Using Minitab to Improve Students' Understanding of the Central Limit Theorem

Kate McGivney
Shippensburg University

Introduction. The Central Limit Theorem (CLT) says the following:

> Let $X_1, X_2, \ldots, X_n$ be a random sample drawn from an arbitrary distribution with a finite mean μ variance σ^2. As $n \to \infty$, the sampling distribution of $\dfrac{\bar{X} - \mu}{\sigma/\sqrt{n}}$ converges to the $N(0,1)$ distribution.

Students in an introductory probability and statistics course typically have a hard time making sense of the Central Limit Theorem. I have tried a number of strategies to help them understand this powerful result. Early on, I would merely state the theorem, attempt to shed insight into its beauty, and then do a few "canned" examples that required the students to apply the theorem. What I found is that, with practice, students could apply this result, but that they failed to have a deep understanding. For the past two semesters, before discussing (or even stating) the Central Limit Theorem, I asked the students to work in small groups on a Minitab lab that was designed, so that with any luck, they would discover the result for themselves. The students were given two days to work on the project and then asked to write a one to two-page paper summarizing what they learned from the experience. I provide an overview of this activity as well as highlight the benefits of using this approach to teach the CLT.

Background Knowledge. Prior to this lab, the students were given a brief introduction to probability, about 3 weeks. The unit began with an overview of the basic axioms and concluded with a discussion on random variables and probability distributions, specifically the uniform, binomial, and normal distributions.

A Sample Activity. After concluding our discussion of random variables, I posed the following problem to the class.

> *Suppose that you have a spinner that is broken into 5 parts, each of equal area. Each section is labeled 1, 2, 3, 4, or 5. Spin the spinner and observe the number that appears. Let X be the selected number.*

What is the probability distribution of X? The students immediately observed that X is a discrete uniform random variable and thus it has the following probability distribution.

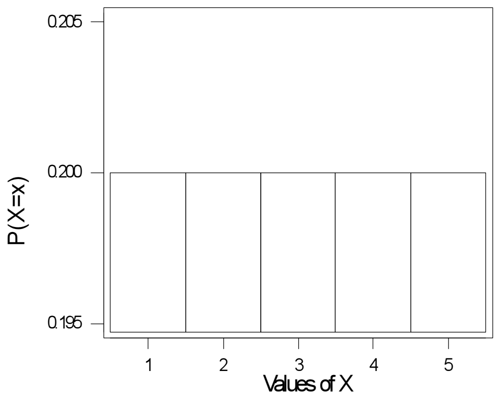

Students then were asked to describe the distribution of X numerically. They computed the expected value, standard deviation, and the minimum and maximum of X and came to the following conclusions:

	Mean	**Standard Deviation**	**Minimum**	**Maximum**
X	3	1.4142	1	5

Next we turned our attention to analyzing the distribution of the sample means of size n. The students were asked to think about spinning a spinner n times and finding the average number spun. We let $\bar{X}_n$ represent the average number based on n spins and noted that it is a random variable. The next question that the students investigated was, "How does this random variable vary?" Using Minitab, the students began by looking at 200 samples, each of size one, drawn from the population described above. (In other words, they were looking at the distribution of 200 sample means each of size 1.)

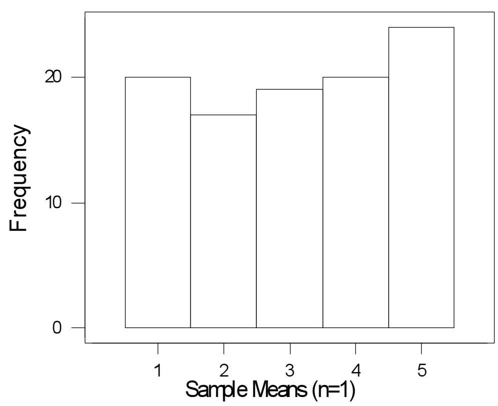

They noticed that the shape of the distribution of $\bar{X}_1$ is very similar to the distribution of X, as is the mean, standard deviation, minimum and maximum values. This result agreed with their intuition.

	Mean	**Standard Deviation**	**Minimum**	**Maximum**
$\bar{X}_1$	3.11	1.463	1	5

I then asked the students to think about spinning the spinner 5 times and averaging the data. I asked them to repeat this experiment 200 times and to describe the approximate sampling distribution of the sample mean, $\overline{X}_5$. They were to think about how this histogram differs from the histogram that describes the approximate distribution of $\overline{X}_1$. Specifically, does the shape look similar? Is the mean roughly the same? Which distribution has more spread? The students continued to look at the distribution of 200 samples, but of increasing size. They experimented with samples of size 25 and 100 and were asked to comment on the commonalties and differences in the distributions. Listed below is a table of results and the corresponding graphs that are based on simulation data.

Approximation distribution (200 samples in each case)	**Mean**	**Standard Deviation**	**Minimum**	**Maximum**
$\overline{X}_5$	3.0420	0.6600	1.20	4.80
$\overline{X}_{25}$	3.0146	0.2840	2.32	3.68
$\overline{X}_{100}$	2.9874	0.1417	2.62	3.42

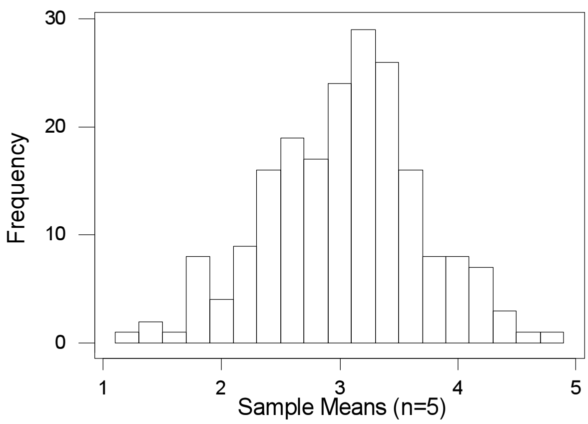

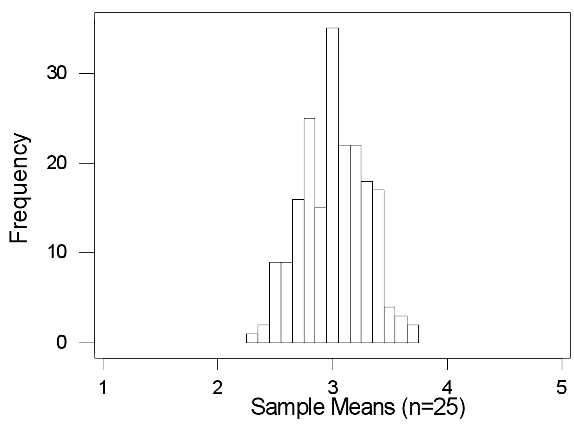

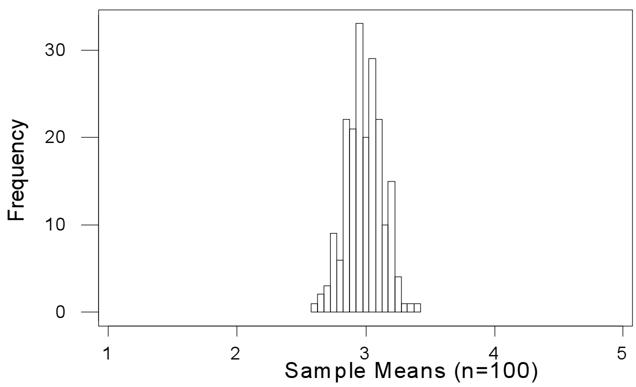

What did the students learn? The majority of the groups discovered the key elements of the Central Limit Theorem, at least for random samples of size n drawn from a population with a discrete uniform distribution. By looking at both graphical and numerical summaries, they discovered that, in each case, the average of the 200 sample means was extremely close to three, the population mean, and that as n increased, the shape of the histograms looked more and more like bell-shaped curves. In terms of analyzing the spread of the data, the students found that as n increased, the spread of the sample means decreased. They noticed that as n increased the differences between the minimum and maximum decreased, as did the standard deviations. Moreover, I asked the students to compute $\sigma/\sqrt{n}$ for each n. They observed that, in each case, the standard deviation of the simulated sample means was very close to $\sigma/\sqrt{n}$.

Extension questions. A natural follow-up question is: "Does the shape of the underlying distribution influence the shape of the sampling distribution of $\bar{X}_n$? " In other words, suppose that the distribution of $X_1, X_2,..., X_n$ is no longer uniform but instead of another other type, say normal or exponential. Generate data according to one of these distributions and consider random samples of size n, what do you notice about the sampling distribution of $\bar{X}_n$? Based on the lab, many students thought that the fact that the underlying distribution was discrete uniform was a necessary condition for the distribution of the sample means to be roughly normal. To dispel this myth, I had the class look at data drawn from both an exponential and normal population, but the problems were put in the context of a "real-world" situation. For instance, the exponential data modeled waiting times for a shuttle bus and the normal data modeled lengths of human pregnancies. The students found it easier to grasp the probability concepts when the problems were stated in context. After experimenting with data drawn from exponential and normal distributions, they quickly saw that the observed results held in these cases too. A second extension question that followed was: "How large does n need to be so that the distribution of the sample means is approximately normal?" The Central Limit Theorem states that the distribution of the sample means is approximately normal "for n large enough". By experimenting with data drawn from three very different distributions: normal, uniform, and exponential, the students saw that the more "non-normal" the underlying population distribution was, the larger n needed to be.

Conclusion. Statistical software packages, such as Minitab, allow students to quickly experiment with a large variety of data drawn from range of probability distributions and to make conjectures regarding the sampling distribution of the sample mean. Using data that models a "real-world" problem and allowing the students to discover the theorem via Minitab provides them with a deeper (and hopefully longer lasting) understanding of the Central Limit Theorem.

References

Tanhame, Ajit and Dunlop, Dorothy (2000), *Statistics and Data Analysis from Elementary to Intermediate*, Prentice Hall.