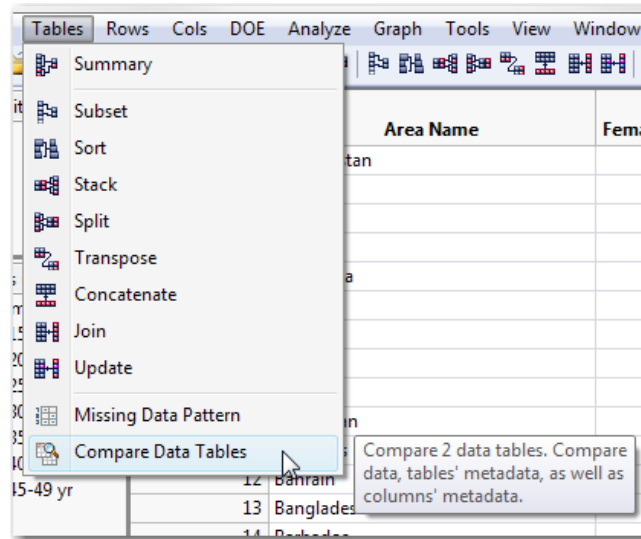
One of the fortunate phenomena of current times is the widespread availability of large, reliable web-accessible databases providing data on a huge variety of subjects. Many of these databases are assembled and maintained by international agencies and governmental entities. Moreover, depending on the client disciplines serviced by the introductory course, it may be as valuable to learn how to navigate the front end of a public database as it is to design and conduct surveys or to carry out experiments.

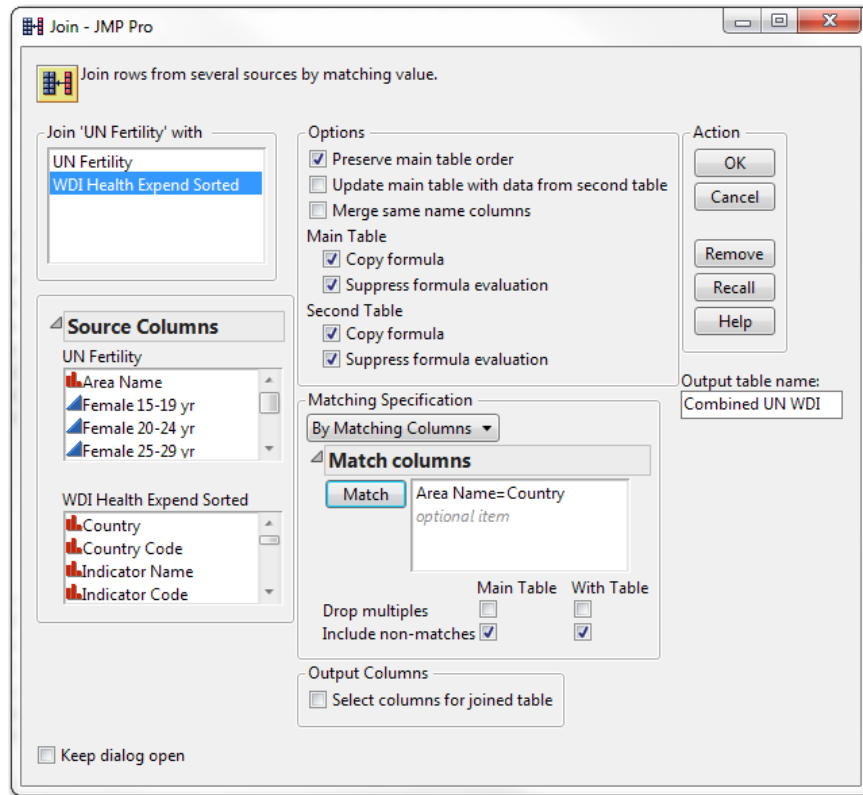
As part of the semester-long team project in my course, I encourage students to find data from two reputable sources, and then prepare it for analysis. I want students to encounter common issues like missing data, inconsistent labeling, varying units, and so on without being completely disabled by the complexity. Even with “major league” web sources, problems abound – particularly for students with no prior background in database systems, data structures or terminology common to data management.

Once forewarned, students can readily make the necessary adjustments, but they need to know enough to watch for problems. Consider a study using data from, for example, the UN Statistics Division’s Gender Info database and The World Bank’s World Development Indicators (WDI).

Specifically, we have per capita health expenditures (current USD) from the World Bank, and Female Fertility Rates for different age groups from the UN. The UN download lists 195 distinct countries with available data from the series of interest; the WDI lists 217 countries. In the UN table, the column listing countries is titled “Area Name” and in the WDI table it is “Country Name.” Moreover, some nations are identified by slightly different names in the two tables. A human reader would readily conclude that the two refer to the same country, but that conclusion is not necessarily obvious to software.

JMP’s table functions provide a “plain language” bridge to the foundations of data management and data cleaning. Students can learn to join two tables while controlling the handling of non-matching keys as well as missing rows. Though it does not use the database terminology of inner and outer joins, check boxes in the Matching Specification panel allow for these operations. At a minimum, the user only needs to specify the tables to merge and the columns to treat as keys—well within the reach of a novice student.



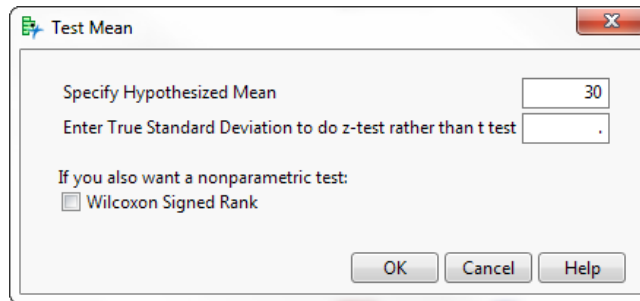


RETHINKING CATEGORIZATIONS OF “ADVANCED” TECHNIQUES

Though there is not universal agreement about the list of topics that belong in an introductory course, there are some topics that most textbook authors seem to identify as beyond the scope. In this section, I consider three more topics that often are absent from introductory textbooks or appear in late chapters or as on-line supplements.

Integration of parametric and non-parametric tests

While most texts specify the assumptions and conditions underlying distribution-based inference, and many indicate that one should “proceed with caution” if assumptions are violated but if non-parametric tests are explained at all, it is typically in a late chapter. In contrast, JMP places the non-parametric options within the context of an analytic goal. From the standpoint of the learner, this is very helpful. Instead of thinking in terms, for instance, of a t-test or a z-test, and what the conditions are for each, consider JMP’s simple-seeming dialog box when a user declares the intention of testing a mean with univariate data. In the unlikely event that the student somehow knows the true standard deviation, JMP will do a z-test. If the student has concerns about the conditions for a t-test, the suitable nonparametric equivalent appears a checkbox. This opens the door to nonparametric testing just in time to learn what it is.

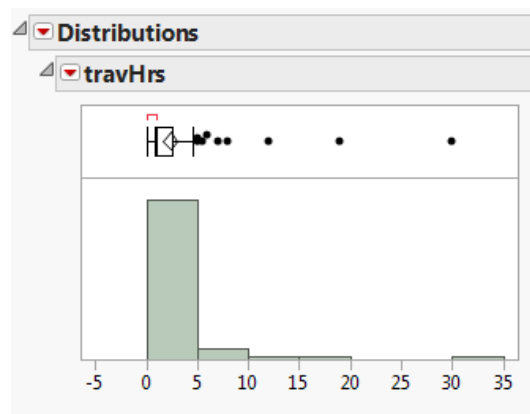


The Fit Y by X platform analogously offers a menu of nonparametric alternatives to the comparison of multiple means. Again, in an introductory course we may not enter into this area in any depth, but the fact that JMP locates it within context is very valuable and presents an instructor with the viable option to move that particular fence.

Bootstrap Confidence Intervals

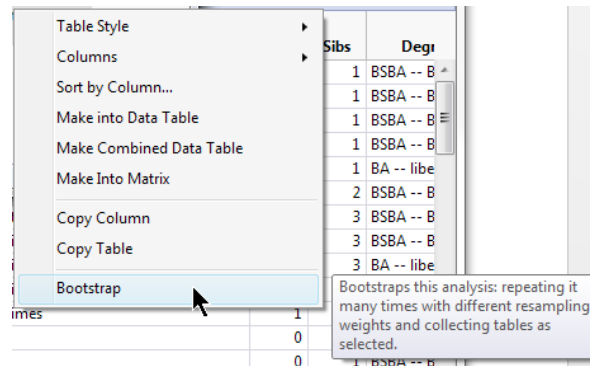
There is probably no concept more challenging than that of sampling distributions. Students are generally comfortable with the idea that we use a sample statistic to estimate a population parameter. However, when we begin to discuss the distribution of all possible sample means and standard error of the sample mean, we see how very fine the line is between comprehension and confusion. Within the statistics education literature, there is a growing consensus that a simulation-based resampling approach is an effective way to introduce the elusive concept of sampling variability [2]. Bootstrap confidence intervals are one common approach, and there are numerous on-line simulators and textbook supplements that provide the technology to take many repeated samples with replacement. Fortunately, instructors with JMP Pro can introduce the bootstrap without leaving the JMP environment.

In my course, I annually survey the students enrolled in our core business statistics course. Because we are a small regional school, most students hail from nearby states. One of the questions on the survey asks how many hours it takes to travel to Stonehill from home. The figure below shows a recent set of responses from just over 100 students.




We have a very skewed sample, and it is the only information we really have about this population. Despite the common guideline of $n > 30$, the Central Limit theorem is probably not going to rescue us here (in fact, there's another JMP simulator that will illustrate the extent to which resampling would generate a symmetric empirical sampling distribution). Fortunately, the

bootstrap option is available by hovering over the Summary Statistics panel in the Distributions report.



In my experience, the bootstrap is not immediately intuitive to many students. Some seem to view it as the computational equivalent of spinning gold from straw, but with careful explanation and a brief demonstration, the idea seems to come across. At any rate, it does provide another answer to the question “what do we do when we can’t rely on the CLT to create a trust-worthy confidence interval?”

The demonstration goes on to show that the distribution of sample means is probably *not* nearly normal in this instance, and that a bootstrap 95% confidence interval is a little wider and to the right of the t-based result.



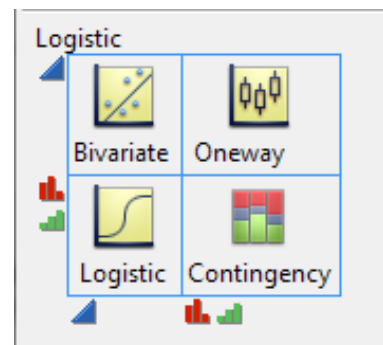
“Easy” Logistic Regression

It is common to organize the bivariate portion of an introductory statistics course along the lines of the schematic familiar to users of the Fit Y by X platform, shown to the lower right.

With some local variations, instructors proceed to deal with three cases:

- Categorical Y, Categorical X – contingency tables and 2-sample proportion and/or Chi-square
- Numeric Y, Categorical X – 2-sample t and ANOVA
- Numeric Y, Numeric X – simple linear regression

As for the fourth case, that is often considered too advanced for an introductory course. Certainly, the leap from a linear model and




straightforward interpretation of coefficients to a logistic model with interpretation of log-odds is a huge one, demanding considerable time and effort. But even if we cannot afford a long stay with the topic, JMP at least makes it very easy to provide some exposure to it. In an era of data mining and classification methods sometimes based on logistic models, there is good reason to whet students' appetites with an interesting example. In this case, the data comes from a series of studies led by Prof. Max Little on monitoring Parkinson's Disease (PD) patients remotely by telephone [15]. Many PD patients live quite far from their neurologists and trekking to the doctor for a routine exam can be quite burdensome. One of the characteristic symptoms of PD is degradation of vocal function, which manifests as tremulous-sounding speech. Little and colleagues used standard electronic measurements of phonation (speech) to see if voice recordings via telephone could be used to monitor the progression of Parkinson's in these patients.

One of the phonation measurements is known as *shimmer*, which refers to fluctuation in a speaker's volume. PD patients tend to have increased shimmer as the disease progresses. Little and colleagues recorded the voices of 32 patients, 23 of whom suffer from PD. The investigators wanted to create a model with a dichotomous dependent variable (PD/no PD) and a continuous predictor, shimmer.

The background narrative here is one that students readily understand, and it provides a captivating topic for a classroom example. In terms of JMP, the platform is one the students know well towards the end of the course, and the "nominal Y/continuous X" completes the quartet of bivariate models.

The logistic model yields some output that may go beyond the introductory course, but the graph that headlines the report nicely illustrates what the logit model does for the analyst. In the figure below, the red markers represent the PD patients, and the blue line is the fitted model.

The most challenging computational operation (for students) is to convert the logistic fit into probability scores for each case, but JMP handles that with an option on the red triangle menu. Thus, even in a first course it is reasonable to introduce the rich potential of logistic models.



“JUST LIKE THE PROS USE!”

Using JMP throughout a full semester course in statistics allows students to develop proficiency with a well-designed mature product that is widely used in industry. Though a single course will not expose students to JMP’s full range of capabilities, it can provide opportunities to gain confidence in their ability to transfer knowledge from one area of the software to another, and to learn about unfamiliar topics. As noted earlier, many of the pedagogical tools embedded in JMP are available elsewhere, but JMP has truly done an outstanding job of combining the tools into a single environment, and continues to do so with each new release. For those students contemplating employment in an analytics-related field, proficiency with JMP is a good resume item.

Outstanding support for undergraduates – documentation, help, screen tips, and more

All software provides internal documentation, but undergraduates typically learn from painful experience that online help is not always helpful. JMP stands out in offering assistance that can speak to students as well as professional users. There is a wide range of books, videos, and tutorials accessible from the Help menu, and JMP continues to develop simulators and teaching demonstrations that not only train users in the software but genuinely develop statistical understanding. All of this keeps students “in the game”, using JMP, building their own understanding of the major statistical ideas, and remaining engaged in the course. Ambitious students can even seek JMP Certification, which in turn is another distinctive entry on a resume.

Removal of the DO NOT ENTER gates by eliminating or lowering barriers to entry

Statistics educators know all too well the reticence that students often have about their required statistics course. Perhaps JMP’s greatest potential contribution is in moving the fence that says “Keep Out”. The intuitive interface and captivating visualizations draw students in, and signal that statistical discovery is accessible and interesting.

JMP’s high degree of interactivity and visual feedback may not only lower the barriers to entry, it may create barriers to exit. Granted, it can’t prevent students from abandoning further study of analytics, but once students learn to put the tool to work and gain confidence in their ability to discover important insights from data, one can reasonably expect growing numbers of students to wonder what comes after the introductory course.

CONCLUSIONS

This is an exciting era in the realms of analytics and statistics education. College faculty are changing the introductory statistics course both in its content and in the strategies used to help students learn about data analysis. We are redefining the scope and content of statistics curricula, and JMP is playing a featured role in these changes. This paper has outlined nine specific ways that JMP is “moving the fences” in undergraduate statistics education.

REFERENCES

- [1] Aliaga, M., Cobb, G., Cuff, C., Garfield, J., Gould, R., Lock, R., Moore, T., Rossman, A., Stephenson, R., Utts, J., Velleman, P., Witmer, J. *Guidelines for assessment and instruction*

- in statistics education: College report*. Alexandria VA: American Statistical Association, 2005. Retrieved December 22, 2013 from <http://www.amstat.org/education/gaise/>.
- [2] Budgett, S., Pfannkuch, M., Regan, M., Wild, C. “Dynamic Visualizations and the Randomization Test.” *Technology Innovations in Statistics Education*, 2005, 7(2). www.escholarship.org/us/item/9dg6h7wb.
- [3] Carver, R. *Practical Data Analysis with JMP, 2nd Ed.*, 2014 Cary NC: SAS Press.
- [4] Centers for Disease Control and Prevention. *National Health and Nutrition Examination Survey*. 2005. Available from <http://www.cdc.gov/nchs/nhanes.htm>.
- [5] Cobb, G., “The Introductory Statistics Course: A Ptolomaic Curriculum?” *Technology Innovations in Statistics Education*, 2005, 1(1). www.escholarship.org/us/item/6hb3k0nz.
- [6] Diez, D., Barr, C., Cetinkaya-Rundel, M. *OpenIntro Statistics, 2nd Ed*, 2013. On-line <http://www.openintro.org/stat/textbook.php>.
- [7] Economist, The “Data Deluge: Special Issue”, February 27, 2010. *The Economist Newspapers Ltd.*, 2010.
- [8] Gould, R. “Statistics and the Modern Student. *International Statistical Review*, 2010, 78(2), 297-315.
- [9] Gould, R., Davis, G., Patel, R., & Esfandiari, M. “Enhancing conceptual understanding with data driven labs. In C. Reading (Ed.), *Data and context in statistics education: Towards an evidence-based society. Proceedings of the Eighth International Conference on Teaching Statistics, Ljubljana, Slovenia*. Voorburg, The Netherlands: International Statistical Institute, 2010. http://www.stat.auckland.ac.nz/~iase/publications/icots8/ICOTS8_C208_GOULD.pdf
- [10] Hogg, R. “Statistical Education: Improvements are Badly Needed”. *The American Statistician*, 1991, 45, 342-343.
- [11] Horton, N., Baumer, B., Kaplan, D., Pruijm, R. “Precursors to the Data Explosion: Teaching How to Compute with Data”. Conference presentation at Joint Statistical Meetings, Montreal, Quebec, 2013. Abstract retrieved January 2, 2014 from <http://www.amstat.org/meetings/jsm/2013/onlineprogram/AbstractDetails.cfm?abstractid=307064>.
- [12] JMP, Version 11. SAS Institute Inc., Cary, NC, 1989–2013.
- [13] Johnson, E., Rust, K., “Population Inferences and Variance Estimation for NAEP Data”. *Journal of Education l Statistics*, 1992, 17(2), 175–190.
- [14] Kish, L. Frankel, M. “Inference from Complex Samples”. *Journal of the Royal Statistical Society. Series B (Methodological)*, 1974, 36, (1), pp. 1–37

- [15] Little, M.A., McSharry, P.E., Hunter, E.J., Ramig, L.O., “Suitability of Dysphonia Measurements for Telemonitoring of Parkinson’s Disease”, 2008, *Nature Proceedings*.
- [16] Manyika, J., Chui, M., Brown, B., Bughin, J., Dobbs, R., Roxburgh, C., Byers, A.H. “Big data: The Next Frontier for Innovation, Competition, and Productivity”. Retrieved August 27, 2011, http://www.mckinsey.com/insights/business_technology/big_data_the_next_frontier_for_innovation.
- [17] Neumann, D., Hood, M., Neumann, M., “Using Real-Life Data When Teaching Statistics: Student Perceptions of this Strategy in an Introductory Statistics Course”. *Statistics Education Research Journal*, 2013, 12(2), 59–70.
- [18] Nottingham, Q., Hawkes, J. *Discovering Business Statistics*. Mt. Pleasant SC: Hawkes Learning Systems, 2013.
- [19] Rosling, Hans. “Wealth & Health of Nations”, Retrieved August 13, 2014, <http://www.gapminder.org/world>.
- [20] Singer, J., Willett, J. “Improving the Teaching of Applied Statistics: Putting the Data Back into Data Analysis”. *The American Statistician*, 1990, 44(3), 223–230.
- [21] The World Bank. *World Bank Open Data*. Accessed December 7, 2013 from <http://data.worldbank.org/>.
- [22] Tukey, J. W. *How Computing and Statistics Affect Each Other*. Paper presented at the Babbage Memorial Meeting: Report of Proceedings, London, 1972.
- [23] United Nations Statistics Division. *UNdata*. Retrieved January 2, 2014 from <http://data.un.org/Explorer.aspx>.
- [24] Verma, V, Scott, C., O’Muircheartaigh. “Sample Designs and Sampling Errors for the World Fertility Survey”, *Journal of the Royal Statistical Society, Series A (General)*. 1980, 143(4), 431–473.