

Statistical Power for Applied Economists

Summer 2019

Nathan Fiala
nvfiala@gmail.com

Prerequisites: This class requires you have taken an introduction to statistics or econometrics class at the masters or PhD level, or equivalent. A masters or PhD level microeconomic theory course is recommended, but not required.

Course description: This class presents the theory and practice behind calculating statistical power in applied economics. We will discuss power in the context of a variety of methods, including randomized control trials, instrumental variables, matching, difference-in-difference and general panel data. If we have time, we will discuss power in the context of behavioral lab experiments.

The class will be practical, with theory and examples included throughout. The examples will come from mostly developing country contexts, though I will try to add Western country examples as well. By the end of the class you will be expected to know how to conduct power analysis and design a study that is appropriately powered.

There is no required textbook, but I expect that students will come prepared to discuss the reading listed below. You should also have looked at the data we will be covering and familiarize yourself with the code.

Main outputs: To receive a grade in the course you will need to produce two main outputs.

First, I expect each person to contribute to the discussions. I will choose randomly someone to discuss one of the readings, so come ready to speak about each of them.

Second, each student will produce their own power calculation report of an existing published study. You will need to identify a published study, obtain the data either through the authors or journals website or by contacting the authors, conduct power analysis on their main outcomes, and write a two-page report detailing the outcomes of this analysis.

Outline: The main topics we will cover include the following:

1. What is power?
2. What determines power?
3. The consequences of underpowered studies
4. Approaches to determining effect sizes and power
5. Multiple hypothesis correction and power in heterogeneity analysis
6. Post hoc power calculations
7. How to report power

Required readings: There are five readings required for the course. Please read each before our first meeting.

1. The power of bias in economics research, by John P. A. Ioannidis, T. D. Stanley and Hristos Doucouliagos
2. Sample size and power calculations, by Andrew Gelman
3. Targeting Next Generations to Change the Common Practice of Underpowered Research, by Rik Crutzen and Gjalt-Jorn Y. Peters
4. Beyond Power Calculations: Assessing Type S (Sign) and Type M (Magnitude) Errors, by Andrew Gelman and John Carlin
5. Power calculation for causal inference in social science: sample size and minimum detectable effect determination, by Eric W. Djimeu and Deo-Gracias Houndolo

In addition, I suggest reading Andrew Gelman's blog (<https://statmodeling.stat.columbia.edu/>). He is a major voice in the concern over power issues in studies.

Data and code: We will use two programs in the course: Stata and Optimal Design (<http://hlmsoft.net/od/>).

The data will come from the paper Christopher Blattman, Nathan Fiala and Sebastian Martinez, "Generating Skilled Self-Employment in Developing Countries: Experimental Evidence from Uganda". Quarterly Journal of Economics, 2014, 129 (2): 697-752. It is available at <https://dataverse.harvard.edu/dataset.xhtml?sessionId=b09da1bd44138069102978fd5c2a?persistentId=doi%3A10.7910/DVN/27898&version=1.0>. Please download the data and code before the start of class. You are not required to read the paper.