Experimental Design and Analysis

Howard J. Seltman

November 1, 2009

Preface

This book is intended as required reading material for my course, Experimental Design for the Behavioral and Social Sciences, a second level statistics course for undergraduate students in the College of Humanities and Social Sciences at Carnegie Mellon University. This course is also cross-listed as a graduate level course for Masters and PhD students (in fields other than Statistics), and supplementary material is included for this level of study.

Over the years the course has grown to include students from dozens of majors beyond Psychology and the Social Sciences and from all of the Colleges of the University. This is appropriate because Experimental Design is fundamentally the same for all fields. This book tends towards examples from behavioral and social sciences, but includes a full range of examples.

In truth, a better title for the course is Experimental Design and Analysis, and that is the title of this book. Experimental Design and Statistical Analysis go hand in hand, and neither can be understood without the other. Only a small fraction of the myriad statistical analytic methods are covered in this book, but my rough guess is that these methods cover 60%-80% of what you will read in the literature and what is needed for analysis of your own experiments. In other words, I am guessing that the first 10% of all methods available are applicable to about 80% of analyses. Of course, it is well known that 87% of statisticians make up probabilities on the spot when they don't know the true values. :)

Real examples are usually better than contrived ones, but real experimental data is of limited availability. Therefore, in addition to some contrived examples and some real examples, the majority of the examples in this book are based on simulation of data designed to match real experiments.

I need to say a few things about the **difficulties of learning** about experimental design and analysis. A practical working knowledge requires understanding many concepts and their relationships. Luckily much of what you need to learn agrees with common sense, once you sort out the terminology. On the other hand, there is no ideal logical order for learning what you need to know, because everything relates to, and in some ways depends on, everything else. So be aware: many concepts are only loosely defined when first mentioned, then further clarified later when you have been introduced to other related material. Please try not to get frustrated with some incomplete knowledge as the course progresses. If you work hard, everything should tie together by the end of the course.

In that light, I recommend that you create your own "concept maps" as the course progresses. A concept map is usually drawn as a set of ovals with the names of various concepts written inside and with arrows showing relationships among the concepts. Often it helps to label the arrows. Concept maps are a great learning tool that help almost every student who tries them. They are particularly useful for a course like this for which the main goal is to learn the relationships among many concepts so that you can learn to carry out specific tasks (design and analysis in this case). A second best alternative to making your own concept maps is to further annotate the ones that I include in this text.

This book is on the world wide web at http://www.stat.cmu.edu/∼hseltman/309/Book/Book.pdf and any associated data files are at http://www.stat.cmu.edu/∼hseltman/309/Book/data/.

One key idea in this course is that you cannot really learn statistics without doing statistics. Even if you will never analyze data again, the hands-on experience you will gain from analyzing data in labs, homework and exams will take your understanding of and ability to read about other peoples experiments and data analyses to a whole new level. I don't think it makes much difference which statistical package you use for your analyses, but for practical reasons we must standardize on a particular package in this course, and that is SPSS, mostly because it is one of the packages most likely to be available to you in your future schooling and work. You will find a chapter on learning to use SPSS in this book. In addition, many of the other chapters end with "How to do it in SPSS" sections.

There are some typographical conventions you should know about. First, in a non-standard way, I use capitalized versions of Normal and Normality because I don't want you to think that the Normal distribution has anything to do with the ordinary conversational meaning of "normal".

Another convention is that optional material has a gray background:

I have tried to use only the minimally required theory and mathematics for a reasonable understanding of the material, but many students want a deeper understanding of what they are doing statistically. Therefore material in a gray box like this one should be considered optional extra theory and/or math.

Periodically I will summarize key points (i.e., that which is roughly sufficient to achieve a B in the course) in a box:

Key points are in boxes. They may be useful at review time to help you decide which parts of the material you know well and which you should re-read.

Less often I will sum up a larger topic to make sure you haven't "lost the forest for the trees". These are double boxed and start with "In a nutshell":

In a nutshell: You can make better use of the text by paying attention to the typographical conventions.

Chapter 1 is an overview of what you should expect to learn in this course. Chapters 2 through 4 are a review of what you should have learned in a previous course. Depending on how much you remember, you should skim it or read through it carefully. Chapter 5 is a quick start to SPSS. Chapter 6 presents the statistical foundations of experimental design and analysis in the case of a very simple experiment, with emphasis on the theory that needs to be understood to use statistics appropriately in practice. Chapter 7 covers experimental design principles in terms of preventable threats to the acceptability of your experimental conclusions. Most of the remainder of the book discusses specific experimental designs and corresponding analyses, with continued emphasis on appropriate design, analysis and interpretation. Special emphasis chapters include those on power, multiple comparisons, and model selection.

You may be interested in my background. I obtained my M.D. in 1979 and practiced clinical pathology for 15 years before returning to school to obtain my PhD in Statistics in 1999. As an undergraduate and as an academic pathologist, I carried

out my own experiments and analyzed the results of other people's experiments in a wide variety of settings. My hands on experience ranges from techniques such as cell culture, electron auto-radiography, gas chromatography-mass spectrometry, and determination of cellular enzyme levels to topics such as evaluating new radioimmunoassays, determining predictors of success in in-vitro fertilization and evaluating the quality of care in clinics vs. doctor's offices, to name a few. Many of my opinions and hints about the actual conduct of experiments come from these experiences.

As an Associate Research Professor in Statistics, I continue to analyze data for many different clients as well as trying to expand the frontiers of statistics. I have also tried hard to understand the spectrum of causes of confusion in students as I have taught this course repeatedly over the years. I hope that this experience will benefit you. I know that I continue to greatly enjoy teaching, and I am continuing to learn from my students.

Howard Seltman
August 2008

Contents