Beyond normal: Understanding power through R Shiny

Christian Stratton, Montana State University

Jennifer Green, Michigan State University

Motivation

Suppose that $X_i \overset{iid}{\sim} \text{Exp}(\theta)$, where θ represents the mean, and we take a random sample of size n from this population. Further suppose that we wish to test the following hypotheses:

$$H_0\colon \theta \le \theta_0$$
$$H_a\colon \theta > \theta_0$$

Consider the following test function:

$$\phi(\boldsymbol{X}) = \begin{cases} 1 & \sum_{i=1}^n X_i > k \\ 0 & \text{else} \end{cases},$$

where k is chosen such that $Pr(\phi(\boldsymbol{X}) = 1|\theta_0) = \alpha$. It can be shown that if $X_i \overset{iid}{\sim} \text{Exp}(\theta)$, then $\sum_{i=1}^n X_i \sim \text{Gamma}(n, \theta)$. Let $T = \sum_{i=1}^n X_i \sim \text{Gamma}(n, \theta)$. To define our test, we seek the value k such that

$$Pr(\phi(\boldsymbol{X}) = 1|\theta_0) = Pr(T > k|\theta_0) = 1 - Pr(T \le k|\theta_0) = \alpha$$

The value of k that satisfies this equation is the $(1-\alpha)^{th}$ quantile of the $\text{Gamma}(n, \theta_0)$ distribution. Letting $\Gamma_{n,\theta_0,1-\alpha}$ denote this value, our test functions becomes

$$\phi(X) = \begin{cases} 1 & \sum_{i=1}^n X_i > \Gamma_{n,\theta_0,1-\alpha} \\ 0 & \text{else} \end{cases}$$

Therefore, our power function is

$$Power(\theta) = Pr(\boldsymbol{X} \in RR) = Pr\left(\sum_{i=1}^n X_i > \Gamma_{n,\theta_0,1-\alpha}\middle|\theta\right) = 1 - Pr\left(\sum_{i=1}^n X_i \le \Gamma_{n,\theta_0,1-\alpha}\middle|\theta\right)$$

□

Motivation

$$Power(\theta) = Pr(\boldsymbol{X} \in RR) = Pr\left(\sum_{i=1}^{n} X_i > \Gamma_{n,\theta_0,1-\alpha}\middle|\theta\right) = 1 - Pr\left(\sum_{i=1}^{n} X_i \leq \Gamma_{n,\theta_0,1-\alpha}\middle|\theta\right)$$

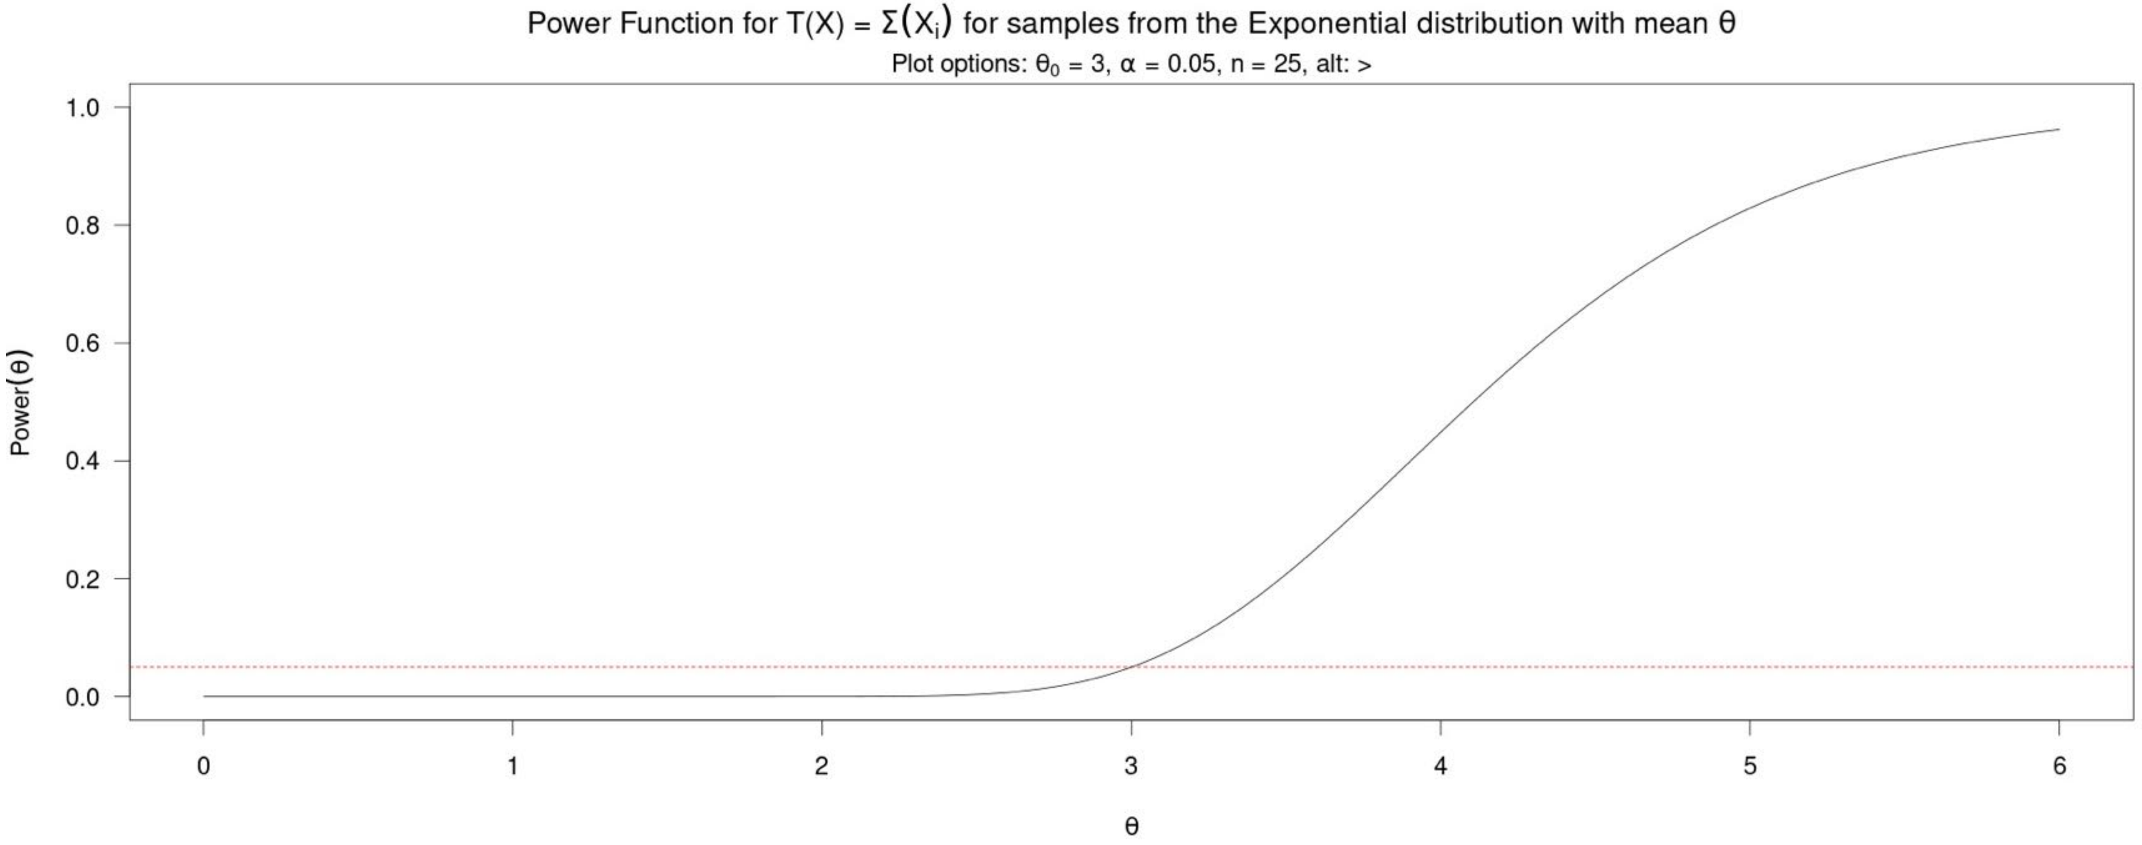

Motivation

- But what does power really mean? How does it connect to other concepts such as sampling distributions?
- Shiny applets can be used to build intuitive relationships with power functions without dedicating class time to multiple derivations
  - Can be accompanied by an activity to focus the exploration

Application goal #1

Understanding how the sample size, null value, alternative hypothesis, significance level, test statistic, and population distribution each affect the power function

The Power of Sampling Distributions

Plots | Derivations

Population Distribution

Exponential

Test Statistic (T(X))

Sum of the X's

Null Value (θ_0)

3

Alternative Hypothesis

Greater than

Alpha Level (α)

0.05

Sample Size (n)

25

Theta (θ)

Introduction

Hello! This application is meant to help visualize power curves, sampling distributions, and the relationship between the two. The panel to the left will allow you to explore the power curve for each of three different population distributions, alternative hypotheses, and test statistics. Furthermore, you may investigate how those curves depend on the significance level, sample size and null value.

This application will also allow you to explore how power is related to the sampling distribution of a test statistic under the null hypothesis and true value of θ. To visualize these distributions for a specific value of θ, simply click the power curve at the desired value of θ, or type the value in the panel. Have fun!

Power Curve

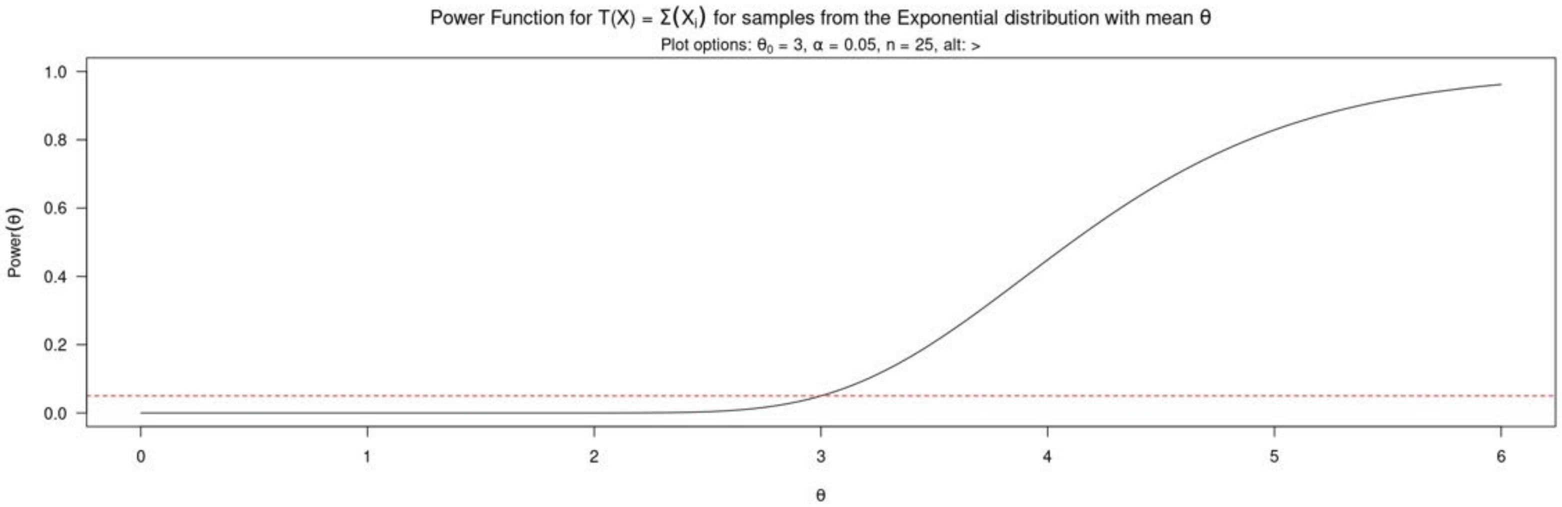

Download this plot

Sampling Distribution

Below are the sampling distributions for the chosen statistic under both the null hypothesis and the true value of θ. The red area corresponds to the significance level and the gray area corresponds to the power; note the relationship between the sampling distribution under θ_0, sampling distribution under θ, significance level and power!

Click a point on the power plot above to visualize the sampling distribution!

Available at: https://shiny.stt.msu.edu/jg/powerapp

Application goal #2

Understanding the relationship between the power function and the sampling distribution of the test statistic under both the null hypothesis and the true value of θ

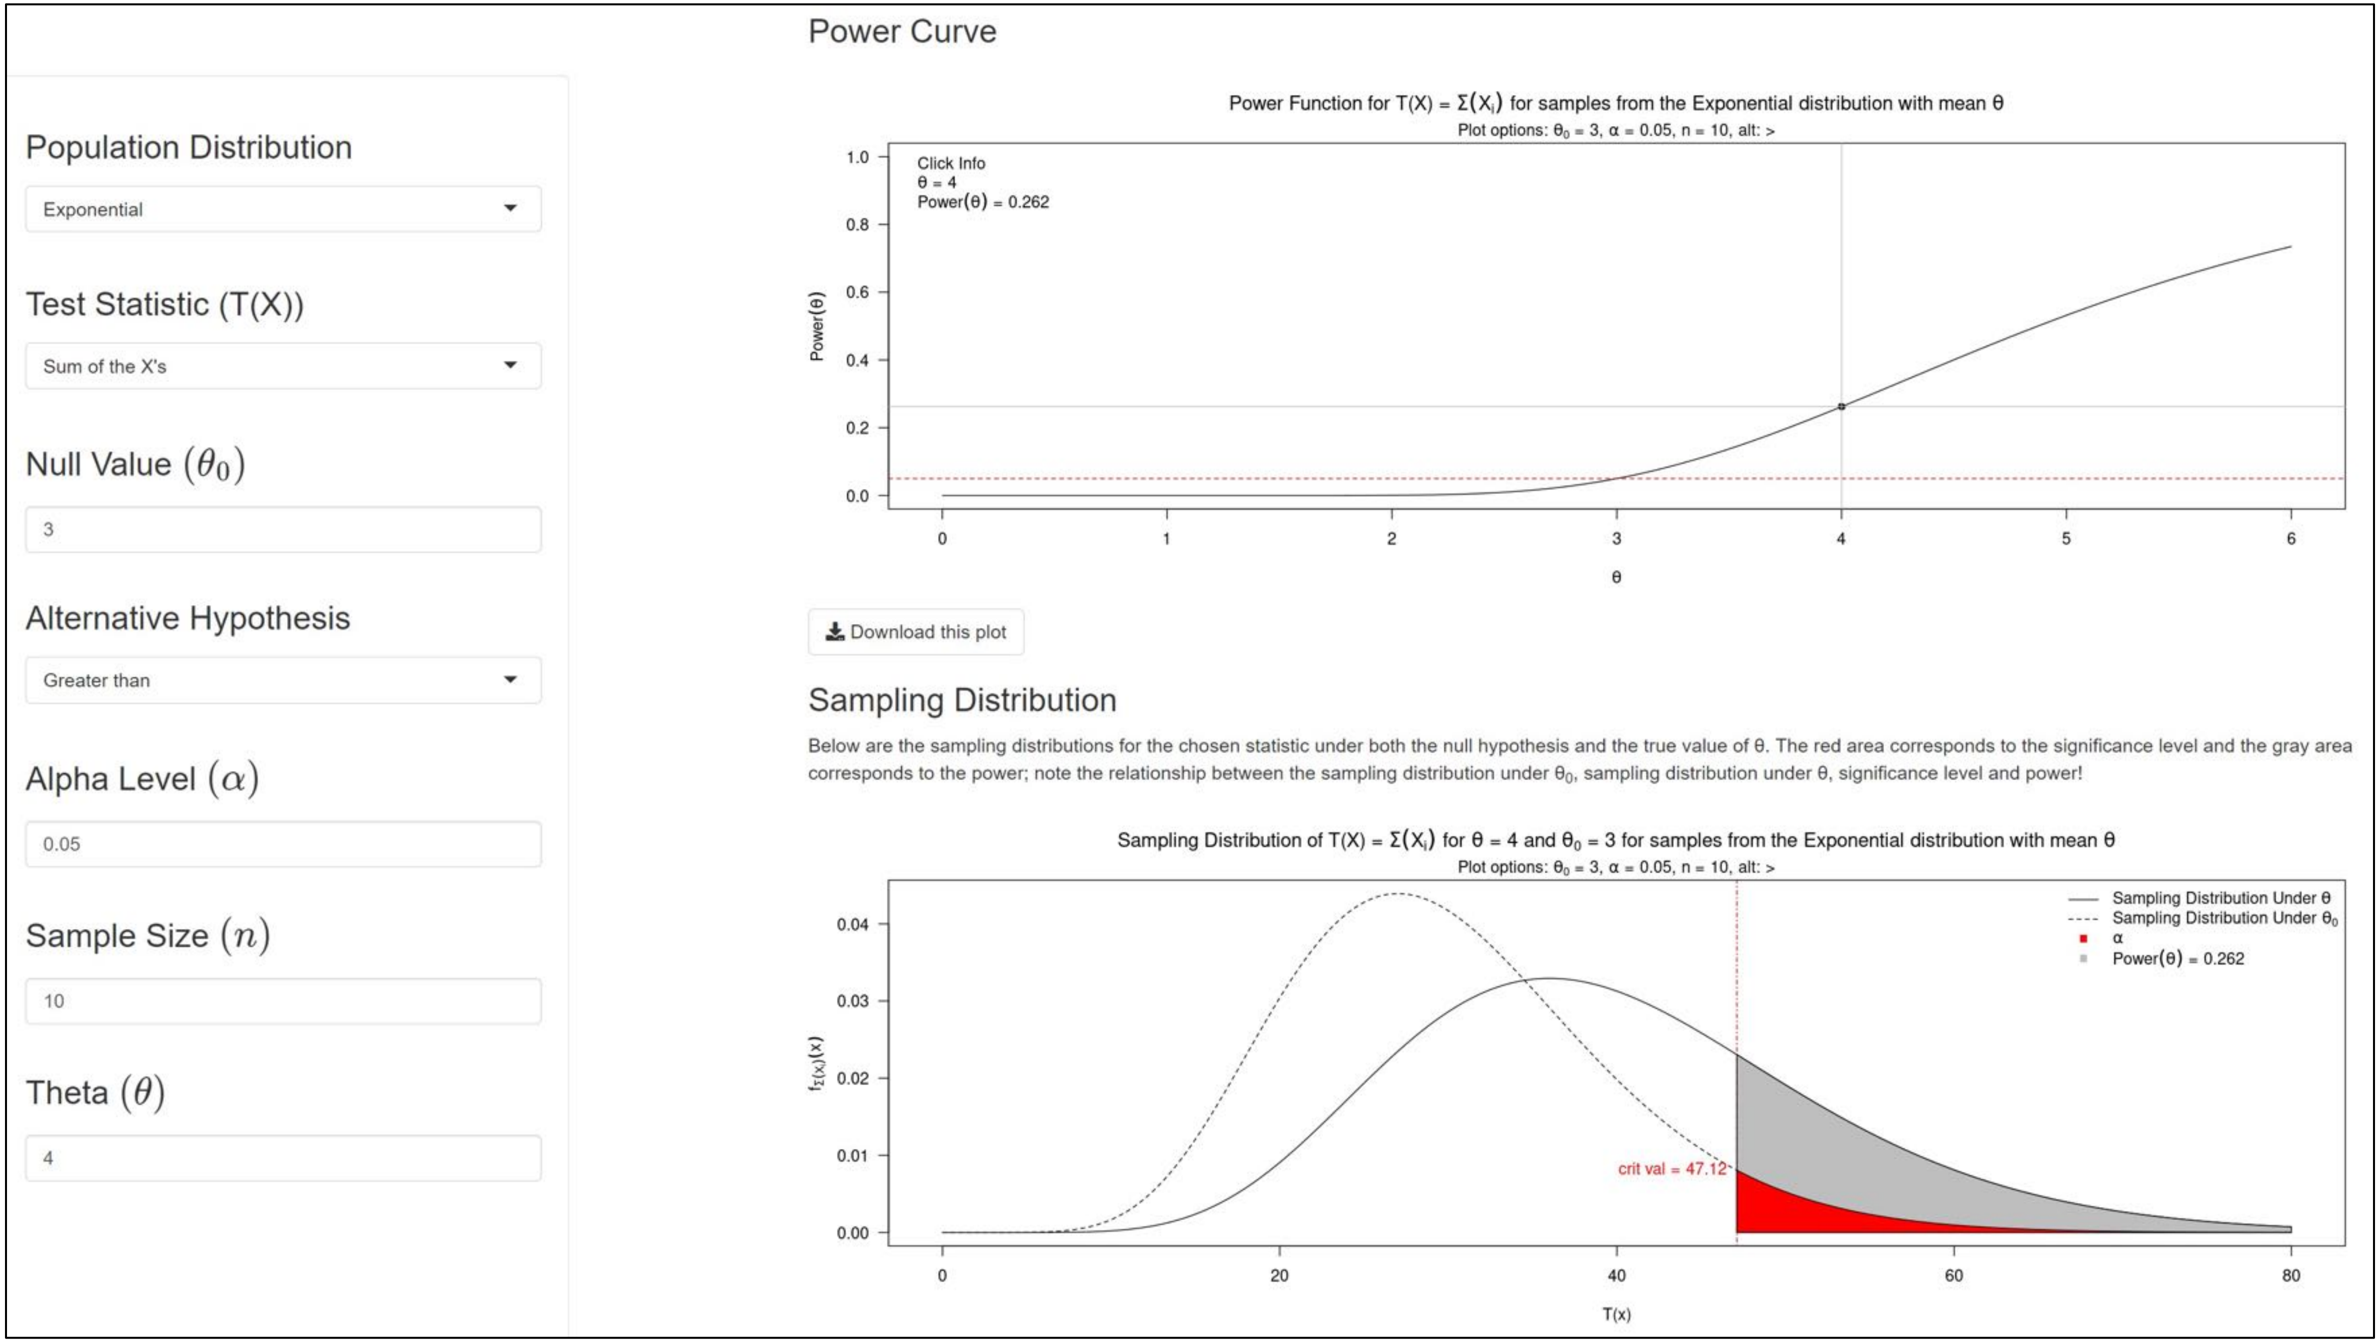

Available at: https://shiny.stt.msu.edu/jg/powerapp

Application goal #3

Understanding how to calculate the power function of a hypothesis test using a variety of strategies, including simulation and normal approximation

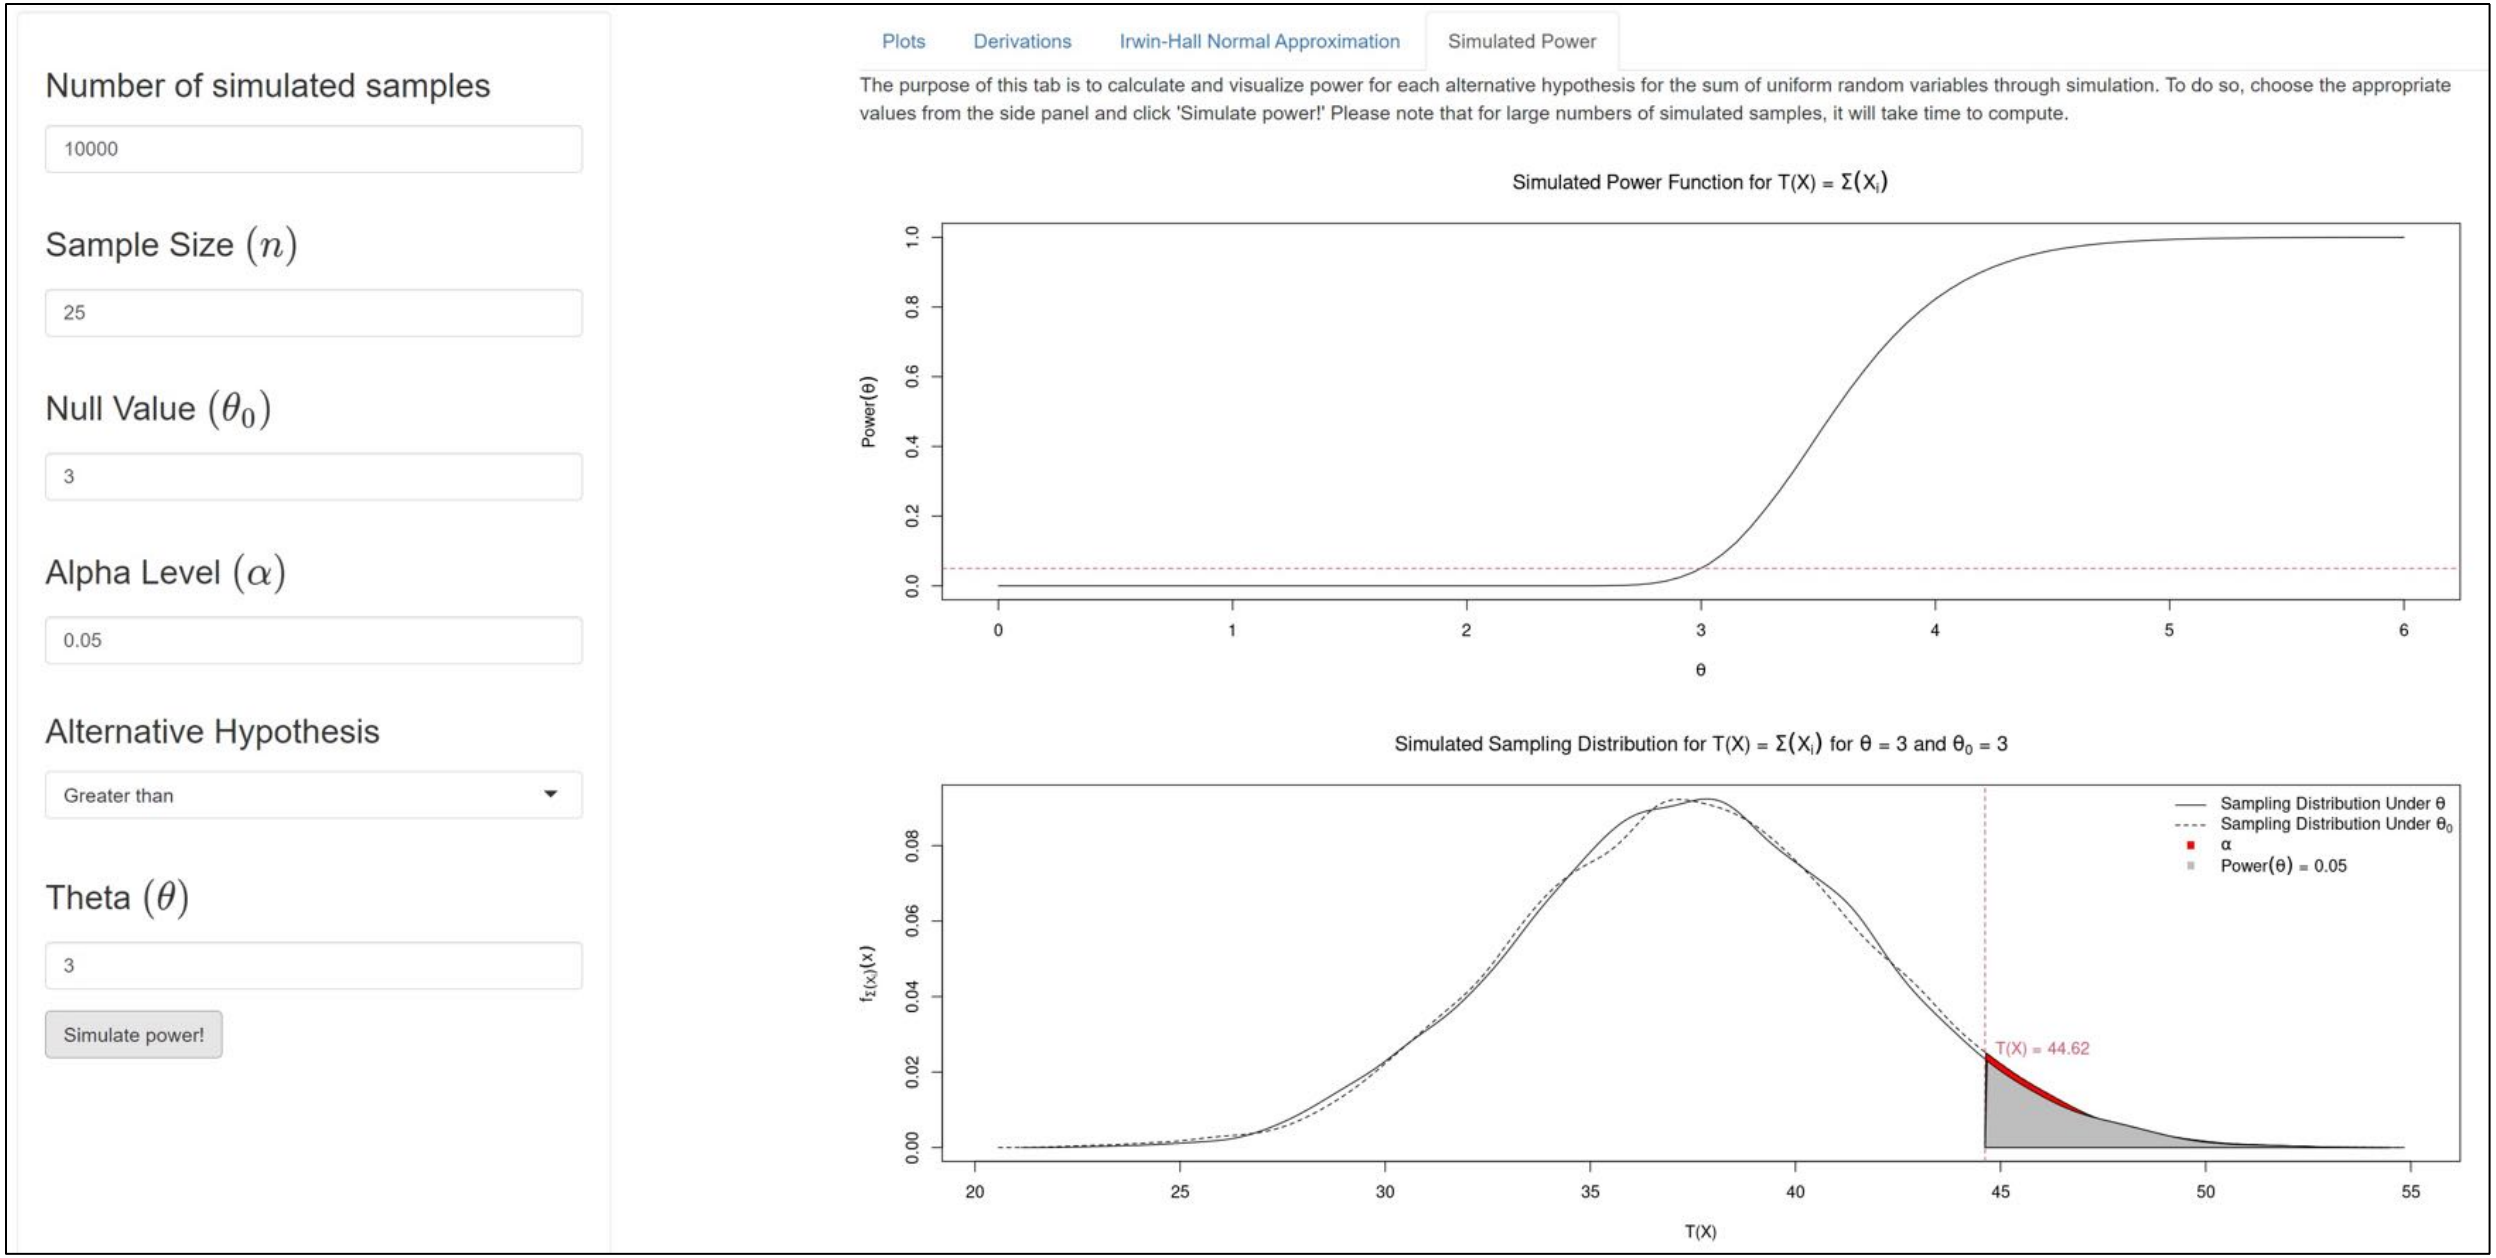

Available at: https://shiny.stt.msu.edu/jg/powerapp

Let's explore!

Using the applet, reflect on how students would engage with the applet and activity.

Relevant links:

- Web application link: https://shiny.stt.msu.edu/jg/powerapp
- Activity link: Power exploration activity

Discussion

- What did you discover and notice?
- How might you use this applet in your classroom?

Additional student insights

- Biased vs unbiased tests
  - uniform distribution, not equal to alternative hypothesis, sample size of one
- Significance level and p-values
- Sufficiency and power
- Numerical instability
  - uniform distribution, sum of random variables

Biased test

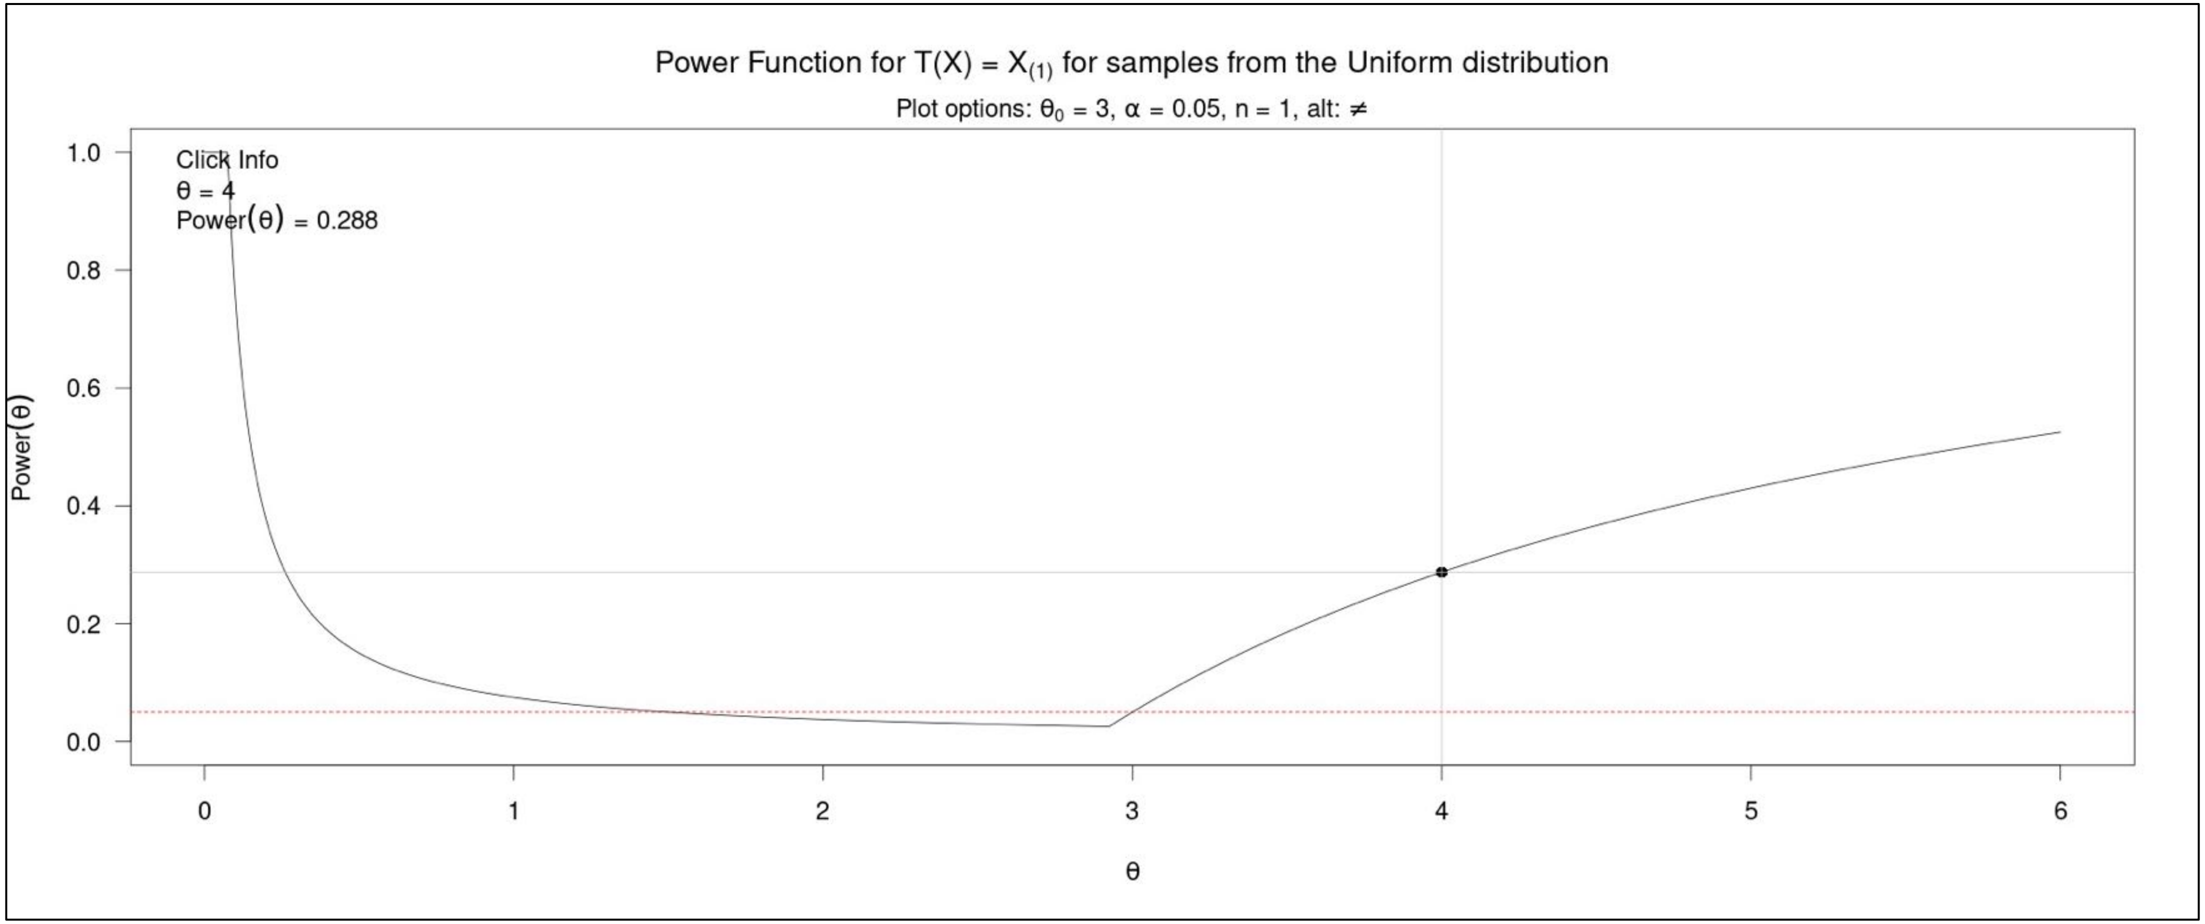

Power Function for T(X) = X(1) for samples from the Uniform distribution
Plot options: θ0 = 3, α = 0.05, n = 1, alt: ≠
Click Info
θ = 4
Power(θ) = 0.288
Power(θ)
θ
1.0
0.8
0.6
0.4
0.2
0.0
0
1
2
3
4
5
6

Numerical instability

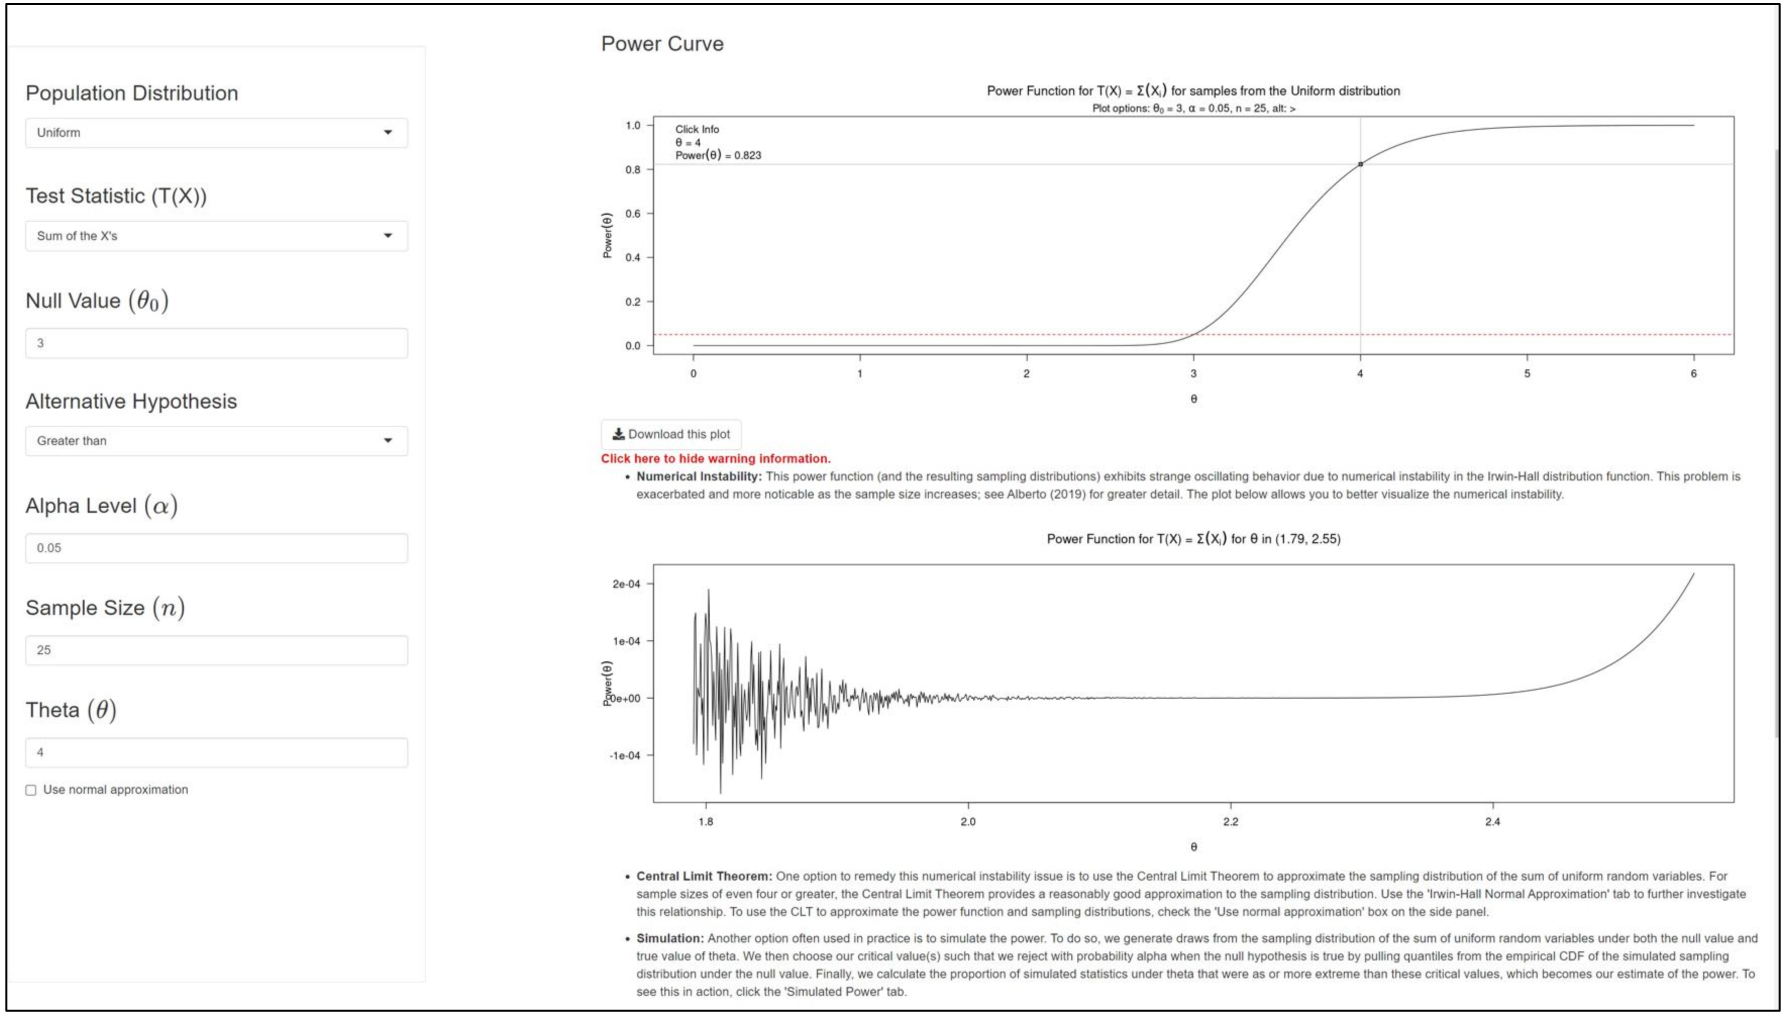

Summary

- Shiny applets can be used to build intuitive relationships with power functions without dedicating class time to multiple derivations
- Use of this web application expands students' opportunities to actively explore and develop conceptual understanding of statistical concepts

Contact Information

Christian Stratton, christianstratton@montana.edu

Jennifer Green, jg@msu.edu