Experimental Economics: Theory and Practice

John A. List

U. Chicago, ANU, and NBER

15 July 2024

Abstract

This document is meant to introduce my forthcoming book, titled "Experimental Economics: Theory and Practice," which is to be published in 2025 by The University of Chicago Press. The document first contains the book's outline followed by a Preface, which summarizes my inspiration for writing the book and my goals and aspirations for choosing the content contained in the book.

I. Experimental Methods in Economics

1. Introduction

1.1 Causal Inference

1.2 The Book's Gameplan

2. A Primer on Economic Experiments

2.1 Four Running Examples

2.2 The Empirical Approach in Economics

2.3 Experiments in Economics

2.3.1 Laboratory Experiments

2.3.2 Field Experiments

2.3.2.1 Seven Criteria that Define Field Experiments

2.3.3 What Parameters do the Various Experimental Types Recover?

2.4 What Experimental Type to Choose?

2.4.1 Control Across the Experimental Spectrum for Identification Purposes

2.4.2 Control Across the Experimental Spectrum for Measurement Purposes

2.4.3 The Ability to Replicate Across the Experimental Spectrum

2.4.4 Control Across the Experimental Spectrum for Inferential Purposes

2.4.5 Control Across the Experimental Spectrum to Ensure External Validity

II. Designing Economic Experiments

5. Optimal Experimental Design

6. Randomization Techniques

III. Violations of Exclusion Restrictions

11.SUTVA: Interference and Hidden Treatments

12.Observability: Non-Random Attrition

13.Complete Compliance: One-Sided and Two-Sided Violations

14.Statistical Independence and Compromised Randomization

IV. Building Scientific Knowledge

15. Building Confidence in (and Knowledge from) Experimental Results

16. Generalizability and Scaling

S4. How To: Conduct Experiments to Measure Preferences, Beliefs, and Constraints

S5. How To: Conduct Experiments to Generate Unconventional Data

Preface

In the mid-1980s I began to frequent baseball card conventions on weekends in Madison, WI. At these conventions both consumers and dealers would buy, sell, and trade sportcards and memorabilia. Before long, I found myself as a sportscard dealer, earning pocket money by doing 'accidental experiments' at these conventions as an undergraduate student at the University of Wisconsin-Stevens Point.

In my studies as an undergraduate, I learned about fascinating theories in sociology, economics, and psychology that I would watch unfolding before my very eyes at these sportscard conventions. Dynamics of family interactions, the psychology of giving something up, and the bargaining and negotiating patterns I observed at these conventions were at once fascinating and at times puzzling because they did not always subscribe to the theories I was learning in my undergraduate courses.

Maybe certain theoretical assumptions were wrong, perhaps the entire theory needed to be discarded? Other questions also had me pondering the usefulness of science: How can the market perform the work of thousands of people without nary a conductor, usher, shepherd, or guide? What is the foundation of these markets and what assumptions were necessary to make them work efficiently? How do these markets clear and set the going price? What psychological role does loss play in negotiations?

Such economic and behavioral questions occupied my nights alongside the other questions that kept me awake most evenings until light started my next day: what is the meaning of life, is the universe deterministic, and what is consciousness?

I would return to my economics courses pondering whether there was a way to add clarity to the market puzzles I was both creating in my mind and observing in the real world. My learnings seemed simple: economics is a science that can contribute to understanding and resolving important theoretical and policy debates. The generality of the economics framework and its ability to hone-in on specific aspects of social problems within quite complicated settings attracted me to the field. The study of how incentives affect choice, and how those choices affect outcomes and society at large attracted me to experimentation.

Indeed, the sportscard world I would frequent nearly every weekend was extremely complicated, thousands of bustling buyers and sellers and hundreds of different prices and other incentives at work. Auctions, bilateral negotiations, price discrimination, collusion, and features of preferences all seemed to be touch points that could be explored. I reasoned that one possible way of adding crisp insights into the agents' behaviors and the institutions of this market was to go beyond the accidental experiments I was conducting and perform controlled experiments with

a purpose, like those I was learning about in chemistry, physics, and biology as an undergraduate student.

I began to implement "experimental" strategies in the late 1980s, rotating from posted price to various auctions to sell my cards, "piloting" new bargaining strategies to buy cards, and sequencing and framing my goods in new ways to explore how those changes impacted sales and profits. Maybe the choice architecture mattered, I reasoned, but the economic theory I was learning did not have that as a key consideration, at least explicitly. Everything seemed to be about preferences, beliefs, and constraints in my courses—I wondered if I could use those primitives to understand the role of architecture.

As I learned about using scientific principles and implementing them, I also earned more in the marketplace. When something worked, I explored why; when I learned about new theories in the classroom, I brought them to my test tube in the field. This was a true two-way street: theory to data and data back to theory.

While satisfying, in hindsight, I could have learned a lot more about the economic science had I understood principles of optimal experimental design, data generation, analysis, establishing causality, building knowledge, and interpretation. Therein lies the genesis of this book. In the early 1990s, I sat in my first course in Experimental Economics at the University of Wyoming as a graduate student. The course contained 100% laboratory experiments, which made sense because that was the lay of the land back then.

The main textbook was written by Doug Davis and Charlie Holt, and the course was supplemented by John Kagel and Al Roth's handbook. These materials were excellent; combined with the supplementary readings an entire new area of study was now at my fingertips. In fact, these materials charted a new course for my academic career and disciplined my thinking. Building on the shoulders of these giants I learned from, I had something different in mind: I wanted to use the sportscard conventions as my laboratory, as my glimpse into the realm of economics. After exhausting learnings in that market, I had wild thoughts of using the world as my lab someday.

My professional research career was started. My career has been importantly complemented by teaching experimental economics to both undergraduate and graduate students. I have taught some version of Experimental Economics since 1996, when I became an Assistant Professor at the University of Central Florida. In these early years the course was mainly describing laboratory experiments, with a few of my own field experiments that I conducted in the early 1990s sprinkled in. By the time I reached the University of Arizona in 2000, half the course revolved around field experiments, albeit most of them working papers in areas as diverse as auction theory, environmental economics, charitable giving, and discrimination.

From 2001-05, I taught both undergraduate and graduate Experimental Economics at the University of Maryland, and through these years added field experiments in areas spanning behavioral economics and finance, gender and race outcomes in markets, education, and experimental methodologies in the field. Now, a Professor at the University of Chicago, my experimental economics courses have evolved into primarily focusing on field experiments with

laboratory experiments sprinkled in. This is merely a matter of taste, as there are plenty of excellent laboratory experiments conducted every year that could fill several semester-long courses. Indeed, this book will relentlessly highlight the wonderful complementarities between laboratory and field experiments.

While over the past decades the classroom has certainly honed my pedagogical strategies, perhaps my experiences in the field have lent the greatest insights into my classroom teachings and the approach in this book, which represents a scientific compilation from every facet of my professional life. This includes the wisdom gained over the past 30 years as a teacher, researcher, White House Presidential Advisor, Chief Economist at Uber, Lyft, and Walmart, and an Economic Advisor to governments around the globe.

My learning has largely been autodidactic and that represents a key impetus for writing this text. In this spirit, my goal in this book is to provide an A-to-Z compendium that will take the student from the ground floor to designing, conducting, analyzing, and interpreting data that they generate themselves. My approach is to highlight causal inference without losing the beauty, science, and art of the experimental approach in economics. We have learned much over the past several decades on how received experimental results depend on various design and implementation choices. In short, throughout the text I will uncompromisingly return to how seemingly innocuous design choices and subtle implementation decisions impact behavior.

In this manner, the book is meant to prepare the scholar to develop and implement their own experimental projects in the laboratory and field. This book represents what I needed in the late 1980s and early 1990s as I toiled around memorabilia markets trying to learn about fundamental economic tenets. I do not want the reader to spend as much time and effort as I have over the past thirty+ years trying to master the experimental approach, with all its nuances. I have made enough experimental mistakes over the years to fill a large tome. This book is structured to eliminate, or at least attenuate, such mistakes, as well as avoid the limitless others that lurk.

The book, while theoretically principled, is designed for the applied researcher, and therefore focuses on practical rather than mathematical results. My presentation is at a level that presupposes the student has had beginning courses on statistical inference and regression. While the book is designed as the main resource for a stand-alone course, I can imagine that it could serve as a nice companion for courses in establishing causality in the social sciences, econometrics, research design, statistics, and any applied research course in general.

I have attempted to write in a manner that serves PhD students, Master's students, and the astute undergraduate student. To do so, I change writing styles systematically: when in formal parts of the book I remain formal. When the ideas demand a more conversational tone, I supply that style of writing. In pilot experiments, I have found that such a change in writing style maximizes engagement of the student.

I have endeavored to choose features that augment the learnings in each chapter. For example, each chapter begins with a *Mind's Eye*, a call to the mental faculties of the student. There is a fair amount of heterogeneity in how the learner forms mental images, but humans have a tendency to learn more deeply when content is complemented by mental pictures. This feature

calls upon that fact. When creating such a learning device, it is important to note that no chapter is the same; some chapters lean theoretical whereas others are quite practical. For example, one chapter may walk us slowly through a framework, teaching the assumptions and restrictions required for proper causal inference. Another may be more applied, providing concrete methods to ensure optimal design or appropriate analysis.

As such, there is no consistent pattern for the *Mind's Eye*. Instead, I choose a distinct color and shape for each type of learning. One feature that is consistent: each *Mind's Eye* begins with a red circle. I view this as the "base" or the central point of consideration in the chapter. The base might connect to any one of five different shapes, defined by the following legend:

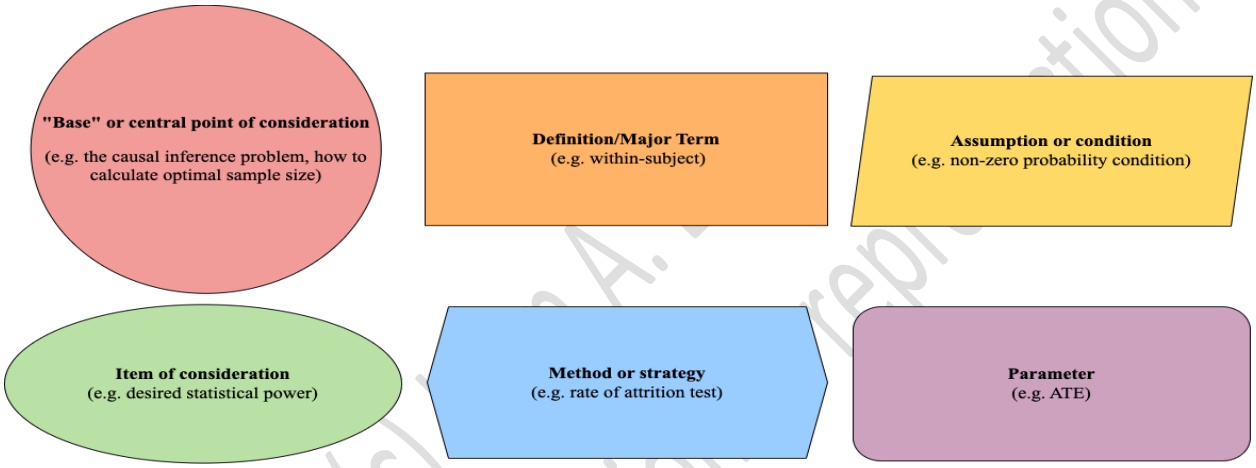

Thus, the *Mind's Eye* is designed in a systematic way to set the tone for the materials that the student should expect to grasp from the chapter. In focus groups, I have learned that with a unique box for each takeaway, the *Mind's Eye* simultaneously provides the student a roadmap for where the chapter is headed while also offering a simple framework that the student can refer back to when reading the chapter.

Each chapter also includes a few running examples from the literature that help the student understand how the learnings of the chapter fit with frontier research. These are drawn primarily from laboratory and field experiments published in premier peer-reviewed economics journals but occasionally I draw upon working papers. My presentation of the research in the running examples is not meant to be exhaustive or as an endorsement of a particular set of findings. Rather, it is meant as an introduction that the student can use to place the chapter's learnings into perspective. That said, I have found that when the student studies the actual paper, the learnings of the chapter become even more readily available. Where induction ends deduction begins.

To complement these running examples, I include *Further Readings* at the end of each chapter. These represent a sampling of my favorite examples of work that is topically related to the contents in the chapter. These are provided to help readers discern where to go next for extensions or useful references in the literature.

Finally, each chapter contains at least one *Choice and Consequence* and one *Letting the Data Speak* learning moment. These are meant to augment the chapter learnings by using current empirical examples to highlight the key insights in the chapter. Specifically, the *Choice and Consequence* boxes take an inductive approach by highlighting how design and implementation choices are, at times, astonishingly important to observed behaviors. These capture the methodological beauty, and art, of the experimental approach. Alternatively, the *Letting the Data Speak* cases emphasize key results that reinforce specific chapter learnings and underscore lessons to be drawn from the materials.

In the end, my hope is that with one reading the student can become not only the teacher, but the master craftswoman (craftsman) in generating data to improve the lives of generations. As economists we have waited too long to use the world as our lab.

Acknowledgements

Academia can be a very lonely place. Fortunately, I have been blessed over the years with hundreds of brilliant coauthors, superb mentors, dozens of friends, and chiliads of bright students. I suspect that each of them will find a piece of what they have taught me in this text. I want to thank you all, which number too great to name one by one. Equally as important to the actual creation of this book is my incredible University of Chicago research team. Justin Holz, Faith Fatchen, Uditi Karna, David Novgorodsky, David Franks, and Michael Cuna were incredible research assistants and thought partners. And, in random order, Faisal Kattan, Francesca Pagnotta, Charles Shi, Jonathan Lambrinos, Lina Ramirez, Ariel Listo, Min Sok Lee, Panos Toulis, Diana Smith, Andrew Simon, Jan Luca Henkel, Max Hippold, Beat Hintermann, Ragan Petrie, Marco Castillo, Anna Dreber, Joshua Dean, three anonymous reviewers, and many others also helped to shape this book. Indeed, every one of them has played an invaluable role in preparing, debating, and finalizing various chapters in the book. I will be forever indebted. Three reviewers, Marco Castillo, and Ragan Petrie provided insightful comments on an earlier version of this book that improved the contents considerably.