Simulating the Central Limit Theorem¹

Marshall A. Taylor
University of Notre Dame

Abstract: Understanding the central limit theorem is crucial for comprehending parametric inferential statistics. Despite this, undergraduate and graduate students alike often struggle with grasping how the theorem works and why researchers rely on its properties to draw inferences from a single unbiased random sample. In this paper, I outline a new Stata package, `sdist`, which can be used to simulate the central limit theorem by generating a matrix of randomly generated normal or non-normal variables and comparing the true sampling distribution standard deviation to the standard error from the first randomly-generated sample. The user also has the option of plotting the empirical sampling distribution of sample means, the first random variable distribution, and a stacked visualization of the two distributions.

Keywords: central limit theorem, simulation, `sdist`, `runiform`, teaching

1 Introduction

Virtually every introductory statistics course discusses the central limit theorem (CLT). Far from simply being an abstract mathematical concept that can be relegated to a footnote or a term in statistics history to which instructors pay a fleeting tribute, the CLT is the bedrock upon which parametric inferential statistics stand. Without the CLT, parametric tests—from the simple *t*-test to advanced statistical models—make little sense as tools for estimating population dynamics. Put simply but seriously, the CLT is what allows frequentist statisticians to do what they do.

Despite the theorem's prevalence in the classroom, there is no guiding method for teaching the subject. The CLT can often be difficult to communicate effectively in a classroom setting with a textbook (Dyck and Gee 1998), and numerous pedagogical tools have been proposed to promote hand-on learning (e.g., Aberson et al. 2000; Dinov, Christou, and Sanchez 2008; Matz and Hause 2008; Price and Zhang 2007; Schoenfelder et al. 2007). A number of these tools involve field exercises, such as having students collect data outside of the classroom. While innovative and

¹ Direct all correspondence to Marshall A. Taylor, Department of Sociology, University of Notre Dame, 4046 Jenkins-Nanovic Hall, Notre Dame, IN, 46556; e-mail: mtaylo15@nd.edu. The package outlined in this paper is available through the Statistical Software Components (SSC) archive. The package—as well as a replication repository for reproducing the simulations in this paper—are available on the author's GitHub account: <https://github.com/Marshall-Soc/SDIST>. The author wishes to thank the editors, anonymous reviewer, and Richard Williams for providing important and helpful advice on the package.

kinesthetically engaging, these strategies can be difficult to implement if the material must be covered quickly (such as in statistics “boot camps” or short courses) or if other contextual factors preclude the instructor from being able to use such involved and time-intensive strategies. Further, given that the power of the CLT becomes more evident as the number of samples (and number of observations within the samples) grows asymptotically, strategies that involve the collection and/or distribution of physical materials are constrained by the number of material items at hand (e.g., M&Ms [Dyck and Gee 1998] or rose blooms [Schoenfelder et al. 2007]). Further, though a number of computer-based tools are available for simulating the CLT (e.g., Caro 2015; Rice Virtual Lab in Statistics n.d.; Dinov et al. 2008), these still tend to either be restrictive in terms of the number of random samples that can be generated or are Web-based and therefore not applicable on machines without an internet connection. Finally, what both material- and computer-based methods lack is hands-on interaction with the statistical software that students will inevitably be using during (and hopefully after) their statistics course.²

In this paper, I outline a simple new package, `sdist`, which can be used to simulate the central limit theorem within the Stata environment. This is accomplished by (1) generating a matrix of randomly-generated normal or non-normal variables, (2) plotting the associated empirical sampling distribution of sample means, (3) comparing the true sampling distribution standard deviation to the standard error from the first randomly-generated variable, and (4) producing a side-by-side comparison of the two distributions. The package allows the student to alter the number of random samples, the number of observations per sample, the type of distribution from which the samples are drawn, and the parameters of the distribution. The code is purposefully kept simple in an attempt to promote student experimentation of the simulator outside of the classroom.

In what follows, I first describe the simulation procedure performed by `sdist`. I then outline the syntax, options, and outputs of the package. I close with an illustration of `sdist` in action.

² The Statistical Consulting Group at the UCLA Institute for Digital Education and Research provides an interactive GUI-based CLT simulator in Stata called `clt` (IDRE n.d). While the `clt` function offers more parent distribution types and is considerably faster than `sdist` when more samples are drawn, it does not offer the ability to visualize or report the mean, standard deviation, or standard error for one of the observed variables which goes into the sampling distribution of sample means. This feature, which `sdist` has, is particularly important when trying to illustrate the asymptotic normality characteristics of the CLT with non-normal variables. The `clt` package also does not appear to be functional on Mac operating systems (or at least newer versions of them).

2 Simulation Procedure

The `sdist` package compares an empirical sampling distribution of means from a matrix of randomly-generated variables (samples) to an empirical distribution of one of the random variables used to generate the sampling distribution. Let $\mathbf{X}$ be an $n \times r$ matrix consisting of n observations and r random variables:

$$\mathbf{X} = \begin{bmatrix} x_{11} & x_{21} & x_{31} & \dots & x_{r1} \\ x_{12} & x_{22} & x_{32} & \dots & x_{r2} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ x_{1n} & x_{2n} & x_{3n} & \dots & x_{rn} \end{bmatrix}, \quad (1)$$

where each x_{ij} cell entry is a randomly-generated number from some distribution (normal or non-normal). Each cell entry is independent of every other cell.

From here, a row vector, $\mathbf{x}$, of sample (column) means is generated from $\mathbf{X}$, where the i^{th} sample mean in the vector is found with:

$$\bar{x}_i = \frac{x_{i1} + x_{i2} + \dots + x_{in}}{n}. \quad (2)$$

This row vector is then transposed to make a column vector $\mathbf{x}^T$ of $\bar{x}_r$ means:

$$\mathbf{x}^T = \begin{bmatrix} \bar{x}_1 \\ \bar{x}_2 \\ \bar{x}_3 \\ \vdots \\ \bar{x}_r \end{bmatrix}. \quad (3)$$

The frequency distribution of $\mathbf{x}^T$ is the empirical sampling distribution of r random sample means, each generated from n observations.

The standard deviation of this sampling distribution, $\sigma_{\bar{x}}$, is then calculated:

$$\sigma_{\bar{x}} = \sqrt{\frac{\sum_{i=1}^r (\bar{x}_i - \mu_{\bar{x}})^2}{r}}. \quad (4)$$

where $\mu_{\bar{x}}$ is the mean of the empirical sampling distribution.

This “true” standard deviation of the sampling distribution is then compared to the standard error estimate from x_1 : $se_{\bar{x}} = s_1/\sqrt{n}$, where s_1 is the standard deviation of x_1 . In addition to $\sigma_{\bar{x}}$ and $se_{\bar{x}}$, a Δ estimate is also reported:

$$\Delta = |\sigma_{\bar{x}} - se_{\bar{x}}|, \quad (5)$$

where Δ is simply the absolute difference between the true standard deviation of the sampling distribution and its standard error estimate—the value of which, per the asymptotic properties of the CLT, will shrink as r and/or n increases.

3 The **sdist** Package

sdist performs the above simulation using a combination of base Stata functions native to at least Stata 13.1. The reliance on base functions is intentional and meant to promote quick classroom implementation. The syntax for **sdist**, outlined below, was also kept simple for this purpose. A consequence of this emphasis on simplicity is that only a selection of graphical parameters from the **histogram** and **graph combine** functions are available for customization. Future updates to the package may increase the number of customizable graphical parameters.

3.1 Syntax

```
sdist [, samples(#) obs(#) type(string) par1(#) par2(#) round(#)
      histplot saveplot1(string) saveplot2(string) replot
      combine lcolor(string) fcolor(string) bckg(string)
      nlcolor(string) nlwidth(#) nlpattern(string) dots]
```

3.2 Options

samples(#) is the number of r random samples to generate. The default is 200.

obs(#) is the number of n observations per sample. The default is 500.

type(string) is the type of distribution from which the random samples should be drawn. The default is **type(uniform)**, which generates random samples from a rectangular uniform distribution. Normal and Poisson distributions are also available, indicated by **type(normal)** and **type(poisson)**, respectively. The distributions are created through calls to Stata's random number generators (StataCorp n.d.1).

par1(#) is the first parameter to be specified depending on the distribution selected in **type()**. Since the default **type()** is the rectangular distribution, the default is the lower end of the $[a, b]$ interval. The samples are generated through the **runiform()** function so the default for a is 0, but this can be changed. If **type(normal)** is selected, this parameter

is the mean, with a default of 0. This parameter does not specify anything if `type(poisson)` is selected; use `par2(#)` to specify the mean of the Poisson distribution instead.

`par2(#)` is the second parameter to be specified depending on the distribution selected in `type()`. Since the default `type()` is the rectangular distribution, the default is the higher end of the $[a, b]$ interval. The samples are generated through the `runiform()` function so the default for b is (an approximation of) 1, but this can be changed. If `type(normal)` is selected, this parameter is the standard deviation, with a default of 1. If `type(poisson)` is selected, this parameter is the mean, also with a default of 0.

`round(#)` is the decimal point to which $\mu_{\bar{x}}$, $\sigma_{\bar{x}}$, $\bar{x}_1$, s_1 , $se_{\bar{x}}$, and Δ should be rounded. The default is 0.001.

`histplot` indicates whether or not histograms of the $\mathbf{x}^T$ and x_1 frequency distributions should be plotted. The plots are generated through Stata's `histogram` command (StataCorp n.d.2). Defaults to no histogram.

`saveplot1(string)` indicates whether the $\mathbf{x}^T$ histogram should be saved and the name for the plot. Defaults to "plot1.gph" if `replot` is specified but not `saveplot1()`. Ignored if `histplot` is not specified. The default is to not save the plot.

`saveplot2(string)` serves the same purpose as `saveplot1()`, but with reference to the x_1 histogram. Defaults to "plot2.gph" if `replot` is specified but not `saveplot2()`. Ignored if `histplot` is not specified. The default is to not save the plot.

`replot` specifies whether or not the saved histograms should replace existing saved histograms in the same directory with the same name. `replot` will default to saving both plots if neither `saveplot1()` nor `saveplot2()` is specified, using "plot1.gph" and "plot2.gph" as the file names, respectively. The default is to not replace plots.

`combine` indicates whether or not the two histograms should be stacked to form a third plot. This is a call to the `graph combine` function (StataCorp n.d.4). Both histograms have to be saved in order for the graphs to be combined, either by specifying both `saveplot1()` and `saveplot2()` simultaneously or by specifying `replot` without either `saveplot` options (though `replot` can still be used in conjunction with both `saveplot1()` and `saveplot2()` if both are specified). Requiring that both `saveplot1()` and

`saveplot2()` or `replot` only be specified prevents the program from erroneously stacking histograms from different simulations. The default is to not stack the plots.

`lcolor` indicates the outline color of the histogram bars. This is a call to the `histogram` function (StataCorp n.d.2). The default is black.

`fcolor` indicates the interior color of the histogram bars. This is a call to the `histogram` function (StataCorp n.d.2). The default is `gs6`.

`bckg` indicates the color of the graph region background. This is a call to the `graphregion(fcolor())` argument within the `histogram` command (StataCorp n.d.3). The default is white.

`nlcolor` indicates the color of the normal curve line. This is a call to the `normopts(lcolor())` argument within the `histogram` command (StataCorp n.d.2). The default is black.

`nlwidth` indicates the thickness of the normal curve line. This is a call to the `normopts(lwidth())` argument within the `histogram` command (StataCorp n.d.2). The default is `.5`.

`nlpattern` indicates the pattern of the normal curve line. This is a call to the `normopts(lpattern())` argument within the `histogram` command (StataCorp n.d.2). The default is solid.

`dots` indicates whether or not the program should show simulation progress using the `_dots` function. The default is no dots.

3.3 Output³

By design, `sdist` will not run if there are any data in memory. If data are in memory, the program will error out and inform the user that they need to clear any data before running `sdist`. After clearing all data and executing the function, `sdist` returns a simple table with $\sigma_{\bar{x}}$, $se_{\bar{x}}$, and Δ . The table also includes a note on the Δ estimate, explaining why the difference between $\sigma_{\bar{x}}$ and $se_{\bar{x}}$ may be as large (or small) as it is (e.g., if r or n are large or small). Below is an example of the output printed to the results pane.

³ The examples in this article were created on a machine running Stata 14.2. The examples will output slightly different results when using Stata versions earlier than 14, since a new random number generator was introduced with Stata 14 (Gopal 2016).

```
. set seed 544

. sdist, histplot samples(500) type(poisson) par2(1) ///
> saveplot1(emp_dist.gph) saveplot2(single_dist.gph) replot combine
```

	sd/se
-----+-----	
sig_Xb	.045
se_Xb	.047
abs(diff)	.002

The difference between sig_Xb and se_Xb is .002. The larger this difference, the poorer the single X variable standard error approximates the standard deviation of the sampling distribution. This may be due to one of two things: a small number of samples and/or a small sample size.

By issuing `histplot`, ready-made plots of the $\mathbf{x}^T$ and x_1 distributions are generated—complete with $\mu_{\bar{x}}$ and $\sigma_{\bar{x}}$ for the sampling distribution and $\bar{x}_1$, s_1 , and $se_{\bar{x}}$ for the variable distribution. An example of the plot is provided in Figure 1 below, which was generated from the above command. While Figure 1 (the combined plot one gets after specifying the `combine` option) is not automatically saved, the two histograms composing it would have been saved since the user specified the `saveplot1()` and `saveplot2()` options. In this case, the hypothetical user chose to name the $\mathbf{x}^T$ histogram “emp_dist.gph” and the x_1 histogram “single_dist.gph.”

Figure 1. The empirical sampling distribution (top) and the first variable distribution (bottom) from a Poisson parent distribution.

If the user had specified `replot` without explicitly saving either of the histograms, `sdist` would have defaulted to the assumption that the user was wanting to save the plots (otherwise they would not have attempted to replace other existing files by issuing `replot` in the first place). Therefore, the two plots would have been saved with the generic names “plot1.gph” and “plot2.gph.” Note, however, that the `combine` argument will never work when the two histograms are not saved in some way (through some combination of `saveplot1()`, `saveplot2()`, or `replot`), since `combine` assumes that the plots have been saved. If `replot` had been issued with `saveplot1()` but not `saveplot2()`, then only the first graph would overwrite any pre-existing file with the same name in the directory. The reverse would also be the case if one issues `replot` with `saveplot2()`.

4 Example

`sdist` requires an empty dataset in order to maximize sample size flexibility and prevent other variables from inadvertently being implemented into the simulation procedure. If the user

attempts to execute the `sdist` command with data loaded in memory, the function will error out with the following explanation:

```
. sdist
Save and/or clear existing data before running -sdist-.
r(4);
```

The random variables and any of their associated classes of objects are wiped from memory after the simulation is completed. The user may also use the `set seed` function before executing `sdist` to recreate any results.

As a hypothetical example, let us say an instructor wishes to illustrate the CLT property of asymptotic normality by comparing the differences between the standard deviation of an empirical sampling distribution of sample means and the standard error from a single sample using two sampling distributions of different sizes: one with an r of 500 and another with an r of 10,000 (both with a sample size of 2,000 and both from a rectangular parent distribution). To run the first simulation, the instructor would simply run `sdist` and set the `samples ( )` parameter to 500 and the `obs ( )` parameter to 2,000. The program would then return the following:

```
. set seed 2144
. sdist, histplot samples(500) obs(2000) replot combine
```

	sd/se
sig_Xb	.007
se_Xb	.006
abs(diff)	.001

The difference between `sig_Xb` and `se_Xb` is .001. The larger this difference, the poorer the single X variable standard error approximates the standard deviation of the sampling distribution. This may be due to one of two things: a small number of samples and/or a small sample size.

The output tells us that the standard deviation of the empirical sampling distribution, $\sigma_{\bar{x}}$, is .007, and the standard error estimate of that standard deviation, $se_{\bar{x}}$, from the first sample in the simulation is .006. The absolute difference between these two numbers, Δ , highlights the fact that the standard error from just one of our samples—which is very much non-normal (see Figure 2 below, which we get after combining `histplot`, `replot`, and `combine`)—is nonetheless a very good estimate of the sampling distribution standard deviation that it seeks to approximate.

The key property of the CLT—i.e., that the sampling distribution of sample means for a variable with a sufficiently large sample size will be approximately normal, regardless of the shape

of the variable's population distribution or observed sample distribution—is visualized in the plots that can be generated through `sdist`. The main parameters of interest from the `sdist` table are also featured in the plot, along with the sampling distribution mean, the individual sample mean, and the individual sample standard deviation.

Figure 2. The empirical sampling distribution and variable distribution from a simulation with $r = 500$ and $n = 2,000$.

Though the example with an r of 500 does a decent job by itself of illustrating how large sampling distributions with sufficient sample sizes approach normality, the CLT property of asymptotic normality is made that much clearer when juxtaposing Figure 2 with the simulation featuring an r of 10,000:

```
. set maxvar 32767

. set seed 4816

. sdist, samples(10000) obs(2000) histplot saveplot1(emp_dist.gph) ///
> saveplot2(single_dist.gph) repplot combine
```

	sd/se
sig_Xb	.006
se_Xb	.006
abs(diff)	0

The difference between `sig_Xb` and `se_Xb` is 0. The larger this difference, the poorer the single `X` variable standard error approximates the standard deviation of the sampling distribution. This may be due to one of two things: a small number of samples and/or a small sample size.

In addition to noting the smoothness around the normal curve with the larger simulation (see Figure 3), the instructor could also point out the diminishing presence of outliers and explain how, when taking averages from multiple samples, the impact of influential observations diminishes as sample means cluster around the center of the sampling distribution—especially when the sampling distribution grows larger.

Figure 3. The empirical sampling distribution and variable distribution from a simulation with $r = 10,000$ and $n = 2,000$.

5 Conclusion

The CLT is fundamental to statistical practice and education. Without quality knowledge of how and why it works, the inferential power of parametric statistics can be difficult to grasp. Though numerous pedagogical tools are available to promote hands-on learning, these strategies are often constrained by the availability of material resources, or, if computer-based, are not featured within the statistical computing environments within which students will inevitably have to gain experience. The package detailed here, `sdist`, addresses both of these shortcomings by

(1) employing simulations that are constrained only by the memory allowances of the user's Stata software, and (2) using a simple syntax structure that promotes in-class and at-home experimentation among students new to the Stata environment.

There are, of course, aspects of `sdist` that can be improved. For instance, since the program requires variable space, larger simulations can only be performed in Stata/SE and/or Stata/MP. At a certain point in the simulation, Stata must be able to hold $r \times 3$ variables in memory. As such, the 10,000 sample example above required space for 30,000 variables—something that Stata/IC and Small Stata cannot provide.⁴ This memory demand also means the processing speed decreases with larger simulations; as such, users are encouraged to pre-run larger simulations prior to instruction. It is hoped that these limitations will be addressed in future iterations of the package.

The CLT is, simultaneously, deceptively simple and deceptively complex. `sdist` can hopefully provide those students and practitioners working within Stata an opportunity to experiment and play around with the theorem, figuring out what exactly makes it tick.

6 References

- Aberson, C. L., D. E. Berger, M. R. Healy, D. J. Kyle, and V. L. Romero. 2000. Evaluation of an interactive tutorial for teaching central limit theorem. *Teaching of Psychology* 27: 289-291.
- Caro, R. J. 2015. Central limit theorem simulator. Accessed February 9, 2017 (https://rpubs.com/RamiroJC/CLT_Slides).
- Dinov, I. D., N. Christou, and J. Sanchez. 2008. Central limit theorem: New SOCR applet and demonstration activity. *Journal of Statistics Education* 16:1-15.
- Dyck, J. L., and N. R. Gee. 1998. A sweet way to teach students about the sampling distribution of the mean. *Teaching of Psychology* 192-195.
- Gopal, K. 2016. How to generate random numbers in Stata. *The Stata Blog*. Accessed January 24, 2018 (<https://blog.stata.com/2016/03/10/how-to-generate-random-numbers-in-stata/>).
- IDRE. N.d. Stata teaching tools: Central limit theorem. <https://stats.idre.ucla.edu/stata/ado/teach/stata-teaching-tools-central-limit-theorem/>.
- Matz, D. C., and E. L. Hause. 2008. “Dealing” with the central limit theorem. *Teaching of Psychology* 35: 198-200.

⁴ The maximum number of random samples Stata/IC can handle is therefore about 682. Small Stata maxes out around 33.

- Price, B. A., and X. Zhang. 2007. The power of doing: A learning exercise that brings the central limit theorem to life. *Decision Sciences* 5: 405-411.
- Rice Virtual Lab in Statistics. N.d. Sampling distributions. Accessed February 9, 2017 (http://onlinestatbook.com/stat_sim/sampling_dist/).
- Schoenfelder, E., R. Olson, M. Bell, and K. Tom. 2007. Stop and smell the roses: An activity for teaching the central limit theorem. *Psychology Learning & Teaching* 6: 80-84.
- StataCorp. N.d.1. Random-number functions. College Station, TX: StataCorp LP. Accessed September 26, 2017 (<https://www.stata.com/manuals13/m-5runiform.pdf>).
- StataCorp. N.d.2. `histogram` – Histograms for continuous and categorical variables. College Station, TX: StataCorp LP. Accessed September 26, 2017 (<https://www.stata.com/manuals13/rhistogram.pdf>).
- StataCorp. N.d.3. `region_options` – Options for shading and outlining regions are controlling graph size. College Station, TX: StataCorp LP. Accessed September 26, 2017 (https://www.stata.com/manuals13/g-3region_options.pdf).
- StataCorp. N.d.4. `graph combine` – Combine multiple graphs. College Station, TX: StataCorp LP. Accessed September 26, 2017 (<https://www.stata.com/manuals13/g-2graphcombine.pdf>).