

Statistical Methods in Online A/B Testing

Statistical Methods in Online A/B Testing

Georgi Z. Georgiev

Statistical Methods in Online A/B Testing

This book is dedicated to explaining the tools of statistical inference and estimation through online controlled experiments; a.k.a. A/B tests. It views them in a risk-management context of balancing the risks and rewards of innovation. With the help of this text, user experience and conversion rate optimization practitioners will be able to harness the power of data-driven decision-making, and enable their business to innovate, while controlling the risk to which it is exposed.

An issue with much of the current statistical theory and practice in online A/B testing is that it is misguided and frequently misinterpreted. It is often applied without a good understanding of the role and limitations of statistical methods, instead blindly copying scientific applications without due consideration of many of the unique features of online business. This book approaches this problem by laying solid statistical foundations, and providing clear definitions and a multitude of practical examples, while constantly keeping an eye on the overarching business goal of A/B testing. By making constant use of the business context, and ample practical examples, this text presents A/B testing statistics in a uniquely useful way.

Georgi Z. Georgiev is the Managing Director of Web Focus LLC, and a veteran web marketer and web developer. His diverse 15-year experience includes owning, developing, and managing dozens of successful online projects, working as an SEO, Google AdWords, Google Analytics, and statistics consultant, as well as delivering training and lectures on multiple seminars and events, including in his capacity as a Google Regional Trainer. He is also developer of statistical tools at Analytics-Toolkit.com, as well as the author of numerous articles and white papers on the topic of statistics in online A/B testing. His most notable works have been “Efficient A/B Testing in Conversion Rate Optimization: The AGILE Statistical Method”, and the first extensive glossary of statistical terms in online A/B testing. His vast experience with online business and statistics positions him uniquely to deliver an accessible book on the topic of statistics applied to online A/B tests.

TABLE OF CONTENTS

MOTIVATION..... 8

WHO IS THE READER OF THIS BOOK? 11

ACKNOWLEDGMENTS 12

NOTATIONS..... 13

1. USING STATISTICS IN BUSINESS 14

 1.1. When are statistics useful?..... 14

 1.2. Statistical inference in business..... 19

 1.3. Primary uses of statistics in online A/B testing..... 21

 1.4. The necessity for counterfactual reasoning 25

 1.5. Establishing causality 29

 1.6. Ruling out alternative explanations for the data..... 34

 1.7. Statistical methods and efficient use of data 35

 1.8. Caveats in using statistical methods in business..... 37

2. ESTIMATING UNCERTAINTY. STATISTICAL SIGNIFICANCE, P-VALUES, OTHER ESTIMATES. 39

 2.1. Substantive hypotheses 39

 2.2. Statistical hypotheses 41

 2.3. Standard deviation and Z-Scores..... 46

 2.4. *p*-value and type I errors..... 54

 2.5. *p*-value: utility and interpretation 59

 2.6. Confidence intervals..... 63

 2.7. Misinterpretations of *p*-values and confidence intervals..... 69

 2.8. *p*-values and confidence intervals in decision-making..... 73

 2.9. Maximum likelihood estimate 75

Statistical Methods in Online A/B Testing

2.10.	Severity	77
3.	STATISTICAL ASSUMPTIONS. ASSESSING MODEL ADEQUACY	80
3.1.	The importance of statistical assumptions	80
3.2.	Probabilistic assumptions of statistical models.....	82
3.3.	Probabilistic assumptions in different practical cases	84
3.4.	Assumptions imposed by the design of the experiment	86
3.5.	Assessing statistical adequacy through statistical tests.....	88
3.6.	Assessing statistical adequacy through A/A tests	89
4.	STATISTICAL POWER AND SAMPLE SIZE CALCULATIONS	92
4.1.	Errors of the second kind (type II errors)	92
4.2.	Statistical power	95
4.3.	The role of statistical power in A/B testing.....	101
4.4.	Minimum effect of interest vs. minimum detectable effect ..	103
4.5.	Underpowered and overpowered tests.....	106
4.6.	Sample size calculations	108
4.7.	False positive and false negative rates.....	118
4.8.	Misunderstandings about statistical power	119
5.	TYPES OF STATISTICAL HYPOTHESES	121
5.1.	One-sided tests and confidence intervals.....	121
5.2.	Misconceptions about one-sided tests.....	125
5.3.	Strong superiority tests	128
5.4.	Non-inferiority tests.....	131
5.5.	<i>p</i> -value notation	135
6.	TESTS WITH MORE THAN ONE VARIANT	137
6.1.	Type I error in an A/B/n test	137

Statistical Methods in Online A/B Testing

6.2.	<i>p</i> -value and CI corrections for testing multiple variants	141
6.3.	Sample size calculations for A/B/n tests	145
6.4.	Dynamically dropping or adding variants.....	149
6.5.	Do factorial designs make sense in online A/B testing?	151
6.6.	Testing the perfect shade of blue	157
7.	SEGMENTATION AND MULTIPLE PERFORMANCE INDICATORS	160
7.1.	The Šidák correction	160
7.2.	Analyses of segments of the sample of an online experiment	161
7.3.	Designs with more than one primary parameter of interest	164
7.4.	Designs with secondary parameters of interest.....	167
8.	WORKING WITH CONTINUOUS DATA	169
8.1.	Standard deviation, <i>p</i> -values, and confidence intervals.....	169
8.2.	Statistical power and sample size calculations	171
8.3.	Is the Normal distribution assumption adequate?	174
8.4.	A workaround for incomplete ARPU data	178
9.	PERCENTAGE CHANGE	181
9.1.	Percentage change (lift) vs. absolute change	181
9.2.	Confidence intervals for percentage change	186
9.3.	<i>p</i> -values for percentage change	187
9.4.	Sample size calculations for percentage change	188
10.	SEQUENTIAL TESTING: CONTINUOUS MONITORING OF DATA	189
10.1.	The issue of repeated significance tests on accumulating data	190

Statistical Methods in Online A/B Testing

10.2.	The sequential probability ratio test	193
10.3.	Fixed analysis time group sequential trials	196
10.4.	Alpha-spending functions and efficacy boundaries	201
10.5.	Beta-spending and futility boundaries	205
10.6.	Expected sample size and efficiency of sequential A/B tests 208	
10.7.	Estimation following a group sequential A/B test.....	218
11.	OPTIMAL SIGNIFICANCE THRESHOLDS AND SAMPLE SIZES 224	
11.1.	Defining “success” for a business experiment	225
11.2.	Costs, benefits, risks, and rewards in A/B testing	229
11.3.	Test parameters and their relationship to costs and benefits 233	
11.4.	Distribution of expected effect sizes	237
11.5.	Calculating risk/reward ratios and key points	242
11.6.	Testing with 50% confidence threshold?	248
11.7.	Inherent cost of A/B testing.....	251
11.8.	Limitations of Risk/Reward calculations	255
12.	EXTERNAL VALIDITY a.k.a. GENERALIZABILITY OF A/B TEST RESULTS.....	257
12.1.	What is external validity (generalizability)?	257
12.2.	Threats to the external validity of online experiments	259
12.3.	Improving the generalizability of A/B test results	263
12.4.	Representative samples and sequential tests.....	265
12.5.	Running multiple concurrent A/B tests	267
13.	MISCELLANEOUS TOPICS.....	273

Statistical Methods in Online A/B Testing

13.1.	Equal or unequal allocation between test groups?	273
13.2.	Holdout groups.....	276
13.3.	Time to event analysis. Hazard ratio	278
13.4.	Meta-analyses of A/B test results	283
13.5.	Adaptive Designs	286
13.6.	A word on multi-armed bandits.....	289
13.7.	Bayesian methods	291
14.	COMMUNICATING STATISTICAL RESULTS.....	296
14.1.	Changing the perception of data variability	296
14.2.	Translating business questions into statistical models.....	299
14.3.	Presenting statistical results	302
	EPILOGUE	315
	REFERENCES.....	316
	INDEX.....	327

MOTIVATION

The most straightforward way to explain why this book exists is via a brief description of my journey from a statistical know-nothing to an author of a book on statistics.

Some years back, I set out to learn more about the application of state-of-the-art scientific methods to the business world of data-driven decision-making. Starting with analyses of observational data, I quickly shifted my focus to online controlled experiments. These are commonly referred to as A/B tests, or split tests.

At the time, I had no formal training in statistics, and only college-level understanding of mathematics, so, to be honest, I didn't even know where to start! Available books on A/B testing barely had anything to say about statistics, so I started reading online blog posts and educational resources from universities, such as online lectures and courses, as well as the odd scientific paper.

In doing so, I had to face an entirely **new jargon** full of counterintuitive terms, such as 'statistical significance', which has little to do with significance, 'statistical power' which has nothing to do with power in the casual sense, 'confidence interval', which has nothing to do with any kind of confidence, and so on. And, above all else, I had to familiarize myself with a notation full of small and capital Greek letters (α , β , γ , δ , θ , μ etc.), which would sometimes mean different things in different contexts, while different letters would also denote the same concept.

To make things worse, there were ample examples of **vague or conflicting information**. There were dozens of definitions for what a p-value is and how it should be interpreted. Almost nobody seemed to care to define what a family of hypotheses is supposed to be when discussing the Family-Wise Error Rate. One source would claim one-tailed tests are preferable, while others would swear by two-tailed tests, and scare you with the heavens coming down on you if you were so reckless as to consider a one-tailed test.

Statistical Methods in Online A/B Testing

Practitioners and academics alike were battling over which approach is best overall, or squabbling over the merits of particular applications - frequentist inference vs. decision-theoretic vs. Bayesian approaches. To make matters even more confusing, there seemed to be noticeable schisms within each school of thought.

Most confusing of all, statistics as such turned out to be very context-dependent - it meant different things in different scientific and business fields. Practitioners in those fields had, over time, developed somewhat separate branches of statistics. Therefore, statistics would mean something different for you depending on whether you come from physics, medicine, social studies, econometrics, environmental studies, or industrial quality control.

It was simply a nightmare attempting to navigate this fractured jungle of jargon, conflicting stances, and math-heavy explanations. Yet, I persevered! And through painstaking reading, practice, implementing/coding methods, and countless simulation runs, I was able to garner a good enough understanding of the matter to begin writing methodological white papers and in-depth articles, to start delivering lectures and courses on statistics in A/B testing, and to become a developer of statistical tools.

From my current position, I see both the immense **value** of statistical methods applied to business risk management, estimation and prediction problems, and the immense **harm** done by improper applications or misguided understanding of those same methods. Thus, in-depth explanations of the practical application of statistical methods, as well as common errors and how to avoid them, are key elements of this work.

Furthermore, in 2019 the difficulties that I went through are about as severe as they were a few years before, despite the valiant efforts of some in the statistics and A/B testing communities. Addressing common mistakes, misconceptions, and misapplications of statistical methods is, therefore, a central part of this work.

My aim with this book is to **carve a clear path through the statistical jungle**, and thus save the reader weeks, months or even years of wandering around in circles, falling into gorges, and crossing rivers, metaphorically speaking! While the book does use the established jargon, each term is explained with painstaking detail and accuracy using the simplest language possible. Math and formulas are kept to a sanitary minimum in order to facilitate reading, while also satisfying the needs of technically-inclined readers, who will also find the detailed references supporting each chapter to be particularly useful. Since this is a book on statistical methods, and not on decision theory per se, the text also sticks to the frequentist error-statistical approach, and only briefly touches on current decision-theoretic and Bayesian methods.

WHO IS THE READER OF THIS BOOK?

This book aims to introduce the complex topic of statistical estimation and inference to readers with somewhere between little and no mathematical and statistical background. The text makes few assumptions, and builds each topic from the ground up, explaining the rationale behind each concept, and following it with a multitude of practical examples from the world of online A/B testing. It contains detailed explanations, so that one can understand the statistical methods deeply enough in order to correctly put them into practice, but steers clear of some of the difficult parts of set theory, calculus, etc., which are typical of many other books on statistics.

A background in conversion rate optimization, operating an online business or mobile app, design of user experiences, or similar, would be helpful for an easier reading, as the primary audience for this work is conversion rate optimization professionals who design, execute, and analyze A/B tests in an online environment. However, due to the similarities with other fields where controlled experimentation is possible and valuable, the book can be a useful guide to A/B testing in areas other than website and mobile application development. Product managers and growth experts should benefit from it regardless of the particular product or service they are focused on.

The overall framing of the presentation is always mindful of the topic of business objectives achieved through statistical methods. This will be useful for those readers who have some experience of statistics in other disciplines, and who are now looking to understand the use of statistical tools in facilitating decision-making through online A/B testing.

However, I must emphasize that the unique research presented in this book, in terms of ideas, models, and simulation results, will be valuable to all readers, regardless of background.

Statistical Methods in Online A/B Testing

This is the end of this free pdf preview of “Statistical Methods in Online A/B Testing”. If you are looking for the eBook version of the book, it can be purchased [on Kobo](#) while the paperback is available [on Amazon](#).

If you are looking for free works of mine, then please consider the vast amount I’ve already shared for free:

- Close to one hundred articles on the [Analytics Toolkit blog](#), full of high-quality information.
- Several hundred entries of an [A/B testing glossary](#) which will get you up to speed with key statistical and experimentation concepts.
- Several [whitepapers](#), available completely for free as well.

You are welcome to read all these free materials before you decide whether it is worth it for you to spend money on the book.

“Statistical methods in online A/B testing” has been the product of many months of effort which built on top of years of experience. If you are looking to obtain its value, you should do so by purchasing it on Amazon or Kobo - the two vendors I have authorized to distribute my work. This will help motivate me make more of my work available to the public one way or another.

Thank you.

Georgi Georgiev

Author of “Statistical Methods in Online A/B Testing”