

Imperial College of Science, Technology and Medicine

Department of Mathematics

Some Statistical and Data Challenges When Building Early-Stage Digital Experimentation and Measurement Capabilities

C. H. Bryan Liu

Submitted in part fulfilment of the requirements for the degree of
Doctor of Philosophy in Modern Statistics and Statistical Machine Learning of
Imperial College London, 3rd October 2023

I confirm that the thesis reports original research work done by myself during the programme of study. The thesis has not been submitted elsewhere for examination for a PhD degree. Part of the thesis is adapted from research papers that I first-authored and published during the programme of study. They are listed in Section 1.4 and stated at the beginning of relevant chapters where applicable. In addition, I have appropriately referenced ideas from others in the thesis and sought the necessary permission to use copyrighted materials.

The thesis is written at a time when large language models (LLMs) are sufficiently advanced to produce “human-like” writing on an arbitrary topic and refine the text based on external prompts. I see much potential with such technology in research and industry practice. That said, the use of LLMs in academic writing remains highly controversial as of 2023, with many unresolved issues on correctness, academic integrity and plagiarism, and publishing ethics [2].

As such, I also confirm that I did not use any LLMs or services derived from them to write this thesis. That said, I used (the non-LLM part of) Grammarly, an online contextual proof-reading platform, to help correct the written text’s spelling, punctuation, and grammar. The platform also flags redundant, overused, or informal phrases and sentences based on set rules to help improve the overall delivery of the thesis. However, the decision to incorporate the suggestions solely rests with me.

The copyright of this thesis rests with the author. Unless otherwise indicated, its contents are licensed under a Creative Commons Attribution 4.0 International Licence (CC BY 4.0).

Under this licence, you may copy and redistribute the material in any medium or format for both commercial and non-commercial purposes. You may also create and distribute modified versions of the work. This on the condition that you credit the author.

When reusing or sharing this work, ensure you make the licence terms clear to others by naming the licence and linking to the licence text. Where a work has been adapted, you should indicate that the work has been changed and describe those changes.

Please seek permission from the copyright holder for uses of this work that are not included in this licence or permitted under UK Copyright Law.

Abstract

Digital experimentation and measurement (DEM) capabilities – the knowledge and tools necessary to run experiments with digital products, services, or experiences and measure their impact – are fast becoming part of the standard toolkit of digital/data-driven organisations in guiding business decisions. Many large technology companies report having mature DEM capabilities, and several businesses have been established purely to manage experiments for others. Given the growing evidence that data-driven organisations tend to outperform their non-data-driven counterparts, there has never been a greater need for organisations to build/acquire DEM capabilities to thrive in the current digital era.

This thesis presents several novel approaches to statistical and data challenges for organisations building DEM capabilities. We focus on the fundamentals associated with building DEM capabilities, which lead to a richer understanding of the underlying assumptions and thus enable us to develop more appropriate capabilities. We address *why* one should engage in DEM by quantifying the benefits and risks of acquiring DEM capabilities. This is done using a ranking under lower uncertainty model, enabling one to construct a business case. We also examine *what* ingredients are necessary to run digital experiments. In addition to clarifying the existing literature around statistical tests, datasets, and methods in experimental design and causal inference, we construct an additional dataset and detailed case studies on applying state-of-the-art methods. Finally, we investigate *when* a digital experiment design would outperform another, leading to an evaluation framework that compares competing designs' data efficiency.

These approaches aim to enable one to run experiments that produce less biased estimates more quickly and adapt to various business constraints. As we maintain the theoretical rigour, we also emphasise applied use cases, interpretable processes/results, and practical tradeoffs, all to ensure that the contributions are accessible to researchers and practitioners from diverse scientific backgrounds.

Acknowledgements

Academic research is never a solo effort, and I would like to acknowledge the many people who have helped, in various ways, for me to deliver this thesis.

First and foremost, I would like to thank my supervisor, Prof Emma McCoy, for her consistent support and guidance throughout the years. She has suggested many worthy research directions, challenged many potential ideas, and helped clarify my exposition on many topics. I would also like to thank Dr Seth Flaxman and Prof Robin Evans for the initial project idea discussions, as well as Prof Niall Adams, Dr Ciara Pike-Burke, Prof Nick Heard, and the anonymous reviewers for their feedback during the initial stages of the PhD.

I am fortunate to have met many collaborators whom I enjoy the many intellectually stimulating discussions and opportunities to publicise the latest results in high-impact venues with: Dr Elaine Bettaney, Dr Ângelo Cardoso, Dr Benjamin Chamberlain, Paul Couturier, Prof Marc Deisenroth, Georg Grob, Dr Duncan Little, and Dr Roberto Pagliari. I have also met many bright minds at Imperial College London and the University of Oxford, many affiliated with the EPSRC CDT in Modern Statistics and Statistical Machine Learning (StatML CDT), whose conversations provided me with new perspectives on the challenges at hand and enabled me to discover new opportunities.

The thesis is one of the many outputs from a fruitful academic-industry collaboration between the StatML CDT at Imperial and Oxford and the AI and Data Science Platform at ASOS.com, with myself also employed at the latter as a data/machine learning scientist. I am blessed to have also worked with and learnt from many talented colleagues at ASOS, including my managers Dr Benjamin Chamberlain, Papinder Dosanjh, and Dawn Rollocks; my teammates in the Customer Understanding (2016–18), Experimentation (2018–19), Prophecy (2019–20), Cassandra (2020), Kernel (2021–22), Experimentation (2022), and the Alchemists (2022–23) teams; and many more unnamed. Among other things, they have ensured that my research remains relevant to industry applications. The research is also generously funded by the two organisations mentioned above (EPSRC grant no. EP/S023151/1), without which I probably would not have taken on the PhD programme at all.

Finally, I would like to thank my friends and family, especially my mum Yvonne, my dad Has-ton, my sister Nicole, and my grandmother Cecilia, for their love and warm encouragement while I took on the ambitious task of studying a PhD while holding an industry job. My PhD years were dominated by the COVID-19 pandemic, and travel restrictions have made visiting them virtually impossible for a few years. Fortunately, the availability of economical video calls has made the isolation much more bearable.

"The true method of knowledge is experiment."

– William Blake

"Negative results are just what I want. They're just as valuable to me as positive results. I can never find the thing that does the job best until I find the ones that don't."

– Thomas A. Edison

"An experiment is a question which science poses to Nature and a measurement is the recording of Nature's answer."

– Max Planck

"Measurement is the first step that leads to control and eventually to improvement. If you can't measure something, you can't understand it. If you can't understand it, you can't control it. If you can't control it, you can't improve it."

– H. James Harrington

Contents

Abstract	3
Acknowledgements	5
1 Introduction	17
1.1 Motivation	17
1.2 Key Research Questions	19
1.3 Contributions and Chapter Organisation	21
1.4 Publications	23
1.5 List of Acronyms and Mathematical Symbols	24
2 What is the Value of Digital Experimentation and Measurement Capabilities?	33
2.1 Motivation	34
2.2 Related work	37
2.3 Mathematical Formulation	38
2.3.1 Modelling values with statistical distributions	39
2.3.2 Key results	41
2.4 Calculating the Expectation	41
2.4.1 Expectation under normal assumptions	43
2.5 Calculating the Variance	45
2.5.1 $\text{Var}(\mathcal{H}_{(r)})/\text{Var}(\mathcal{L}_{(s)})$	48
2.5.2 $\text{Var}(\mathcal{V}_{\mathcal{I}(r)})/\text{Var}(\mathcal{V}_{\mathcal{J}(s)})$	48

2.5.3	$\text{Var}(\mathcal{W}_1)/\text{Var}(\mathcal{W}_2)$	49
2.5.4	$\text{Var}(\mathcal{D})$	50
2.6	Empirical Verification	53
2.7	Case Study	57
2.7.1	e-Commerce companies	57
2.7.2	Marketing companies	58
2.8	Empirical Extensions	60
2.8.1	Valuation under independent t -distributed assumptions	60
2.8.2	Partial estimation/measurement noise reduction	62
2.9	A Brief Recap	64
3	Statistical Testing	67
3.1	Motivation	67
3.2	Pre-matter and Considerations	70
3.3	Alternative Approaches to Formal Tests	72
3.3.1	The inter-ocular trauma test	73
3.3.2	Effect size	73
3.3.3	Confidence Intervals	74
3.4	Null Hypothesis Significance Tests	77
3.5	Examples of NHSTs	79
3.5.1	z -test	79
3.5.2	t -tests	81
3.5.3	The binomial test of proportions	87
3.6	Key Quantities in NHST	89
3.6.1	p -value	89
3.6.2	Type I and II errors	92

3.6.3	Significance level	93
3.6.4	Power	95
3.6.5	Required sample size	98
3.6.6	Minimum detectable effect	100
3.7	Non-parametric Tests	101
3.7.1	Mann-Whitney U test	102
3.7.2	χ^2 goodness-of-fit test	103
3.8	Sequential Tests	105
3.9	Bayesian Tests	109
3.10	A Brief Recap	113
4	Datasets for Digital Experiments	115
4.1	Motivation	115
4.2	A Taxonomy for Digital Experiment Datasets	118
4.3	Public Online Controlled Experiment Datasets	120
4.4	Matching Dataset Taxonomy with Statistical Tests	123
4.4.1	Effect size and Welch's t -test	123
4.4.2	Experiments with adaptive stopping	124
4.4.3	Non-parametric tests	126
4.5	A Novel Dataset for Experiments with Adaptive Stopping	127
4.5.1	Potential use cases	128
4.5.2	Ethical considerations	131
4.6	A Brief Recap	132
5	Recent Advances in Digital Experimentation Methods	135
5.1	Motivation	135
5.2	From Correlation to Causation – The Potential Outcomes Framework	136

5.3	Classes and Examples of Digital Experiments	139
5.3.1	Randomised controlled trials	139
5.3.2	Quasi-experiments	142
5.3.3	Natural experiments	143
5.3.4	Special case – Experiments for personalisation strategies	144
5.4	Advances in Online Controlled Experiments	145
5.4.1	Advanced statistical methods	146
5.4.2	Other topics	151
5.5	A Spotlight on Dependent Responses in e-Commerce Experiments	154
5.5.1	Beware of Hidden Inflation in Measurement Uncertainty	156
5.5.2	Re-estimating the SE using bootstrap	157
5.5.3	Impact on OCE decisions	160
5.5.4	Lessons Learned During Deployment	161
5.6	Advances in Quasi-experiments and Natural Experiments	164
5.6.1	CATE and Conditional Exchangeability Revisited	165
5.6.2	Methods and Examples	167
5.6.3	A mini spotlight on geo-experiments	169
5.7	A Brief Recap	172
6	An Evaluation Framework for Personalisation Strategy Experiment Designs	173
6.1	Motivation	173
6.2	Evaluation Framework	175
6.2.1	User grouping	176
6.2.2	Experiment setups	178
6.2.3	Evaluation criteria	179
6.3	Comparing Experiment Setups	180

6.3.1	Actual and minimum detectable effect sizes	181
6.3.2	Is dilution always bad?	184
6.3.3	When is dual control more effective?	186
6.4	Empirical Verification	189
6.5	A Brief Recap	191
7	Conclusion	193
7.1	Summary of Thesis Achievements	193
7.2	Future Work	195
	References	199
A	Miscellaneous Mathematical Results	223
A.1	Probability in Generating Specific Order Statistic Pairs	223
A.1.1	Representing the probability as a certain number of successes in multiple Bernoulli trials	227
A.1.2	Modelling the success count distribution using beta-binomials	229
A.1.3	Estimation under normal assumptions	230
A.1.4	Some basic validations	232
A.2	Calculations Related to Dilution in Personalisation Strategy Experiment Designs	236
A.2.1	The first criterion	237
A.2.2	The second criterion	238
A.3	Calculations Related to Dual Control in Personalisation Strategy Experiment Designs	244
A.3.1	The master inequality	245
A.3.2	Simplifying RHS by assuming similar response variances	248
A.3.3	Simplifying LHS by assuming similar sample sizes	249
A.3.4	Required number of users for dual control to emerge superior	251

List of Tables

1.1	List of mathematical symbols and the quantity/concept represented by them.	26
2.1	The number of statistical tests with the centred 95% confidence interval containing the derived theoretical value for each quantity of interest.	55
4.1	Results from the first ever survey of OCE datasets.	122
4.2	Comparing the number of reject / not reject verdicts given by Welch's t -test and an mSPRT at the end of an experiment for all four metrics.	131
5.1	Pitfalls in running an OCE end-to-end as described in nine selected works. . .	153
5.2	Summary of the three online retail/e-commerce transaction datasets featured in this paper (post data cleaning).	156
5.3	The ratio between the bootstrap standard error (SE) estimates and the vanilla sample SE estimates for ASP.	163
6.1	All group-scenario combinations in our evaluation framework for personalisation strategy experiments.	177
6.2	Number of evaluations where the theoretical value of the quantities falls between the 95% bootstrap confidence interval for each experiment setup.	190
A.1	Mean Kullback-Leibler (KL) divergence between the empirical marginal distribution of $\mathbb{P}(\mathcal{I}(r) = \mathcal{J}(s))$ and distributions produced from different modelling approaches.	235

List of Figures

1.1	Illustration of an A/B test.	18
2.1	Prioritising projects according to their value under different noise levels in value estimation.	35
2.2	The generative model in the ranking under lower uncertainty problem in plate notation.	40
2.3	The generative model in the ranking under lower uncertainty problem when two distinct noise levels are involved in plate notation.	46
2.4	Relationship between different variances/covariances used to calculate the variance of the value gained when the estimation noise is reduced.	47
2.5	Relationship between different quantities in a three-item generative model under the ranking under lower uncertainty problem.	52
2.6	Histogram of the theoretical quantity's percentile rank compared to the empirical samples across multiple statistical tests.	56
2.7	The value gained by having some digital experimentation and measurement capabilities $\mathcal{D}$ under different capacity M in the case study on 6,700 e-commerce experiments reported by Browne and Johnson [30].	59
2.8	The value gained by having some digital experimentation and capabilities $\mathcal{D}$ under different capacity M in the case study on 184 marketing experiments reported by Johnson et al. [136].	61
2.9	The value gained by having some some digital experimentation and measurement capabilities under normal assumptions and t -distributed assumptions. .	63

2.10	The near-linear relationship between the proportion of business propositions with their value estimated under a lower estimation noise and the improvement in mean true value of the selected propositions under normal assumptions.	65
4.1	Distribution of p -values attained by the 99 OCEs in the ASOS Digital Experiment Dataset using Welch's t -tests, split by decision metrics.	129
4.2	Change in p -values in a mixed sequential probability ratio test and posterior belief in the null hypothesis in a Bayesian hypothesis test during the experiment for five experiments selected at random.	130
5.1	The onion model for different types of experiments.	140
5.2	Kernel density estimation plots showing a user's transaction value/size against that of their next transaction in three online retail transaction datasets.	157
5.3	The ratio between the one-way bootstrap standard error (SE) estimates and the vanilla sample SE estimates for different business metrics.	158
5.4	Trajectories of the one-way bootstrap standard error (SE) estimates and the vanilla SE estimates under different dataset-metric combinations.	159
5.5	The actual test power of and the coverage of the centred confidence interval from a two-tailed z -test against different standard error ratios.	160
6.1	Venn diagram of the user groups in our evaluation framework for personalisation strategy experiment designs.	176
6.2	Experiment setups overlaid on the user grouping Venn diagram in Figure 6.1.	177
A.1	Figure 2.3 and RHS of Figure 2.5 restated.	224
A.2	Heatmap showing the empirical distribution of $\mathbb{P}(\mathcal{I}(r) = \mathcal{J}(s))$ across different r and s .	225
A.3	The marginal distribution of $\mathbb{P}(\mathcal{I}(r) = \mathcal{J}(s))$ across different s for a given r .	226
A.4	The fitted marginal distributions overlaid on Figure A.3.	234
A.5	The fitted marginal distributions overlaid on the empirical marginal distributions of $\mathbb{P}(\mathcal{I}(r) = \mathcal{J}(s))$ ($N = 10$).	235

Chapter 1

Introduction

1.1 Motivation

The value of making data-driven or data-informed decisions has become increasingly clear in recent years [187, 256]. Key to making data-driven decisions is the ability to accurately measure a given choice's impact and experiment with possible alternatives. We define digital experimentation and measurement (DEM) capabilities as the knowledge and tools necessary to run experiments (controlled or otherwise) with different digital products, services, or experiences and measure their impact. The capabilities may be an online controlled experiment framework, a team of analysts, or a system capable of performing machine learning-aided causal inference.

The simplest example of a digital experiment is what is commonly known as an A/B test [155]. Suppose we are interested in whether offering free delivery to users of an e-commerce website will lead to more of them making a purchase (i.e., an increased *conversion rate* in business speak). We set up an experiment where incoming users are randomly split into two groups, where one group is shown a "free delivery" banner on the website (the *treatment*), while the other acts as the *control* – being shown the original website without any mention of free delivery (see Figure 1.1). We calculate a *decision metric* (here, the conversion rate) for both groups based on responses from their members and compare the decision metrics using a statistical test to draw causal statements about the treatment. The approach is popular in digital organisations, with the largest technology companies having reported running hundreds or




Figure 1.1: Illustration of an A/B test (a randomised controlled trial with two parallel groups) designed to test the claim that “offering free delivery to users of an e-commerce website will lead to a larger proportion of the said users making a purchase.” Incoming users are split randomly into two groups, where one group (A) acts as the control, and the other group (B) is shown a “free delivery” banner on the website as the treatment.

thousands of experiments at any given time [149, 260, 288].

From a statistical point of view, digital experimentation and measurement is essentially the application of experimental design and causal inference methods in a digital setting. Likewise, the free delivery experiment mentioned above is an online randomised controlled trial. Readers will agree that the underlying experimental design and causal inference methods have been established and frequently applied in agriculture, medicine, and economics for decades.

That said, the arrival of the Web and, subsequently, the ability to run an experiment end-to-end online have brought different opportunities and challenges. The differences are well documented in [156], which necessitate the development of a different set of tools and processes exclusively for digital experiments. In some cases, they also require a complete rethink of how we approach experiment design and analysis. Under such premises, we examine the statistical and data challenges one faces when building DEM capabilities.

Guiding Principles Before we state the research questions in detail (see section below), we outline the three guiding principles that lead us to the questions. Firstly, similar to many contemporaries in digital experimentation, we seek methods that enable one to make better decisions more quickly and thus drive business/organisational growth. The methods may be those that produce less biased estimates, generate measurements more quickly, or give alternatives that can measure the impact of an intervention under various constraints in practice (e.g., inability to perform user-level random assignment in the “free delivery” example above).

Secondly, we seek interpretable processes and results alongside performance to maintain engagement from those beyond our field. We, humans, value the understanding of why a decision is made as such, and the use of black-box models or overly complex procedures in decision-making often faces additional barriers in gaining trust from stakeholders despite its potential to improve performance [237].

Lastly, and perhaps more importantly, we seek to make the knowledge and tools involved in digital experimentation accessible, thus ensuring widened participation.¹ Many recent advances in digital experimentation are proposed by researchers and practitioners who operate in an environment with mature capabilities. Focusing on those works may give the impression that mature capabilities are the norm and that we only need a little work to explain the basics on top of what is already available. Both are not true – the DEM capabilities of many individuals and organisations are still in their infancy. These individuals and organisations will benefit from us continuing to consolidate and clarify the building blocks in digital experiments, e.g., statistical tests, datasets, and decision metrics.

1.2 Key Research Questions

The thesis aims to address digital experimenters’ and organisations’ challenges as they build DEM capabilities. As the thesis title suggests, most of the work concerns statistical and data challenges one encounters at the early stage of the process. It also touches upon topics en-

¹ Clearly, a PhD thesis is a poor medium to engage with a non-technical audience. We thus limit the scope of “widened participation” to include researchers and practitioners with diverse academic backgrounds.

countered by those with established capabilities, enabling readers to take away something wherever they are on the experimentation maturity spectrum [90]. We ask the following questions:

Why should one engage in digital experimentation and measurement? While running experiment(s) to evaluate a hypothesis is second nature to scientists, it is unnatural in the business world. Business decision-makers prefer having business cases, ideally complete with a cost-benefit analysis. The ability to quantify the benefits of acquiring DEM capabilities (the knowledge and tools necessary to experiment and measure) will strengthen our business case. It enables us to speak in the language of the business and thus stand a better chance of garnering their support.

What ingredients do we require to run experiments successfully in a digital setting? Committing to investing in DEM capabilities is the first of many steps to reap the benefits. To run their first experiment successfully, organisations require, at a minimum, a clear hypothesis and evaluation criteria, detailed procedures to deliver different treatments to and collect responses from experiment participants, plus basic knowledge in data processing and statistical testing. They would also require knowledge of advanced mathematical and statistical methods, powerful computer systems with sufficient computing power and data storage capacity, and a growth mindset organisational culture that embraces scientific, data-driven discoveries to scale their experiment operations. Each topic, whilst important, warrants extensive consideration. The thesis will focus on statistical testing, data, and advanced methods. We will provide a brief overview and pointers to relevant work for other topics for those interested.

When would an experiment design outperform another? In the parable of two woodcutters, the more experienced woodcutter emerged victorious in a woodchopping contest by spending most of their time sharpening their axe. In the same spirit, it is vital to continuously improve our tools, one of which is experiment designs, as we engage in experimentation efforts. While there are many qualities that we can use to compare two experiment designs, in digital experimentation, we usually care about *data efficiency*. This is represented by the

minimum number of experiment participants required to reach a statistically sound conclusion.² Increasing the data efficiency (i.e., requiring fewer experiment participants) enables experimenters to make decisions sooner, leading to better organisational performance as one can deploy improvements and roll back harmful changes more quickly.

1.3 Contributions and Chapter Organisation

The thesis is divided into seven chapters, each detailing contributions to a different statistical and data topic necessitated by developing early-stage DEM capabilities. The broad coverage of topics means the thesis embeds the background knowledge required for each topic within the individual chapters.

The remainder of *Chapter 1* covers housekeeping items: a list of publications from the research programme and a list of mathematical symbols and the quantities/concepts they represent.

Chapter 2 introduces the ranking under lower uncertainty problem as a novel model to value DEM capabilities. This arises from the observation that DEM capabilities reduce measurement uncertainty on the value of individual business propositions, improving the inherent value of business propositions chosen during prioritisation. We derive the expected gain and the variance of our gain estimate and provide case studies based on large-scale meta-analyses. This addresses the first research question and enables one to build a quantitative business case to justify investment in the capabilities.

Chapter 3 provides an introduction to statistical testing in digital experimentation that is both mathematically rigorous and accessible to researchers and practitioners with different backgrounds. It starts from the very beginning of a statistical test (i.e., specifying the statistical hypotheses), provides a detailed treatment on the popular null hypothesis significance testing (NHST) framework, and involves other more advanced test paradigms (e.g., sequential and Bayesian testing) and alternatives. It aims to provide a sufficient theoretical grounding to appreciate what is happening behind the scenes while highlighting the pitfalls in test design and interpretation that can trap even the most experienced.

² This is closely related to optimal design [8] and the asymptotic efficiency of a statistical test [254] in the statistical theory literature. We will approach the challenge from a more applied point of view.

Chapter 4 describes a systemic investigation of datasets in digital experiments. We create the first ever taxonomy for digital experiment datasets and compile the first ever survey for publicly available online controlled experiment datasets. We also map the relationship between the dataset taxonomy and statistical tests discussed in Chapter 3. The taxonomy and survey also identify gaps in dataset availability, leading to the development of the ASOS Digital Experiments Dataset. This first real, multi-experiment time series dataset enables the design and running of experiments with adaptive stopping.

Chapter 5 reviews existing challenges in digital experimentation and the established and emerging methods that address such challenges. We also complement the review by viewing some of the reported methodological advances through a practical lens. This is done by interleaving the review with case studies on randomised controlled trials with dependent responses and quasi-experiments with geographical regions as treatment units. They provide evidence on the extent of reported challenges and practical considerations when implementing relevant methods. Together with Chapters 3 and 4, these chapters address the second research question.

Chapter 6 develops an evaluation framework for personalisation strategy experiment designs. These experiment designs face unique challenges, such as low test power and partial non-compliance from users, issues that we will cover in Chapters 3 and 5. The evaluation framework enables one to compare two experiment setups given the circumstances they face with their target audience, thus addressing the third research question. We also derive interpretable rules of thumb from the framework to enable experimenters to compare typical experiment setups quickly.

We conclude in *Chapter 7* with an outline of promising investigation pathways arising from each workstream. In addition, we make all the code used for experiments and simulations described in the thesis, as well as the ASOS Digital Experiment Dataset, publicly available. Readers can find the links to the code repositories and the dataset at the bottom of this page³

³ The link to code repositories and the dataset are as follows:

- Code repository associated with the ranking under lower uncertainty framework (as described in Chapter 2): https://github.com/liuchbryan/ranking_under_lower_uncertainty;
- ASOS Digital Experiments Dataset and accompanying datasheet (as described in Chapter 4): <https://osf.io/64jsb/>. The OSF project also embeds a code repository with experiments on the dataset (<https://github.com/liuchbryan/oce-dataset/>);

and again in the footnotes of the relevant chapters.

1.4 Publications

This thesis contains research that has been published during PhD study. We list where readers can find them in the thesis, together with details of the associated conference and journal papers, below:

- Chapter 2 and Appendix A.1: C. H. B. Liu and B. P. Chamberlain (2019). What is the Value of Experimentation & Measurement? In: *2019 IEEE International Conference on Data Mining (ICDM '19)*, 1222–1227. DOI: 10.1109/ICDM.2019.00151.
- Chapter 2 and Appendix A.1: C. H. B. Liu, B. P. Chamberlain and E. J. McCoy (2020). What is the Value of Experimentation and Measurement? *Data Science and Engineering* 5, 152–167. DOI: 10.1007/s41019-020-00121-5. Part of Special Issue: Highly-Rated Short Papers of ICDM 2019.
- Chapter 4: C. H. B. Liu, Â. Cardoso, P. Couturier and E. J. McCoy (2021). Datasets for Online Controlled Experiments. In: *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks (NeurIPS Datasets and Benchmarks '21)*. URL: <https://datasets-benchmarks-proceedings.neurips.cc/paper/2021/file/274ad4786c3abca69fa097b85867d9a4-Paper-round2.pdf>.
- Chapter 5 (Section 5.5): C. H. B. Liu and E. J. McCoy (2023). Measuring e-Commerce Metric Changes in Online Experiments. In: *Companion Proceedings of the ACM Web Conference 2023 (WWW '23 Companion)*. DOI: 10.1145/3543873.3584654.
- Chapter 6, Appendix A.2, and Appendix A.3: C. H. B. Liu and E. J. McCoy (2020). An Evaluation Framework for Personalization Strategy Experiment Designs. *arXiv*: 2007.11638

-
- Code repository and results on the two publicly available datasets associated with e-commerce experiments with dependent responses (as described in Section 5.5): <https://github.com/liuchbryan/oce-ecomm-abv-calculation>; and
 - Code repository associated with the evaluation framework for personalisation strategy experiment designs (as described in Chapter 6): https://github.com/liuchbryan/experiment_design_evaluation.

[stat.ME]. Presented in *AdKDD '20 Workshop* (in conjunction with KDD '20), awarded Best Student Paper.

1.5 List of Acronyms and Mathematical Symbols

Readers may encounter the following acronyms recurring in multiple thesis sections. We will define the acronyms in the main text at least once before referring to the acronyms.

- ATE: Average treatment effect
- CATE: Conditional average treatment effect
- CI: Confidence interval
- CDF: Cumulative density function
- CLT: Central limit Theorem
- CVR: Conversion rate
- DEM: Digital experimentation and measurement (usually followed by the word “capability/capabilities”)
- MDE: Minimum detectable effect
- NHST: Null hypothesis significance test(ing)
- OCE: Online controlled experiment
- PDF: Probability density function
- RCT: Randomised controlled trial

We also list the mathematical symbols and the quantities/concepts they represent below. Each entry is indexed by the leading symbol and follows the following format:

(Applicable chapter(s)) *Symbol(s)* – (Type) Quantity/concept the symbol represents.

Applicable chapter(s) This states the thesis chapter(s) that a notation applies to. The notations used in Chapter 2 also apply to Appendix A.1, and those used in Chapter 6 also apply to Appendices A.2 and A.3.

Symbol(s) The thesis covers topics traditionally belonging to multiple statistical disciplines. Thus, it is inevitable to see some notation clash. We follow the order set out below when deciding which set of notations takes precedence:

1. Elementary statistical constructs found in introductory statistical texts (e.g., X for a random variable, $f_X(\cdot)$ for the probability density function of X , and $\mathbb{E}(X) = \mu_X$ for the expected value of X);
2. Notation specific to statistical testing (e.g., H_0 and H_1 for null and alternate hypotheses, α for the significance level, β for the power);
3. Common notation in other sub-fields, including causal inference, order statistics, and econometrics; and
4. Specific notation introduced by individual research articles.

Notations with lower precedence are generally assigned a new symbol to minimise confusion. That said, in a couple of cases, the same symbol may represent two different quantities/concepts in the same chapter – their meaning should be apparent given the context surrounding its use.

Type We also abbreviate the types of quantities and concepts in the tables as follows:

- calc.: calculated quantity,
- const.: constant,
- dist.: distribution,
- func.: function,
- id.: identifier, and
- r.v.: random variable.

Table 1.1: List of mathematical symbols and the quantity/concept represented by them (seven pages). See Section 1.5 for further information.

Leading Symbol	Quantity/Concept Represented
A/a	(Chapter 5) $\mathcal{A}$ – (r.v.) Treatment Assignment variable in the potential outcomes framework, with realised values denoted a. (Chapter 6) A (and $A1, A2$) – (id.) An Analysis group in a personalisation strategy experiment, often referenced together with B.
B/b	(Chapter 3) B – (calc. / r.v.) Binomial test of proportion statistic. (Chapters 3–4) $BF_{n,m}$ – (calc.) Bayes Factor (in a Bayesian test, upon observing first n and m items in the first and second sample, respectively). (Chapter 6) B (and $B1, B2$) – (id.) An analysis group in a personalisation strategy experiment, often referenced together with A.
C/c	(All chapters) $\text{Cov}(\cdot, \cdot)$ – (func.) Covariance. (Chapter 2) c – (const.) Adjustment to the quantile of an order statistic. (Chapter 2) c – (const.) Risk-free value in the Sharpe ratio. (Chapter 3) C^2 – (calc. / r.v.) χ^2 goodness-of-fit test statistic. (Chapter 5) $\mathcal{C}$ – (r.v.) Covariate(s) under the potential outcomes framework. (Chapter 6) $C0, C1, C2, C3$ – (id.) User groups in a personalisation strategy experiment that acts as Control. (Appendix A.1) C – (r.v.) Count.
D/d	(Chapter 2) $\mathcal{D}$ – (r.v.) Difference in value of items selected in the ranking under lower uncertainty problem. (Chapters 3–4) d – (calc.) Cohen’s d.
E/e	(All chapters) $\mathbb{E}(\cdot)$ – (func.) Expected value. (Chapter 2) $\mathcal{E}_i / \mathcal{E}_{(r)}$ – (r.v.) Estimated value of item i / rth-ranked item under estimation uncertainty in the ranking under lower uncertainty problem, often referenced together with $\mathcal{V}_i / \mathcal{V}_{I(r)}$. (Chapters 3) E_i – (r.v.) Expected frequency for category i in a χ^2 goodness-of-fit test. (Chapters 3–4) $E_{n,m}$ – (calc.) Effective sample size (in a Bayesian test, upon observing first n and m items in the first and second sample, respectively).

Leading Symbol	Quantity/Concept Represented
F/f	(All chapters) $F_X(\cdot)$ – (func.) Cumulative density function (CDF) of an r.v. X. (All chapters) $f_X(\cdot)$ – (func.) Probability density function (PDF) of an r.v. X. (All chapters) $f'_X(\cdot)$ – (func.) Derivative of the PDF of an r.v. X. (All chapters) $f(\cdot \cdot)$ – (func.) Conditional PDF.
G/g	(Chapter 5) $g(\cdot)$ – (func.) A generic function that maps the treatment and covariate(s) to the observed/potential response under the potential outcomes framework. (Chapter 6) G – (id.) A generic user Group-scenario combination in a personalisation strategy experiment. See Table 6.1 for possible values.
H/h	(Chapter 2) $\mathcal{H}_i / \mathcal{H}_{(r)}$ – (r.v.) Estimated value of item i / rth-ranked item under high estimation noise in the ranking under lower uncertainty problem, often referenced together with $\mathcal{V}_i / \mathcal{V}_{\mathcal{I}(r)}$ and $\mathcal{L}_j / \mathcal{L}_{(s)}$. (Chapters 3–6) H_0, H_1 – (id.) Null and alternate hypotheses. (Chapters 3–4) H – (id.) The mixing distribution in a mixed Sequential Probability Ratio Test (mSPRT). (Chapters 3–4) H – (id.) A generic statistical hypothesis in a Bayesian test.
I/i	(All chapters) i – (id.) Index identifying an r.v. within a set of i.i.d. r.v.'s, often referenced together with j. (Chapter 2) $\mathcal{I}(\cdot)$ – (func.) Index function that maps the rank of an r.v. within a set of independently identically distributed (i.i.d.) r.v.'s $\mathcal{E} / \mathcal{H}$ to its index, often referenced together with $\mathcal{J}(\cdot)$. (Chapter 6) $I1, I2, I\phi, I\psi$ – (id.) User groups in a personalisation strategy experiment under Intervention. (Chapter A.1) $\mathbb{I}$ – (r.v.) Indicator variable, often with the condition expressed as a subscript.

Leading Symbol	Quantity/Concept Represented
J/j	(All chapters) j – (id.) Index identifying an r.v. within a set of i.i.d. r.v.'s, often referenced together with i. (Chapter 2) $\mathcal{J}(\cdot)$ – (func.) Index function that maps the rank of an r.v. within a set of independently identically distributed (i.i.d.) r.v.'s. $\mathcal{L}$ to its index, often referenced together with $\mathcal{I}(\cdot)$.
L/l	(Chapter 2) $\mathcal{L}_i / \mathcal{L}_{(s)}$ – (r.v.) Estimated value of item i / sth-ranked item under low estimation noise in the ranking under lower uncertainty problem, often referenced together with $\mathcal{V}_i / \mathcal{V}_{\mathcal{J}(s)}$ and $\mathcal{H}_i / \mathcal{H}_{(r)}$. (Appendix A.1) $\mathcal{L}^*$ – (r.v.) $\mathcal{L}_{\mathcal{I}(r)}$ normalised by the mean and variance of $\mathcal{L}_j$ under normal assumptions.
M/m	(Chapter 2) M – (const.) Number of items one can select (i.e., capacity) in the ranking under lower uncertainty problem. (Chapters 3–6) m – (const.) Number of items in (or size of) a statistical sample Y, either total or that observed so far within a sequential/Bayesian test setting; often referenced together with n.
N/n	(All chapters) $\mathcal{N}(\cdot, \cdot)$ – (dist.) A normal distribution, with its mean and variance as parameters. (Chapter 2) N – (const.) Number of items available for selection in the ranking under lower uncertainty problem. (Chapters 3–6) n – (const.) Number of items in (or size of) a statistical sample X, either total or that observed so far within a sequential/Bayesian test setting; often referenced together with m. (Chapter 6) $n_{\mathcal{G}}$ – (const.) Number of samples in a user/analysis group $\mathcal{G}$, where $\mathcal{G} \in \{0, 1, 2, 3, A, B, A1, A2, B1, B2\}$, in a personalisation strategy experiment.
O/o	(Chapters 3) O_i – (r.v.) Observed frequency for category i in a χ^2 goodness-of-fit test. (Chapter 6) $O(\cdot)$ – (func.) Big-O notation describing a function's limiting behaviour.
P/p	(Chapter 2) $\mathbb{P}(\cdot)$ – (func.) Probability. (Chapters 3–6) p – (calc.) p-value of a null hypothesis significance test.

Leading Symbol	Quantity/Concept Represented
R/r	(Chapter 2) r – (id.) Rank of an r.v. (in order statistics terms) within a set of r.v.'s $\mathcal{E} / \mathcal{H}$, often referenced together with s. (Chapter 6) R – (id.) A personalisation strategy experiment setup, often referenced together with S.
S/s	(Chapter 2) s – (id.) Rank of an r.v. (in order statistics terms) within a set of r.v.'s $\mathcal{L}$, often referenced together with r. (Chapters 3) S_n – (r.v.) Sequential probability ratio test (SPRT) statistic. (Chapters 3–4) $S(\cdot, \cdot)$ – (func.) Sign function (in a Mann-Whitney U test). (Chapters 3–6) s_X^2 – (calc.) Sample variance of a sample X. (Chapter 6) S – (id.) A personalisation strategy experiment Setup, often referenced together with R.
T/t	(All chapters) t_ν – (dist.) A Student's t distribution with ν degrees of freedom. (All chapters) $T_\nu(\cdot)$ – (func.) CDF of a Student's t-distribution with ν d.f. (All chapters) $t_{\nu,q}$ – (calc.) qth quantile of a Student's t-distribution with ν d.f. (Chapter 2) T_i / T'_i – (r.v.) Student's t-distributed r.v.'s. (Chapters 3–6) T – (calc. / r.v.) Test statistic (generic, or that of a t-test). (Chapter 4) $t, t + 1, \dots$ – (id.) Points in time (in statistical tests with adaptive stopping). (Appendix A.1) $T(\cdot, \cdot)$ – (func.) Owen's T-function.
U/u	(Chapters 2 & 6) $U(a, b)$ – (dist.) A uniform distribution with bounds a and b. (Chapters 3–4) U – (calc. / r.v.) Mann-Whitney U-test statistic.
V/v	(All chapters) $\text{Var}(\cdot)$ – (func.) Variance. (Chapter 2) $\mathcal{V}_i, \mathcal{V}_{\mathcal{I}(r)}, \mathcal{V}_{\mathcal{J}(s)}$ – (r.v.) True Value of items $i, \mathcal{I}(r)$, and $\mathcal{J}(s)$ in the ranking under lower uncertainty problem, often referenced together with $\mathcal{E}_i / \mathcal{E}_{(r)}, \mathcal{H}_i / \mathcal{H}_{(r)}$, and $\mathcal{L}_j / \mathcal{L}_{(s)}$. (Chapters 3–4) V^2 – (const.) Hyperparameter, in a Bayesian test, which represents the variance of the effect size prior.

Leading Symbol	Quantity/Concept Represented
W/w	(Chapter 2) $\mathcal{W}, \mathcal{W}_1, \mathcal{W}_2$ – (r.v.) Average true value of items selected under some estimation uncertainty, high estimation noise, and low estimation noise, respectively, in the ranking under lower uncertainty problem. (Chapters 3–4) $W_{n,m}$ – (calc. / r.v.) Wald test statistic (in a sequential/Bayesian test, upon observing first n and m items in the first and second sample, respectively).
X/x	(Chapters 3–6) X_i – (r.v.) Item i in a sample X, often referenced together with Y_j in two-sample testing. (Chapters 3–6) $\bar{X}$ – (r.v.) Sample mean of a sample X. (Chapters 3–6) $\bar{X}_n$ – (r.v.) Sample mean of a sample X (in a sequential/Bayesian test, upon observing the first n items in the sample). (Chapters 3–6) $\bar{x}$ – (calc.) Realised sample mean of a sample X. (Chapters 5) $X(a) / X_i(a)$ – (r.v.) Potential response/outcome of X/X_i under the treatment assignment $a \in \{0, 1, \dots\}$ in the potential outcomes framework, often referenced together with $Y(a) / Y_j(a)$.
Y/y	(Chapters 3–6) Y_j – (r.v.) Item j in a sample Y, often referenced together with X_i in two-sample testing. (Chapters 3–6) $\bar{Y}$ – (r.v.) Sample mean of a sample Y. (Chapters 3–6) $\bar{Y}_n$ – (r.v.) Sample mean of a sample Y (in a sequential/Bayesian test, upon observing the first n items in the sample). (Chapters 5) $Y(a) / Y_j(a)$ – (r.v.) Potential response/outcome of Y/Y_j under the treatment assignment $a \in \{0, 1, \dots\}$ in the potential outcomes framework, often referenced together with $X(a) / X_i(a)$.
Z/z	(Chapters 3–6) z_q – (calc.) qth quantile of a standard normal r.v. (Chapter 3) Z – (calc. / r.v.) Standard score / z-test statistic. (Chapter 6) z – (id.) Mathematical shorthand – the difference between two normal quantiles (see Equation (6.14)). (Appendix A.1) Z – (r.v.) A standard normal r.v.

Leading Symbol	Quantity/Concept Represented
α	(Chapters 3–6) α – (const.) Significance level of a statistical test. (Appendix A.1) $\alpha_{\mathcal{P}}$ – (const.) Beta distribution parameter, often referenced together with $\beta_{\mathcal{P}}$.
β	(Chapters 3–6) β_{θ} – (const.) Power of a statistical test under a specific hypothesis with parameter θ . (Appendix A.1) $\beta_{\mathcal{P}}$ – (const.) Beta distribution parameter, often referenced together with $\alpha_{\mathcal{P}}$.
Δ/δ	(All chapters) Δ – (calc. / r.v.) Effect size. (Chapters 3–4) Δ° – (calc. / r.v.) Effect size standardised by the pooled variance in a Bayesian test. (Chapters 3–4) $\delta_{n,m}$ – (calc. / r.v.) “Test statistic” of a Bayesian test, upon observing the first n and m items in the first and second sample, respectively. (Chapter 6) Δ_S – (calc.) Actual effect size of a personalisation strategy setup S .
ϵ	(Chapter 2) $\epsilon, \epsilon_1, \epsilon_2$ – (r.v.) Estimation noise.
η	(Chapter 6) η – (id.) Mathematical shorthand (see Equation (6.14)).
Θ/θ	(All chapters) θ – (const.) Parameter of interest in a hypothesis test. (All chapters) θ_0 – (const.) Value of the parameter of interest under the null hypothesis. (All chapters) θ^* – (const.) A specific parameter of interest. (Chapter 6) θ_S^* – (const.) Minimum detectable effect size of a personalisation strategy experiment setup S .
Λ/λ	(Chapters 3–4) Λ_n^{H,θ_0} / $\tilde{\Lambda}_n^{H,\theta_0}$ – (calc. / r.v.) One-sample / two-sample mixed sequential probability ratio test (mSPRT) statistic with mixing distribution H and null hypothesis parameter θ_0 .
μ	(All chapters) μ_X – (const.) Mean of an r.v., a population, or a group X .
ν	(All chapters) ν – (const.) Degrees of freedom (d.f.).
Ξ/ξ	(Chapter 6) ξ – (id.) Mathematical shorthand (see Equation (6.14)).

Leading Symbol	Quantity/Concept Represented
Π/π	(Chapter 3) $\pi_X - (\text{const.})$ Success rate of a Bernoulli distributed r.v. X . (Chapter 6) $\pi_{\min} - (\text{const.})$ Minimum test power.
ρ	(All chapters) $\rho_{XY} - (\text{func.})$ Correlation between two r.v.'s X and Y .
Σ/σ	(All chapters) $\sigma_X - (\text{const.})$ Standard deviation of an r.v., a population, or a group X . (All chapters) $\sigma_X^2 - (\text{const.})$ Variance of an r.v., a population, or a group X .
τ	(Chapters 3–4) $\tau^2 - (\text{const.})$ Hyperparameter in a mixed sequential probability ratio test (mSPRT), representing the variance of the mixture distribution.
Φ/ϕ	(All chapters) $\Phi(\cdot) - (\text{func.})$ Cumulative density function (CDF) of a standard normal r.v. (All chapters) $\phi(\cdot) / \phi(\cdot; \mu, \sigma^2) - (\text{func.})$ Probability density function (PDF) of a standard normal r.v. / a normal r.v. with mean μ and variance σ^2 .
χ	(Chapter 3) $\chi_\nu^2 - (\text{dist.})$ A χ^2 -distribution with ν d.f.

Chapter 2

What is the Value of Digital Experimentation and Measurement Capabilities?

From the thesis title, it is reasonable for readers to expect content covering methodological advances in statistical tests, experiment design, and causal inference. While we will address such mainstream topics in later chapters of the thesis, we first present a ranking under lower uncertainty problem as an application of order statistics to motivate our topics. As the chapter title suggests, addressing the problem enables us to value digital experimentation and measurement (DEM) capabilities quantitatively, which, in turn, helps us to answer *why* we require such capabilities in the first place.

This chapter is adapted from the research paper “*What is the Value of Experimentation & Measurement?*” It was first presented at the 2019 IEEE International Conference on Data Mining (ICDM) [174] and later extended and published in *Data Science and Engineering* as part of a special issue on Highly-Rated Short Papers of ICDM 2019 [175]. In Sections 2.1 and 2.2, as well as introducing related work, we will motivate how the requirement to value DEM capabilities arises and leads to our upcoming ranking under lower uncertainty problem. Readers who are more interested in the mathematical formulation of the ranking under lower uncertainty problem can skip straight to Section 2.3, though they may find the chapter organisation at the end of Section 2.1 useful.


2.1 Motivation

The value of DEM capabilities is currently best reflected in the success of organisations that have adopted and advocated for them in the past decade. Many major technology companies report having mature infrastructure for online controlled experiments (OCEs, e.g., Google [260], LinkedIn [288], and Microsoft [149]) or are heavily investing in state-of-the-art techniques (e.g., Airbnb [161], Netflix [286], and Yandex [222]), or both. Amazon [123], Facebook [105], and Uber [20] have also reported the use of various causal inference techniques to measure the incrementality of advertising campaigns. Several start-ups (e.g. Optimizely [134] and Qubit [30]) have also recently been established purely to manage OCEs for businesses.


While mature DEM capabilities can quantify the value of a business proposition, it remains a substantial challenge to “measure the measurer” – to quantify the value of the DEM capabilities themselves. To the best of our knowledge, no work addresses the question, “Should we invest in DEM capabilities?” or how to value these capabilities when we first present our ideas. The lack of prior work makes it hard to build a compelling business case to justify investment in the related personnel and infrastructure. We address this problem by calculating both the expected value and the risk, allowing the Sharpe ratio [246] for a DEM capability to be calculated and compared to other potential investments.

The value created by DEM capabilities comes from the following three sources:

1. **Recognising value:** DEM capabilities enable one to attribute value to a digital product, business proposition or service. They also prevent damage from business propositions that have a negative value. This is important for dynamic organisations with large numbers of business propositions, as the damage caused by individual rollouts can be compartmentalised and contained in a similar fashion, similar to unit and integration testing in software development.
2. **Prioritisation:** Without DEM capabilities, one prioritises based on back-of-envelope estimates or gut feel, which has high uncertainty. DEM reduces the magnitude of the noise arising from estimation, enabling prioritisation based on estimates closer to the true values and improving long-term decision-making (see Figure 2.1).



(a) High estimation noise



(b) Low estimation noise

Figure 2.1: Prioritising four projects (the fruits) according to their value (x-axes). The semi-opaque icons represent the projects' true value, and the solid icons represent possible project value estimates under some level of uncertainty (horizontal lines) in the estimation process. (Top) Under a noisy estimation process, projects with a low true value (e.g., project apple) may appear to have a high value and be prioritised erroneously. (Bottom) DEM reduces the estimation noise, enabling a better prioritisation with value estimates closer to the truth. The figures include icons designed by Smashicons from Flaticon.com.

3. **Optimisation:** DEM capabilities allow one to evaluate large numbers of variants against each other and the best to be selected efficiently. Without such capabilities, we can still experiment with different business propositions sequentially, though it is slow and introduces noise from the changing environment.

Quantifying the value of a DEM capability is relatively straightforward once it is in place. For example, the value of Item 1 comes from rolling back negative business propositions: we can calculate it by summing the negative contributions of unsuccessful business propositions. Likewise, we can calculate the value of Item 3 by summing the difference between the maximum and the mean value for each variant over the business propositions. Indeed, these are the approaches taken by [182] when valuing LinkedIn's experimentation platform using past experiment data.

The approaches above are not feasible for organisations building DEM capabilities from scratch. One can attempt to estimate the values from Items 1 and 3 using generic value distribution across business propositions, which is given across industries in [136] and [30], though these estimates are generally no better than back-of-envelope calculations. Without an accurate valuation, organisations remain in a chicken and egg situation – no valuation, no investment; no investment, no capability; no capability, no valuation.

Given the above, we see quantifying the value of Item 2 as the less explored yet more compelling approach and the subject of the remainder of this chapter. DEM capabilities improve prioritisation by reducing uncertainty in the value estimates of each business proposition. This is a form of ranking under uncertainty, a well-studied problem in statistics and operations research. However, in all previous work, the variance is assumed to be a fixed constant or changed without measuring the corresponding change in the value of the ranked items. Here, we wish to understand the value of ranking under *lower* uncertainty through DEM capabilities.

Our contribution is as follows. We

1. Specify the first model that values the contribution of a DEM capability in terms of better prioritisation due to reduced estimation noise for business propositions (Sections 2.3 and 2.4);

2. Derive the variance of our estimator, allowing one to calculate a Sharpe ratio to guide organisations considering an investment in DEM (Section 2.5); and finally
3. Provide two case studies based on large-scale meta-analyses that reflect how one can apply our model to real-world practice (Section 2.7) and two extensions that open the door to future work in this area (Section 2.8).

2.2 Related work

There is a vast literature on using controlled or natural experiments in a digital technology context. In addition to that mentioned in Section 2.1, there are works dedicated to running trustworthy online controlled experiments [74], choosing good metrics [126] and designing experiments where samples are dependent due to external confounders [18, 19].¹ However, these works all assume the existence of DEM capabilities, and to the best of our knowledge, no literature that helps organisations justify the acquisition of DEM capabilities exists. While [182] asked a similar question as this chapter does, they aimed to value an experimentation platform, i.e., an existing DEM capability. We believe that filling this gap is necessary for widespread adoption and that increased participation will accelerate the development of the field.

This chapter is related to existing work in statistics and operations research, particularly on decision-making under uncertainty, which has been extensively studied since the 1980s. Notable work includes proposals for additional components in a decision maker's utility function [22], alternate risk measures [289], and a general framework for decision-making with incomplete information (i.e., uncertainty) [278]. These works assume the inability to change the noise associated with estimation or measurement (or both).

The sub-problem of ranking under uncertainty has also attracted considerable attention, partially due to the advent of large databases and the requirement to rank results with a certain ambiguity in relevance [255]. While [298] measured the influence of noise levels in their work, they focused on the quality of the ranks themselves but not the value associated with the ranks.

¹ We will describe the works in greater detail in Chapter 5.

The project selection problem is a related topic in optimisation, where the goal is to find the optimal set of business propositions using mixed integer linear programming, possibly under uncertainty. Work in this domain generally seeks methods that cope with existing risk/noise [186], and to the best of our knowledge, no work considers the value of reducing risk. While [245] have discussed lowering the uncertainty level during the selection process, they refer to the uncertainty of decision parameters instead of the general noise level.

2.3 Mathematical Formulation

We formulate the ranking under lower uncertainty problem, which has a wide variety of important applications in its own right [251, 298], as follows. Consider the scenario where we select M business propositions from N candidates, where $M < N$. The Estimated value of each business proposition is given by $\mathcal{E}_i = \mathcal{V}_i + \epsilon_i$, where $\mathcal{V}_i$ are the *true yet unobserved* Values estimated with error ϵ_i . The business propositions are labelled in ascending order of estimated value $\mathcal{E}_i$ to get the order statistics $\mathcal{E}_{(1)}, \mathcal{E}_{(2)}, \dots, \mathcal{E}_{(N)}$, and we select the M business proposition with the highest estimated values: $\mathcal{E}_{(N-M+1)}, \mathcal{E}_{(N-M+2)}, \dots, \mathcal{E}_{(N)}$. We are interested in the true value of the selected business propositions, given by

$$\mathcal{V}_{\mathcal{I}(N-M+1)}, \mathcal{V}_{\mathcal{I}(N-M+2)}, \dots, \mathcal{V}_{\mathcal{I}(N)}, \quad (2.1)$$

where $\mathcal{I}(\cdot)$ denotes the index function that maps the ranking to the index of the business proposition. Note one should not confuse the set in Expression (2.1) with the set $\{\mathcal{V}_{(N-M+1)}, \mathcal{V}_{(N-M+2)}, \dots, \mathcal{V}_{(N)}\}$ – the latter denotes the top M business propositions ranked by their true value and are likely to be different [298].

We define the mean true value of the M selected business propositions as

$$\mathcal{W} \triangleq \frac{1}{M} (\mathcal{V}_{\mathcal{I}(N-M+1)} + \mathcal{V}_{\mathcal{I}(N-M+2)} + \dots + \mathcal{V}_{\mathcal{I}(N)}) , \quad (2.2)$$

where the best prioritisation maximises $\mathcal{W}$.²

² Readers can internalise the $\mathcal{W}$ notation by drawing parallels between the name of the letter in French (*double vé*) and Spanish (*doble ve*) – both literally meaning “double v” – and the fact that we use $\mathcal{W}$ to denote an aggregation of two or more $\mathcal{V}$.

Part of the value of ranking under lower uncertainty, which DEM capabilities help bring, arises from the observation that $\mathcal{W}$ increases when the magnitude of the uncertainties arising from estimation (ϵ_i) decreases. We are interested in the value gained by reducing estimation uncertainty **without changing the set of business propositions** (i.e., retaining all $\mathcal{V}_i$), as the true value of the business propositions does not depend on the measurement method used:

$$\mathcal{D} \triangleq \mathcal{W} \big|_{\text{lower noise}} - \mathcal{W} \big|_{\text{higher noise}} . \quad (2.3)$$

Obtaining the value gained ($\mathcal{D}$) enables us to compare DEM capabilities with other potential investments. The Sharpe ratio is a standard quantity used in finance to compare different investments. This standard score-like ratio represents the risk-adjusted expected excess return on investment compared to a risk-free investment [246]. In our case, $\mathbb{E}(\mathcal{D})$ represents the expected return on investment of DEM capabilities, $\text{Var}(\mathcal{D})$ represents the risk, and we can specify the return on investment of the risk-free investment as a constant c . This enables us to calculate the Sharpe ratio as

$$\frac{\mathbb{E}(\mathcal{D}) - c}{\sqrt{\text{Var}(\mathcal{D})}} . \quad (2.4)$$

2.3.1 Modelling values with statistical distributions

To value a DEM capability using such a generic framework that one can apply in many different ways across diverse organisations, it is first necessary to make some simplifying assumptions about the statistical properties of the business propositions under consideration. We assume the value of the business propositions ($\mathcal{V}_i$) and the estimation noises (ϵ_i) are randomly distributed:

$$\begin{aligned} \mathcal{V}_i &\stackrel{\text{i.i.d.}}{\sim} F_{\mathcal{V}}(\cdot), \text{ where } \mathbb{E}(\mathcal{V}_i) = \mu_{\mathcal{V}}, \text{Var}(\mathcal{V}_i) = \sigma_{\mathcal{V}}^2, \\ \epsilon_i &\stackrel{\text{i.i.d.}}{\sim} F_{\epsilon}(\cdot), \text{ where } \mathbb{E}(\epsilon_i) = \mu_{\epsilon}, \text{Var}(\epsilon_i) = \sigma_{\epsilon}^2, \end{aligned} \quad (2.5)$$

where $\mathcal{V}_i \perp \epsilon_j \forall i, j$ (see Figure 2.2).

We note two special cases, one when both the value and the noises are assumed to be nor-




Figure 2.2: The generative model in the ranking under lower uncertainty problem in plate notation. $\mathcal{V}_i$ represents the true, unobserved values of the items to be ranked. $\mathcal{E}_i$ represents the observed values under some estimation noise level σ_ϵ^2 .

mally distributed:

$$\mathcal{V}_i \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(\mu_{\mathcal{V}}, \sigma_{\mathcal{V}}^2), \quad \epsilon_i \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(\mu_{\epsilon}, \sigma_{\epsilon}^2); \quad (2.6)$$

and the other when both the value and the noises are assumed to follow some Generalized Student's t -distributions:

$$\begin{aligned} \mathcal{V}_i &= \mu_{\mathcal{V}} + \sigma_{\mathcal{V}} \sqrt{\frac{(\nu - 2)}{\nu}} T_i, & T_i &\stackrel{\text{i.i.d.}}{\sim} t_{\nu}, \\ \epsilon_i &= \mu_{\epsilon} + \sigma_{\epsilon} \sqrt{\frac{(\nu - 2)}{\nu}} T'_i, & T'_i &\stackrel{\text{i.i.d.}}{\sim} t_{\nu}, \end{aligned} \quad (2.7)$$

where t_{ν} is a standard Student's t -distribution with ν degrees of freedom. The location and scaling parameters ensure $\mathcal{V}_i$ and ϵ_i have the mean and variance specified in Expression (2.5).

These two cases are particularly relevant as meta-analyses compiled on the results of 6,700 e-commerce [30] and 432 marketing experiments [136], respectively, indicate the uplifts measured by the experiments, and hence the value of the business propositions under some estimation noise, exhibit the following properties:

1. They can be positive or negative,
2. They are usually clustered around an average instead of uniformly spreading across a specific range, and
3. The distributions are heavy-tailed.

The normal assumptions cover the first two properties only, yet enable one to draw on the wealth of results in order statistics and Bayesian inference related to normal distributions to get started. The t -distributed assumptions also cover property 3, though valuation under such assumptions is more complicated as t -distributions do not have conjugate priors.

We will include the valuation under t -distributed assumptions under the general case for brevity. We will, however, present empirical results in Section 2.8.1, showing that the value gained under t -distributed assumptions has a higher mean and variance, thus demonstrating that the model can capture the “higher risk, higher reward” concept.

2.3.2 Key results

In the following two sections, we will derive the expected value and variance for $\mathcal{W}$, the mean true value of the top M business propositions selected after being ranked by their estimated value (as defined in Equation (2.2)), as well as the expected value and the variance of $\mathcal{D}$, the value gained when the estimation noise is reduced. This enables us to calculate the Sharpe ratio defined in Expression (2.4).

We will also provide two key insights. Firstly, the expected mean true value of the selected business propositions ($\mathcal{W}$) increases when the estimation noise (σ_ϵ^2) decreases, and the relative increase in value depends on how much noise we can reduce. Secondly, when M is small, reducing the estimation noise may not lead to a statistically significant improvement in the true value of the business propositions selected. As a result, improvements in prioritisation driven by DEM may only be justified for larger organisations.

2.4 Calculating the Expectation

We first derive the expected value for $\mathcal{D}$. This requires the expected values of, in order:

1. $\mathcal{E}_{(r)}$ – the *estimated* value of the r^{th} ranked business proposition in estimated value;
2. $\mathcal{V}_{\mathcal{I}(r)}$ – the *true* value of the r^{th} ranked business proposition in estimated value; and
3. $\mathcal{W}$ – the mean of the *true* value for the M most valuable business propositions ranked by their estimated values.

To obtain the expected value for $\mathcal{E}_{(r)}$, we apply a result in Section 4.6 of [56], which approximates the expected value of the order statistics $\mathcal{E}_{(r)}$ using the quantile function of $\mathcal{E}_i$:

$$\mathbb{E}(\mathcal{E}_{(r)}) \approx F_{\mathcal{E}}^{-1}\left(\frac{r - c}{N - 2c + 1}\right), \quad (2.8)$$

where $F_{\mathcal{E}}^{-1}(\cdot)$ denotes the quantile function for $\mathcal{E}_i$ and c is a constant correcting the quantile of the ranks.³ The formula results from applying the first-order delta method on the underlying compound beta- $F_{\mathcal{E}}^{-1}$ distribution [223]. Such estimation will inevitably incur a small bias, particularly for extreme order statistics (i.e., r close to one or N) and when N is small. However, as verified empirically in Section 2.6, we consider such an estimation accurate enough for our application.⁴ One can also include higher-order terms in the corresponding Taylor expansion to improve estimation accuracy, though we believe the additional complexity and, thus, reduced interpretability in the resultant formula is not worth the marginal gain in accuracy.

We then obtain the expected value of $\mathcal{V}_{\mathcal{I}(r)}$ by applying Equation 6.8.3a of [56]:

$$\begin{aligned} \mathbb{E}(\mathcal{V}_{\mathcal{I}(r)}) &= \mu_{\mathcal{V}} + \rho_{\mathcal{V}\mathcal{E}} \sigma_{\mathcal{V}} \mathbb{E}\left(\frac{\mathcal{E}_{(r)} - (\mu_{\mathcal{V}} + \mu_{\epsilon})}{\sqrt{\sigma_{\mathcal{V}}^2 + \sigma_{\epsilon}^2}}\right) \\ &= \mu_{\mathcal{V}} + \frac{\sigma_{\mathcal{V}}^2}{\sigma_{\mathcal{V}}^2 + \sigma_{\epsilon}^2} (\mathbb{E}(\mathcal{E}_{(r)}) - (\mu_{\mathcal{V}} + \mu_{\epsilon})), \end{aligned} \quad (2.9)$$

where $\rho_{\mathcal{V}\mathcal{E}} = \sqrt{\sigma_{\mathcal{V}}^2} / \sqrt{\sigma_{\mathcal{V}}^2 + \sigma_{\epsilon}^2}$ is the correlation between $\mathcal{V}_i$ and $\mathcal{E}_i$.

Equation (2.9) shows that decreasing the estimation noise σ_{ϵ}^2 will increase $\mathbb{E}(\mathcal{V}_{\mathcal{I}(r)})$ for any $r > (N + 1) \cdot F_{\mathcal{E}}(\mu_{\mathcal{V}} + \mu_{\epsilon})$. It follows that the mean true value of the top M business propositions, selected according to their estimated value, will generally increase with a lower estimation noise. We show this by applying the expectation function to $\mathcal{W}$ defined in Equation (2.2)

³ Many values for c were proposed. Early works advocate $c = 0$ or $c = 0.5$ depending on how one approaches continuity correction between ranks and quantiles [119]. [27] proposes a compromise value of $c = \frac{3}{8}$ based on a tabulation of c required to yield the correct expected value for different r and N . [119] expands the tabulation and proposes using a different c for each N , with the values of c hovering around 0.4. For simplicity, we take $c = 0.4$ for all N in all our calculations.

⁴ [119] also noted for normal order statistics, the maximum error is 0.018 (i.e., less than one percent relative to the value of extreme order statistics) for some choice of c .

to obtain

$$\begin{aligned}\mathbb{E}(\mathcal{W}) &= \frac{1}{M} \sum_{r=N-M+1}^N \mathbb{E}(\mathcal{V}_{\mathcal{I}(r)}) \\ &= \mu_{\mathcal{V}} + \frac{\sigma_{\mathcal{V}}^2}{\sigma_{\mathcal{V}}^2 + \sigma_{\epsilon}^2} \left[\left(\frac{1}{M} \sum_{r=N-M+1}^N \mathbb{E}(\mathcal{E}_{(r)}) \right) - (\mu_{\mathcal{V}} + \mu_{\epsilon}) \right].\end{aligned}\quad (2.10)$$

We finally consider the improvement when we reduce the estimation noise from $\sigma_{\epsilon}^2 = \sigma_1^2$ to $\sigma_{\epsilon}^2 = \sigma_2^2$. This will be the expected value gained by having better DEM capabilities:

$$\begin{aligned}\mathbb{E}(\mathcal{D}) &= \mathbb{E}(\mathcal{W} | \sigma_{\epsilon}^2 = \sigma_2^2) - \mathbb{E}(\mathcal{W} | \sigma_{\epsilon}^2 = \sigma_1^2) \\ &= \frac{\sigma_{\mathcal{V}}^2}{\sigma_{\mathcal{V}}^2 + \sigma_2^2} \left[\left(\frac{1}{M} \sum_{r=N-M+1}^N \mathbb{E}(\mathcal{E}_{(r)} | \sigma_{\epsilon}^2 = \sigma_2^2) \right) - (\mu_{\mathcal{V}} + \mu_{\epsilon}) \right] - \\ &\quad \frac{\sigma_{\mathcal{V}}^2}{\sigma_{\mathcal{V}}^2 + \sigma_1^2} \left[\left(\frac{1}{M} \sum_{r=N-M+1}^N \mathbb{E}(\mathcal{E}_{(r)} | \sigma_{\epsilon}^2 = \sigma_1^2) \right) - (\mu_{\mathcal{V}} + \mu_{\epsilon}) \right].\end{aligned}\quad (2.11)$$

2.4.1 Expectation under normal assumptions

In the special case where $\mathcal{E}_i$ are normally distributed (with mean $\mu_{\mathcal{V}} + \mu_{\epsilon}$ and variance $\sigma_{\mathcal{V}}^2 + \sigma_{\epsilon}^2$), the expected value for the normal order statistics $\mathcal{E}_{(r)}$ is approximately

$$\mathbb{E}(\mathcal{E}_{(r)}) \approx \mu_{\mathcal{V}} + \mu_{\epsilon} + \sqrt{\sigma_{\mathcal{V}}^2 + \sigma_{\epsilon}^2} \Phi^{-1} \left(\frac{r - c}{N - 2c + 1} \right), \quad (2.12)$$

where $\Phi^{-1}(\cdot)$ denotes the quantile function of a standard normal distribution [27]. It is worth noting that decreasing the estimation noise σ_{ϵ}^2 will decrease $\mathbb{E}(\mathcal{E}_{(r)})$ for any $r > \frac{N+1}{2}$, appearing to lower the average value of the top M business propositions. This is a common pitfall – we are not optimising the combined estimated value of the selected business propositions. What matters is the true yet unobserved value of that proposition, $\mathcal{V}_{\mathcal{I}(r)}$, as shown below.

For $\mathcal{V}_{\mathcal{I}(r)}$, we can simplify Equation (2.9) by substituting Equation (2.12). We can also evaluate from first principles by noting a standard result in Bayesian inference, which states that the posterior distribution of $\mathcal{V}_i$ once $\mathcal{E}_i$ is observed is also normally distributed with mean

$$\mu_{\mathcal{V}_i | (\mathcal{E}_i = e)} = \frac{\sigma_{\mathcal{V}}^2}{\sigma_{\mathcal{V}}^2 + \sigma_{\epsilon}^2} (e - \mu_{\epsilon}) + \frac{\sigma_{\epsilon}^2}{\sigma_{\mathcal{V}}^2 + \sigma_{\epsilon}^2} \mu_{\mathcal{V}}, \quad (2.13)$$

and applying the law of iterated expectations to obtain⁵

$$\mathbb{E}(\mathcal{V}_{\mathcal{I}(r)}) = \mathbb{E}(\mathbb{E}(\mathcal{V}_{\mathcal{I}(r)} | \mathcal{E}_{(r)})) \approx \mu_{\mathcal{V}} + \frac{\sigma_{\mathcal{V}}^2}{\sqrt{\sigma_{\mathcal{V}}^2 + \sigma_{\epsilon}^2}} \Phi^{-1} \left(\frac{r - c}{N - 2c + 1} \right). \quad (2.14)$$

Here, decreasing the estimation noise σ_{ϵ}^2 will lead to an increase in $\mathbb{E}(\mathcal{V}_{\mathcal{I}(r)})$ for any $r > \frac{N+1}{2}$.⁶

The value of business propositions chosen ($\mathcal{W}$) under normal assumptions then evaluates to

$$\mathbb{E}(\mathcal{W}) \approx \mu_{\mathcal{V}} + \frac{\sigma_{\mathcal{V}}^2}{\sqrt{\sigma_{\mathcal{V}}^2 + \sigma_{\epsilon}^2}} \frac{1}{M} \sum_{r=N-M+1}^N \Phi^{-1} \left(\frac{r - c}{N - 2c + 1} \right). \quad (2.15)$$

This is done by substituting Equation (2.14) into Equation (2.10). Note the absence of μ_{ϵ} in Equation (2.15), which suggests that systematic bias in estimation will not affect the true value of the chosen business propositions in the normal case.

Finally, the expression for the expected value of $\mathcal{D}$ when we reduce the estimation noise from $\sigma_{\epsilon}^2 = \sigma_1^2$ to σ_2^2 is much neater under normal assumptions, as many terms cancel out in Equation (2.11), leading to

$$\mathbb{E}(\mathcal{D}) \approx \left(\frac{\sigma_{\mathcal{V}}^2}{\sqrt{\sigma_{\mathcal{V}}^2 + \sigma_2^2}} - \frac{\sigma_{\mathcal{V}}^2}{\sqrt{\sigma_{\mathcal{V}}^2 + \sigma_1^2}} \right) \frac{1}{M} \sum_{r=N-M+1}^N \Phi^{-1} \left(\frac{r - c}{N - 2c + 1} \right). \quad (2.16)$$

If we further assume that $\mu_{\mathcal{V}} = 0$ (i.e., the true value of the business propositions centres around zero), then the relative gain is entirely dependent on $\sigma_{\mathcal{V}}^2$, σ_1^2 , and σ_2^2 :

$$\frac{\mathbb{E}(\mathcal{D} | \mu_{\mathcal{V}} = 0)}{\mathbb{E}(\mathcal{W} | \sigma_{\epsilon}^2 = \sigma_1^2, \mu_{\mathcal{V}} = 0)} = \frac{\sqrt{\sigma_{\mathcal{V}}^2 + \sigma_1^2}}{\sqrt{\sigma_{\mathcal{V}}^2 + \sigma_2^2}} - 1. \quad (2.17)$$

To calculate the relative improvement in prioritisation delivered by DEM under these assumptions, we plug the following into Equation (2.17) to obtain an estimate of how much one will gain from acquiring such capabilities:

1. The estimated spread of the values ($\sigma_{\mathcal{V}}^2$),
2. The estimated deviation of the current estimation process (σ_1^2), and

⁵ $\mathcal{E}_{(r)} \equiv \mathcal{E}_{\mathcal{I}(r)}$, as we rank the business propositions by their estimated values.

⁶ While $\mathcal{V}_{\mathcal{I}(r)}$ are no longer normally distributed with the ranking information, it remains valid to use the law of iterated expectations to obtain the expected value for all r .

3. The estimated deviation to the actual value upon acquisition of DEM capabilities (σ_2^2).

For example, if Example Company Ltd's project values are spread with a standard deviation of 1 unit and their current estimation has a standard error of 0.5 units, then by acquiring an A/B test framework that is capable of measuring with an error of 0.4 units, the company gains 3.8% of extra value simply due to the ability to prioritise with more accurate measurements under normal assumptions.

We conclude this section by suggesting how one may estimate σ_V^2 , σ_1^2 , and σ_2^2 , especially when they have yet to build any DEM capabilities. Clearly, it is impossible to recommend specific values for the three parameters as organisations seeking to build DEM capabilities come in all shapes and sizes. That said, one may consider using industry averages published in meta-analyses [30, 136] when estimating σ_V^2 . For σ_1^2 , they may solicit value estimates for several business propositions from current decision makers and take the variance over such estimates.⁷ Finally, they may leverage the user numbers and the variance of user responses that form the business metric to estimate σ_2^2 under an A/B test.⁸ We also encourage one to explore parameter values around their estimates for robustness, as we will show in Section 2.7 when we provide two case studies.

2.5 Calculating the Variance

It is crucial to understand both the expected gain and the risk or uncertainty of ranking under lower uncertainty to make effective investment decisions. Therefore, having derived the expected value in Equation (2.11), we address the investment risk given by the variance of $\mathcal{D}$ in this section.

The variance calculation features new challenges in addition to that identified in the section

⁷ It does not matter if the decision makers provide value estimates far from the actual value of the business propositions. In fact, one need not know the business propositions' actual value. What matters here is the spread of the estimates.

⁸ For example, in an A/B test measuring conversion rate (CVR) uplift using a practical t -test, the estimation noise (σ_ϵ^2 or σ_2^2) is the variance of the difference in CVR between two variants under a nil null hypothesis. This equals $2 \cdot \frac{p(1-p)}{n/2} = \frac{4p(1-p)}{n}$, where p is the base CVR and n is the total number of users across both groups. See Chapter 3 for further discussions on statistical tests, particularly Sections 3.5.2 and 3.5.3 on practical t -tests and how they are applied to business metrics based on binary responses such as CVR.




Figure 2.3: The generative model in the ranking under lower uncertainty problem when two distinct noise levels are involved in plate notation. When we change the noise level of our ranking under uncertainty setup from σ_1^2 to σ_2^2 (see Figure 2.2), we obtain two sets of observed values, $\mathcal{H}_i$ and $\mathcal{L}_i$, for each noise level.

above, the most prominent of which concerns the interactions between quantities generated under different estimation noise levels. While these interactions do not affect the expected value, they influence the variance via the covariance terms. Failure to account for the covariance terms may lead to a large error in the variance estimate.

To address the challenges, we first extend the notation to clarify the interactions. Instead of a single set of noise shown in Section 2.3, we define two sets of random noise with different noise levels:

$$\begin{aligned} \epsilon_{1i} &\stackrel{\text{i.i.d.}}{\sim} F_{\epsilon_1}(\cdot), \text{ where } \mathbb{E}(\epsilon_{1i}) = \mu_\epsilon, \text{Var}(\epsilon_{1i}) = \sigma_1^2, \text{ and} \\ \epsilon_{2i} &\stackrel{\text{i.i.d.}}{\sim} F_{\epsilon_2}(\cdot), \text{ where } \mathbb{E}(\epsilon_{2i}) = \mu_\epsilon, \text{Var}(\epsilon_{2i}) = \sigma_2^2. \end{aligned} \quad (2.18)$$

Similar to how the problem is modelled in Section 2.3, we assume all true item values ($\mathcal{V}$) and noises (ϵ_1 and ϵ_2) are independent, i.e., $\mathcal{V}_i \perp \epsilon_{1j}$, $\mathcal{V}_i \perp \epsilon_{2j}$, and $\epsilon_{1i} \perp \epsilon_{2j} \forall i, j$. The estimated value of each item is then given by

$$\begin{aligned} \mathcal{H}_i &= \mathcal{V}_i + \epsilon_{1i} \quad \text{under Higher noise level } \sigma_1^2 \text{ and} \\ \mathcal{L}_i &= \mathcal{V}_i + \epsilon_{2i} \quad \text{under Lower noise level } \sigma_2^2. \end{aligned} \quad (2.19)$$

The setup is illustrated in Figure 2.3.




Figure 2.4: Relationship between different variances/covariances used to calculate the variance of $\mathcal{D}$, the value gained when the estimation noise is reduced. An arrow from quantity A to B means the value of B depends on the value of A.

Having obtained two sets of estimated values, we rank and trace the corresponding indices for each set separately. For $\mathcal{H}$, we denote $\mathcal{H}_{(r)}$ as the r^{th} order statistic of $\mathcal{H}_i$, the estimated value of the r^{th} ranked item under noise level σ_1^2 , followed by $\mathcal{V}_{\mathcal{I}(r)}$ as the concomitant [56] of $\mathcal{H}_{(r)}$, i.e., the true value of the r^{th} item ranked by its estimated value. We repeat the process for $\mathcal{L}$: we denote $\mathcal{L}_{(s)}$ as the s^{th} order statistic of $\mathcal{L}_i$ and $\mathcal{V}_{\mathcal{J}(s)}$ as the concomitant of $\mathcal{L}_{(s)}$.⁹

We also define the mean true value of the top M items, ranked by their estimated value, under both noise levels as

$$\mathcal{W}_1 = \frac{1}{M} \sum_{r=N-M+1}^N \mathcal{V}_{\mathcal{I}(r)}, \quad \mathcal{W}_2 = \frac{1}{M} \sum_{s=N-M+1}^N \mathcal{V}_{\mathcal{J}(s)}, \quad (2.20)$$

where $\mathcal{W}_1$ is the mean true value under σ_1^2 , and $\mathcal{W}_2$ is the mean true value under σ_2^2 . Finally, we denote the difference between the mean true values as $\mathcal{D} \triangleq \mathcal{W}_2 - \mathcal{W}_1$.

Deriving the variance is similar to deriving the expectation – one has to obtain the variances for (in order) $\mathcal{H}_{(r)}/\mathcal{L}_{(s)}$, $\mathcal{V}_{\mathcal{I}(r)}/\mathcal{V}_{\mathcal{J}(s)}$, $\mathcal{W}_1/\mathcal{W}_2$, and $\mathcal{D}$. The relationship between these quantities is shown in Figure 2.4.

⁹ $\mathcal{I}(\cdot)$ and $\mathcal{J}(\cdot)$ are the index functions that map the ranking to the index for $\mathcal{H}$ and $\mathcal{L}$, respectively.

2.5.1 $\text{Var}(\mathcal{H}_{(r)})/\text{Var}(\mathcal{L}_{(s)})$

We apply a result from [54], which states that the variance of $\mathcal{H}_{(r)}$ and $\mathcal{L}_{(s)}$ can be approximated as

$$\text{Var}(\mathcal{H}_{(r)}) \approx \frac{r(N-r+1)}{(N+1)^2(N+2)} \frac{1}{\left(f_{\mathcal{H}}\left(F_{\mathcal{H}}^{-1}\left(\frac{r}{N+1}\right)\right)\right)^2}, \quad (2.21)$$

$$\text{Var}(\mathcal{L}_{(s)}) \approx \frac{s(N-s+1)}{(N+1)^2(N+2)} \frac{1}{\left(f_{\mathcal{L}}\left(F_{\mathcal{L}}^{-1}\left(\frac{s}{N+1}\right)\right)\right)^2}, \quad (2.22)$$

where $f(\cdot)$ and $F^{-1}(\cdot)$ denote the probability density function and quantile function of the corresponding r.v., respectively. As is the case with calculating the expected value (see Section 2.4), the formula results from applying the first-order delta method [171], with the same strength and weakness considerations. In the special case where $\mathcal{V}_i$, ϵ_{1i} , and ϵ_{2i} are all normally distributed, the variances are

$$\text{Var}(\mathcal{H}_{(r)}) \approx \frac{r(N-r+1)}{(N+1)^2(N+2)} \frac{\sigma_X^2 + \sigma_1^2}{\left(\phi\left(\Phi^{-1}\left(\frac{r}{N+1}\right)\right)\right)^2}, \quad (2.23)$$

$$\text{Var}(\mathcal{L}_{(s)}) \approx \frac{s(N-s+1)}{(N+1)^2(N+2)} \frac{\sigma_X^2 + \sigma_2^2}{\left(\phi\left(\Phi^{-1}\left(\frac{s}{N+1}\right)\right)\right)^2}, \quad (2.24)$$

where $\phi(\cdot)$ is the probability density function, and $\Phi^{-1}(\cdot)$ is the quantile function of a standard normal distribution.

2.5.2 $\text{Var}(\mathcal{V}_{\mathcal{I}(r)})/\text{Var}(\mathcal{V}_{\mathcal{J}(s)})$

The variance for $\mathcal{V}_{\mathcal{I}(r)}$ is obtained using properties of the concomitants of order statistics [55]:¹⁰

$$\text{Var}(\mathcal{V}_{\mathcal{I}(r)}) = \sigma_{\mathcal{V}}^2 \left(\rho_{\mathcal{V}\mathcal{H}}^2 \frac{\text{Var}(\mathcal{H}_{(r)})}{\sigma_{\mathcal{V}}^2 + \sigma_1^2} + 1 - \rho_{\mathcal{V}\mathcal{H}}^2 \right) = \frac{\sigma_1^2 \sigma_{\mathcal{V}}^2}{\sigma_{\mathcal{V}}^2 + \sigma_1^2} + \frac{\sigma_{\mathcal{V}}^4}{(\sigma_{\mathcal{V}}^2 + \sigma_1^2)^2} \text{Var}(\mathcal{H}_{(r)}), \quad (2.25)$$

$$\text{Var}(\mathcal{V}_{\mathcal{J}(s)}) = \sigma_{\mathcal{V}}^2 \left(\rho_{\mathcal{V}\mathcal{L}}^2 \frac{\text{Var}(\mathcal{L}_{(s)})}{\sigma_{\mathcal{V}}^2 + \sigma_2^2} + 1 - \rho_{\mathcal{V}\mathcal{L}}^2 \right) = \frac{\sigma_2^2 \sigma_{\mathcal{V}}^2}{\sigma_{\mathcal{V}}^2 + \sigma_2^2} + \frac{\sigma_{\mathcal{V}}^4}{(\sigma_{\mathcal{V}}^2 + \sigma_2^2)^2} \text{Var}(\mathcal{L}_{(s)}), \quad (2.26)$$

¹⁰ $\mathcal{H}/\mathcal{L}$ and $\mathcal{V}$ in this thesis correspond to X and Y in [55].

where $\rho_{\mathcal{V}\mathcal{H}} = \sigma_{\mathcal{V}} / \sqrt{\sigma_{\mathcal{V}}^2 + \sigma_1^2}$ denotes the correlation between $\mathcal{V}_i$ and $\mathcal{H}_i$, and $\rho_{\mathcal{V}\mathcal{L}} = \sigma_{\mathcal{V}} / \sqrt{\sigma_{\mathcal{V}}^2 + \sigma_2^2}$ denotes the correlation between $\mathcal{V}_i$ and $\mathcal{L}_i$.

In the multivariate normal case, the same result can also be obtained from first principles. We first recall Equation (2.13), which describes the mean of $\mathcal{V}_i$ once we observe $\mathcal{E}_i$ (now $\mathcal{H}_i$ or $\mathcal{L}_i$), and note its variance counterpart is

$$\sigma_{\mathcal{V}_i | \mathcal{H}_i}^2 = \frac{\sigma_1^2 \sigma_{\mathcal{V}}^2}{\sigma_{\mathcal{V}}^2 + \sigma_1^2}, \quad \sigma_{\mathcal{V}_i | \mathcal{L}_i}^2 = \frac{\sigma_2^2 \sigma_{\mathcal{V}}^2}{\sigma_{\mathcal{V}}^2 + \sigma_2^2}. \quad (2.27)$$

Using the results in Equations (2.13) and (2.27), we apply the law of total variance to obtain¹¹

$$\begin{aligned} \text{Var}(\mathcal{V}_{\mathcal{I}(r)}) &= \mathbb{E}(\text{Var}(\mathcal{V}_{\mathcal{I}(r)} | \mathcal{H}_{(r)})) + \text{Var}(\mathbb{E}(\mathcal{V}_{\mathcal{I}(r)} | \mathcal{H}_{(r)})) \\ &\approx \frac{\sigma_1^2 \sigma_{\mathcal{V}}^2}{\sigma_{\mathcal{V}}^2 + \sigma_1^2} + \frac{\sigma_{\mathcal{V}}^4}{\sigma_{\mathcal{V}}^2 + \sigma_1^2} \frac{r(N-r+1)}{(N+1)^2(N+2)} \frac{1}{(\phi(\Phi^{-1}(\frac{r}{N+1})))^2}, \end{aligned} \quad (2.28)$$

$$\begin{aligned} \text{Var}(\mathcal{V}_{\mathcal{J}(s)}) &= \mathbb{E}(\text{Var}(\mathcal{V}_{\mathcal{J}(s)} | \mathcal{L}_{(s)})) + \text{Var}(\mathbb{E}(\mathcal{V}_{\mathcal{J}(s)} | \mathcal{L}_{(s)})) \\ &\approx \frac{\sigma_2^2 \sigma_{\mathcal{V}}^2}{\sigma_{\mathcal{V}}^2 + \sigma_2^2} + \frac{\sigma_{\mathcal{V}}^4}{\sigma_{\mathcal{V}}^2 + \sigma_2^2} \frac{s(N-s+1)}{(N+1)^2(N+2)} \frac{1}{(\phi(\Phi^{-1}(\frac{s}{N+1})))^2}. \end{aligned} \quad (2.29)$$

2.5.3 $\text{Var}(\mathcal{W}_1) / \text{Var}(\mathcal{W}_2)$

To derive the variance of $\mathcal{W}_1$, we require the covariance between $\mathcal{H}_{(r)}$ and $\mathcal{H}_{(s)}$, as well as that between $\mathcal{V}_{\mathcal{I}(r)}$ and $\mathcal{V}_{\mathcal{I}(s)}$ $\forall r, s$. The same goes for $\mathcal{W}_2$, where we require the covariance between $\mathcal{L}_{(r)}$ and $\mathcal{L}_{(s)}$, as well as that between $\mathcal{V}_{\mathcal{J}(r)}$ and $\mathcal{V}_{\mathcal{J}(s)}$ $\forall r, s$. This is necessary as the terms of $\mathcal{W}_1$ (see Equation (2.2)), which result from removing noise from successive order statistics, are highly correlated.

Equation 4.6.5 of [56] provided a formula to estimate the covariance between $\mathcal{H}_{(r)}$ and $\mathcal{H}_{(s)}$

¹¹ Similar to what has been mentioned in Footnotes 5 and 6, $\mathcal{H}_{(r)} \equiv \mathcal{H}_{\mathcal{I}(r)}$ and $\mathcal{L}_{(s)} \equiv \mathcal{L}_{\mathcal{J}(s)}$, since we rank the business propositions by their estimated values. Moreover, while $\mathcal{V}_{\mathcal{I}(r)}$ and $\mathcal{V}_{\mathcal{J}(s)}$ are no longer normally distributed with the ranking information, using the law of total variance to obtain the variance for all r and s remains valid.

and between $\mathcal{L}_{(r)}$ and $\mathcal{L}_{(s)}$ for any $r < s \leq N$:¹²

$$\text{Cov}(\mathcal{H}_{(r)}, \mathcal{H}_{(s)}) \approx \frac{r(N-s+1)}{(N+1)^2(N+2)} \frac{1}{f_{\mathcal{H}}(F_{\mathcal{H}}^{-1}(\frac{r}{N+1})) f_{\mathcal{H}}(F_{\mathcal{H}}^{-1}(\frac{s}{N+1}))}, \quad (2.30)$$

$$\text{Cov}(\mathcal{L}_{(r)}, \mathcal{L}_{(s)}) \approx \frac{r(N-s+1)}{(N+1)^2(N+2)} \frac{1}{f_{\mathcal{L}}(F_{\mathcal{L}}^{-1}(\frac{r}{N+1})) f_{\mathcal{L}}(F_{\mathcal{L}}^{-1}(\frac{s}{N+1}))}. \quad (2.31)$$

To obtain the covariance between $\mathcal{V}_{\mathcal{I}(r)}$ and $\mathcal{V}_{\mathcal{I}(s)}$ for any $r, s \leq N$, we again refer to [55] (Equation 2.3d):

$$\text{Cov}(\mathcal{V}_{\mathcal{I}(r)}, \mathcal{V}_{\mathcal{I}(s)}) = \rho_{\mathcal{V}\mathcal{H}}^2 \sigma_{\mathcal{V}}^2 \frac{\text{Cov}(\mathcal{H}_{(r)}, \mathcal{H}_{(s)})}{\sigma_{\mathcal{V}}^2 + \sigma_1^2} = \frac{\sigma_{\mathcal{V}}^4}{(\sigma_{\mathcal{V}}^2 + \sigma_1^2)^2} \text{Cov}(\mathcal{H}_{(r)}, \mathcal{H}_{(s)}), \quad (2.32)$$

$$\text{Cov}(\mathcal{V}_{\mathcal{J}(r)}, \mathcal{V}_{\mathcal{J}(s)}) = \rho_{\mathcal{V}\mathcal{L}}^2 \sigma_{\mathcal{V}}^2 \frac{\text{Cov}(\mathcal{L}_{(r)}, \mathcal{L}_{(s)})}{\sigma_{\mathcal{V}}^2 + \sigma_2^2} = \frac{\sigma_{\mathcal{V}}^4}{(\sigma_{\mathcal{V}}^2 + \sigma_2^2)^2} \text{Cov}(\mathcal{L}_{(r)}, \mathcal{L}_{(s)}). \quad (2.33)$$

Equation (2.32) and (2.33) affirms the claim that $\mathcal{V}_{\mathcal{I}(r)}$ are positively correlated $\forall r$. Unlike $\mathcal{V}_i$, which are independent by definition, they become correlated under the presence of ranking information.

We can now state the variance of $\mathcal{W}_1$ and $\mathcal{W}_2$. Applying the variance function to Equation (2.20), we get

$$\text{Var}(\mathcal{W}_1) = \frac{1}{M^2} \left(\sum_{r=N-M+1}^N \text{Var}(\mathcal{V}_{\mathcal{I}(r)}) + \sum_{r=N-M+1}^N \sum_{s=r+1}^N 2 \cdot \text{Cov}(\mathcal{V}_{\mathcal{I}(r)}, \mathcal{V}_{\mathcal{I}(s)}) \right), \quad (2.34)$$

$$\text{Var}(\mathcal{W}_2) = \frac{1}{M^2} \left(\sum_{s=N-M+1}^N \text{Var}(\mathcal{V}_{\mathcal{J}(s)}) + \sum_{r=N-M+1}^N \sum_{s=r+1}^N 2 \cdot \text{Cov}(\mathcal{V}_{\mathcal{J}(r)}, \mathcal{V}_{\mathcal{J}(s)}) \right), \quad (2.35)$$

where the constituent variances and covariances are derived in Equations (2.25)/(2.26) and (2.32)/(2.33), respectively.

2.5.4 Var($\mathcal{D}$)

Finally, we derive the variance of $\mathcal{D}$. In addition to the variance of $\mathcal{W}_1$ and $\mathcal{W}_2$ derived, respectively, in Equations (2.34) and (2.35), we require the covariance between these two

¹² For $r > s$, simply swap r and s as covariance functions are symmetrical.

terms. This, in turn, requires the covariance between $\mathcal{H}_{(r)}$ and $\mathcal{L}_{(s)}$ and that between $X_{\mathcal{I}(r)}$ and $X_{\mathcal{J}(s)}$.

The covariance between $\mathcal{H}_{(r)}$ and $\mathcal{L}_{(s)}$ can be derived using results in [56]:

$$\begin{aligned} \text{Cov}(\mathcal{H}_{(r)}, \mathcal{L}_{(s)}) &= \rho_{\mathcal{V}\mathcal{H}}\rho_{\mathcal{V}\mathcal{L}} \text{Cov}(\mathcal{V}_{(r)}, \mathcal{V}_{(s)}) \\ &= \frac{\sigma_{\mathcal{V}}^2}{\sqrt{\sigma_{\mathcal{V}}^2 + \sigma_1^2}\sqrt{\sigma_{\mathcal{V}}^2 + \sigma_2^2}} \frac{r(N-s+1)}{(N+1)^2(N+2)} \frac{1}{f_{\mathcal{V}}(F_{\mathcal{V}}^{-1}(\frac{r}{N+1})) f_{\mathcal{V}}(F_{\mathcal{V}}^{-1}(\frac{s}{N+1}))}, \end{aligned} \quad (2.36)$$

where $f_{\mathcal{V}}(\cdot)$ and $F_{\mathcal{V}}^{-1}(\cdot)$ are the probability density function and quantile function for $\mathcal{V}_i$, respectively.

Deriving the covariance between $\mathcal{V}_{\mathcal{I}(r)}$ and $\mathcal{V}_{\mathcal{J}(s)}$ is perhaps the most challenging problem within the work, as they take two forms depending on the indices:

$$\text{Cov}(\mathcal{V}_{\mathcal{I}(r)}, \mathcal{V}_{\mathcal{J}(s)}) = \begin{cases} \text{Var}(\mathcal{V}_{\mathcal{I}(r)}) = \text{Var}(\mathcal{V}_{\mathcal{J}(s)}) & \text{if } \mathcal{I}(r) = \mathcal{J}(s) \\ \frac{\sigma_{\mathcal{V}}^2}{\sigma_{\mathcal{V}}^2 + \sigma_1^2} \frac{\sigma_{\mathcal{V}}^2}{\sigma_{\mathcal{V}}^2 + \sigma_2^2} \text{Cov}(\mathcal{H}_{(r)}, \mathcal{L}_{(s)}) & \text{if } \mathcal{I}(r) \neq \mathcal{J}(s) \end{cases}, \quad (2.37)$$

where the second case is a standard Bayesian inference result.

The problem arises as the r^{th} ranked $\mathcal{H}$ and the s^{th} ranked $\mathcal{L}$ can be generated by the same $\mathcal{V}_i$ for some i . It will not arise if we have only $\mathcal{H}$ or $\mathcal{L}$ (see Figure 2.5 for an example). In this case, when we consider the covariance of the concomitants $\mathcal{V}_{\mathcal{I}(r)}/\mathcal{V}_{\mathcal{J}(s)}$, we have to take into account both the existing variance of $\mathcal{V}_i$ and the ranking information provided by $\mathcal{H}_{(r)}$ and $\mathcal{L}_{(s)}$. If the order statistics are generated by different $\mathcal{V}$, we only need to consider the latter as we assumed $\mathcal{V}_i$ to be independent and thus uncorrelated.

As we are interested in the overall behaviour, we only need to derive the two cases on the RHS of Equation (2.37) and weigh them using the probability that $\mathcal{I}(r) = \mathcal{J}(s)$ without worrying about which case applies to each (r, s) pair. The first case (when $\mathcal{I}(r) = \mathcal{J}(s)$) can




Figure 2.5: Relationship between different quantities in a three-item generative model under the ranking under lower uncertainty problem. $\mathcal{V}_i$, $\mathcal{H}_i/\mathcal{L}_i$, and $\mathcal{H}_{(r)}/\mathcal{L}_{(s)}$ represent the true value, the unranked noisy estimates, and the ranked noisy estimates of the items, respectively, for $i, r, s \in \{1, 2, 3\}$. (L) Under one estimation noise level, $\exists$ a bijection between $\mathcal{V}_i$ and $\mathcal{H}_{(r)}$ within a set of generated samples. (R) With two noise levels, $\mathcal{H}_{(r)}$ and $\mathcal{L}_{(s)}$ may be generated by the same $\mathcal{V}_i$ for some r and s .

be evaluated using the law of total variance with multiple conditioning random variables:

$$\begin{aligned}
 \text{Var}(\mathcal{V}_{\mathcal{I}(r)}) &= \text{Var}(\mathcal{V}_{\mathcal{J}(s)}) \\
 &= \mathbb{E}(\text{Var}(\mathcal{V}_{\mathcal{I}(r)}|\mathcal{H}_{(r)}, \mathcal{L}_{(s)})) + \mathbb{E}(\text{Var}(\mathbb{E}(\mathcal{V}_{\mathcal{I}(r)}|\mathcal{H}_{(r)}, \mathcal{L}_{(s)})|\mathcal{H}_{(r)})) + \text{Var}(\mathbb{E}(\mathcal{V}_{\mathcal{I}(r)}|\mathcal{H}_{(r)})) \\
 &= \frac{\sigma_{\mathcal{V}}^2 \sigma_1^2 \sigma_2^2}{\sigma_{\mathcal{V}}^2 \sigma_1^2 + \sigma_{\mathcal{V}}^2 \sigma_2^2 + \sigma_1^2 \sigma_2^2} + \left(\frac{\sigma_{\mathcal{V}}^2 \sigma_1^2}{\sigma_{\mathcal{V}}^2 \sigma_1^2 + \sigma_{\mathcal{V}}^2 \sigma_2^2 + \sigma_1^2 \sigma_2^2} \right)^2 \text{Var}(\mathcal{L}_{(s)}) + \left(\frac{\sigma_{\mathcal{V}}^2}{\sigma_{\mathcal{V}}^2 + \sigma_1^2} \right)^2 \text{Var}(\mathcal{H}_{(r)}).
 \end{aligned} \tag{2.38}$$

The second case can be derived by substituting Equation (2.36) into Equation (2.37).

For the weighting probability $\mathbb{P}(\mathcal{I}(r) = \mathcal{J}(s))$, we see its derivation as an interesting and potentially important problem in its own right, yet to the best of our knowledge, no proper treatment was given to the problem. In this work, we approximate the probability using beta-binomial distributions, with parameters derived from quantities calculated above. Without distracting readers from the main question of quantifying the value and risk of DEM capabilities, we relegate the detailed discussion on approximating the quantity to Appendix A.1.

With the three components for the covariance between $\mathcal{V}_{\mathcal{I}(r)}$ and $\mathcal{V}_{\mathcal{J}(s)}$ in place, we can finally derive $\text{Cov}(\mathcal{W}_1, \mathcal{W}_2)$ and $\text{Var}(\mathcal{D})$ by applying the covariance and variance functions to the

definitions of, respectively, $\mathcal{W}_1/\mathcal{W}_2$ and $\mathcal{D}$ (see (2.20)) to obtain

$$\text{Cov}(\mathcal{W}_1, \mathcal{W}_2) = \frac{1}{M^2} \sum_{r=N-M+1}^N \sum_{s=N-M+1}^N \text{Cov}(\mathcal{V}_{\mathcal{I}(r)}, \mathcal{V}_{\mathcal{J}(s)}), \quad (2.39)$$

$$\text{Var}(\mathcal{W}) = \text{Var}(\mathcal{W}_2 - \mathcal{W}_1) = \text{Var}(\mathcal{W}_1) + \text{Var}(\mathcal{W}_2) - 2 \text{Cov}(\mathcal{W}_1, \mathcal{W}_2), \quad (2.40)$$

where the first two terms on the RHS of Equation (2.40) are that derived in Equations (2.34) and (2.35).

We conclude this section by observing that M and N influence $\text{Var}(\mathcal{D})$ considerably. In particular, $\text{Var}(\mathcal{D})$ is generally large when M and N are small with other parameters fixed. This is crucial as even in cases where $\mathbb{E}(\mathcal{D})$ is positive, the limited capacity of an organisation to introduce new business propositions may mean that the Sharpe ratio defined in Section 2.3 (see Expression (2.4)) may not be high enough to justify investment in a DEM capability.

The exact threshold where an organisation should consider acquiring such capabilities depends on multiple factors, including their size (which affects M), the size of their backlog (N), the nature of their work ($\mu_{\mathcal{V}}$ and $\sigma_{\mathcal{V}}^2$), and how good they were at estimation (σ_1^2). Thus, we refrain from providing a one-size-fits-all recommendation but give examples in Section 2.7.

2.6 Empirical Verification

Having performed theoretical calculations for the expectation and variance of the value DEM capabilities deliver through enhanced prioritisation, we verify those calculations using simulation results.¹³

We verify the result derived in Sections 2.4 and 2.5 empirically, in particular under the normal assumptions. We run multiple *statistical tests* for each quantity of interest – the mean and variance of $\mathcal{W}_1/\mathcal{W}_2$ and $\mathcal{D}$, as well as the covariance between different pairs of order statistics and their concomitants.

¹³ All code used in the experiments, case studies and extensions are available on GitHub: https://github.com/liuchbryan/ranking_under_lower_uncertainty.

In each *statistical test*, we randomly select and fix the value of the parameters, i.e., N , M , μ_V , μ_ϵ , σ_V^2 , σ_1^2 , σ_2^2 , r , and s (the latter two for the covariance of the order statistics only). We select N uniformly on a log scale, with the log scale favoured due to its better resemblance to the organisation size distribution in real life and bounds chosen to represent the usual number of items/propositions an organisation will consider during prioritisation. We also select μ_V , μ_ϵ , σ_V^2 , and σ_1^2 randomly from a uniform distribution (on a linear scale), with bounds chosen that are reflective of the usual business metrics an organisation uses to prioritise their items/propositions and their unit of measurement (e.g., percentage points change, revenue impact in hundred thousand pounds).¹⁴

Each parameter is drawn independently from other parameters where possible, with the only exceptions being when the range of a parameter is restricted (i.e., $M < N$, $\sigma_2^2 < \sigma_1^2$, and $N - M + 1 \leq r, s, \leq N$ for selected items/propositions). To draw M , we draw a uniformly random multiplier between 0.01 and 0.8, then apply the multiplier to N .¹⁵ We apply a similar process to draw σ_2^2 , though we use a different multiplier range and apply the multiplier on σ_1 (the standard deviation) before squaring.¹⁶ Finally, we draw r and s independently and uniformly amongst integers between $N - M + 1$ and N (both inclusive).

Upon selecting and fixing the value of the parameters, we compare the theoretical value of the quantity of interest to the centred 95% confidence interval (CI) generated from multiple *empirical samples*. If the derivations above are exact, the 95% CI should contain the theoretical value in around 95% of the statistical tests. The histogram of the percentile rank of the theoretical quantity among the empirical samples should also follow a uniform distribution [259].

Each *empirical sample* is generated using one of the following two methods depending on the quantity we are evaluating:

a) Bootstrap resampling – We use the method to generate a sample for the mean/variance of $\mathcal{W}_1/\mathcal{W}_2$ and $\mathcal{D}$. We first generate the initial samples for $\mathcal{W}_1$, $\mathcal{W}_2$, and $\mathcal{D}$ by performing

¹⁴ Denoting $U(a, b)$ as a uniform distribution with bounds a and b , we draw $N \sim 10^{U(1, 3.5)}$, $\mu_V, \mu_\epsilon \sim U(-10, 10)$ and $\sigma_V^2, \sigma_1^2 \sim U(0.3, 10)^2$.

¹⁵ In other words, we draw $M \sim N \times U(0.01, 0.8)$, subjected to rounding with the floor function and a minimum value of one. Multipliers above 0.8 are ignored, as it is unrealistic for organisations to have such a high relative capacity.

¹⁶ We draw $\sigma_2^2 \sim (\sigma_1 \times U(0.1, 0.99))^2$, subjected to a minimum value of 0.2². Multipliers below 0.1 and above 0.99 are ignored, as we consider such an estimation noise reduction unrealistic and of negligible value, respectively.

Table 2.1: The number of statistical tests with the centred 95% confidence interval containing the derived theoretical value for each quantity of interest. If a theoretical value derived in Sections 2.4 or 2.5 is exact, its associated 95% CI should contain the theoretical value in 95% of the statistical tests.

Quantity	# within CI	# statistical tests	% within CI
$\mathbb{E}(\mathcal{W})$	3,991	4,428	90.13%
$\mathbb{E}(\mathcal{D})$	4,162	4,428	93.99%
$\text{Var}(\mathcal{W})$	3,390	4,555	74.42%
$\text{Cov}(\mathcal{H}_{(r)}, \mathcal{L}_{(s)})$	4,663	4,940	94.39%
$\text{Cov}(\mathcal{V}_{\mathcal{I}(r)}, \mathcal{V}_{\mathcal{J}(s)})$	4,730	4,940	95.75%

10,000 *simulation runs* (see below). We then resample the initial samples and calculate the mean/variance of the resample to obtain an empirical mean/variance sample. Finally, we repeat the resampling 2,000 times to obtain a representative empirical distribution for the mean/variance.

b) Sampling for order statistics – The bootstrapping approach is unlikely to work on the covariance between the order statistics (e.g., $\mathcal{H}_{(r)}$ and $\mathcal{L}_{(s)}$) and their concomitants (e.g., $\mathcal{V}_{\mathcal{I}(r)}$ and $\mathcal{V}_{\mathcal{J}(s)}$), as the ranking information may not be preserved during resampling. Hence, for these quantities, we opt for a more naïve sampling approach. We generate 200 samples for $\mathcal{H}_{(r)}$, $\mathcal{L}_{(s)}$, $\mathcal{V}_{\mathcal{I}(r)}$, and $\mathcal{V}_{\mathcal{J}(s)}$ via the same number of *simulation runs*, and calculate the covariance between these quantities to obtain an empirical sample. The process is repeated 1,000 times to yield a representative empirical distribution for the covariance.

Finally, in each *simulation run*, we obtain one sample for each of $\mathcal{H}_{(r)}/\mathcal{L}_{(s)}$, $\mathcal{V}_{\mathcal{I}(r)}/\mathcal{V}_{\mathcal{J}(s)}$, $\mathcal{W}_1/\mathcal{W}_2$, and D w.r.t. the parameters via the following four-step process:¹⁷

1. Take N samples from $\mathcal{N}(\mu_{\mathcal{V}}, \sigma_{\mathcal{V}}^2)$ as V_i ;
2. Take N samples from $\mathcal{N}(\mu_{\epsilon}, \sigma_1^2)$, and sum the i^{th} -indexed sample with $V_i \forall i$ to obtain H_i :
 - (a) Rank all H_i , take the value of the r^{th} -ranked H_i as $H_{(r)}$ and its index as $I(r)$,
 - (b) Take the value of the $I(r)^{\text{th}}$ -indexed sample of V_i as $V_{I(r)}$,

¹⁷ Identifiers in monospace refer to variables used in software packages, which correspond to the random variables used in Sections 2.3–2.5.




Figure 2.6: Histogram of the theoretical quantity's percentile rank compared to the empirical samples across multiple statistical tests. If the theoretical value derived in Section 2.4 is exact, the histogram should show a uniform distribution with probability mass around the black line [259].

- (c) Obtain the indices of the M largest samples of the ranked H_i ,
 - (d) Calculate the mean of V_i , where i is in the set of indices in Step 2c, as W_1 ;
3. Take N samples from $\mathcal{N}(\mu_\epsilon, \sigma_\epsilon^2)$, and sum the i^{th} -indexed sample with $V_i \forall i$ to obtain L_i :
- (a) Rank all L_i , take the value of the s^{th} -ranked L_i as $L_{(s)}$ and its index as $J(s)$,
 - (b) Take the value of the $J(s)^{\text{th}}$ -indexed sample of V_i as $V_{J(s)}$,
 - (c) Obtain the indices of the M largest samples of the ranked L_i ,
 - (d) Calculate the mean of V_i , where i is in the set of indices in Step 3c, as W_2 ;
4. Take the difference between W_2 obtained in Step 3d and W_1 from Step 2d to get D .

We show the results in Table 2.1 and Figure 2.6. We observed that the 95% CI of the quantities $\mathbb{E}(\mathcal{W})$, $\mathbb{E}(\mathcal{D})$, $\text{Var}(\mathcal{W})$, $\text{Cov}(\mathcal{H}_{(r)}, \mathcal{L}_{(s)})$, and $\text{Cov}(\mathcal{V}_{\mathcal{I}(r)}, \mathcal{V}_{\mathcal{J}(s)})$ contain the derived theoretical value for roughly 90%, 94%, 74%, 94%, and 96% of the times respectively. While these numbers are expected for the expectations and covariances, considering they are approximations, they are on the low side for the variances. Upon further investigation, we realised that most out-of-CI cases have a theoretical variance below the CI, suggesting a slight underestimate in our variance derivation. We believe that this is due to the omission of higher-order terms when using the formulas in [54], leading to a small (less than one percent relative) bias,

which is more apparent for smaller N and M . Otherwise, we are satisfied with the soundness of the derived quantities.

2.7 Case Study

“What do e-commerce/marketing companies gain by acquiring digital experimentation and measurement capabilities?”

It is difficult to verify any model that seeks to ascertain the value of DEM capabilities with real-life data. This is because of the inability to observe the true value of a digital product / business proposition / service without any measurement error and the lack of published measurements from organisations. The closest proxies are meta-analyses, including those compiled by Browne and Johnson [30] and Johnson et al. [136], which contain statistics on the measured uplift (in relative %) over a large number of e-commerce and marketing experiments for many organisations.

The information presented by the two groups of researchers, which we describe in more detail below, is sufficient for us to ask the following question: If the same organisation conducted all the experiments presented by Browne and Johnson / Johnson et al., how much value did the DEM capabilities add due to improved prioritisation? We present results under normal assumptions in this section and will revisit the question when we discuss the model under t -distributed assumptions in Section 2.8.1.

2.7.1 e-Commerce companies

In [30], Browne and Johnson reported running 6,700 A/B tests in e-commerce companies, with an overall effect in relative conversion rate (CVR) uplift centred at around zero, and the 5% and 95% percentiles at around $\pm[1.2\%, 1.3\%]$. We then divide the range by $z_{0.95} \approx 1.645$, the 95th percentile of a standard normal, to estimate that the distribution reported has a standard deviation of around 0.75%. Based on this information, we take $\mu_V = 0$ and $\sigma_V^2 = (0.6\%)^2$, considering that the reported distribution incorporated some estimation noise, and hence, the spread of the true values should be slightly lower.

Given an A/B test on CVR uplift run by the most prolific organisations (e.g. one with five million visitors and a 5% CVR) carries an estimation noise of around $(0.28\%)^2$,⁸ we explore the scenarios where we reduce the noise level from $\sigma_1^2 = \{(1\%)^2, (0.8\%)^2, (0.6\%)^2\}$ to $\sigma_2^2 = \{(0.8\%)^2, (0.6\%)^2, (0.4\%)^2\}$, representing different levels of estimation abilities before and after the acquisition of DEM capabilities for companies of various sizes. We also calculate the value gained under different M (from 10 to 2,000) to simulate organisations with different yet realistic capacities while fixing $N = 6700$ (# experiments). Finally, we set $\mu_\epsilon = 0$ as we do not assume any systematic bias during estimation in this case.

We show the results in Figure 2.7, which shows the relationship between different M and the value gained under different magnitudes of estimation noise reduction. One can observe that the expected gain in value decreases when M increases. This is expected: as one increases their capacity, they will run out of the most valuable work and have to settle for less valuable work that has many acceptable replacements with similar value, limiting the value DEM capabilities bring.

We can also see an inverse relationship between the size of M and the uncertainty of the value gained. As a result, while the expected value gain decreases with increasing M , the uncertainty drops quicker, such that at some M , we will see a statistically significant increase in value gained or an acceptable Sharpe ratio that justifies investment in DEM capabilities or both. The specific value that tips the balance is heavily dependent on individual circumstances.

2.7.2 Marketing companies

In the second case study, we repeat the process applied to e-commerce in Section 2.7.1 for the marketing experiments described in [136]. In that work, Johnson et al. reported running 184 marketing experiments that measured relative CVR uplift, with a mean relative uplift of 19.9% and standard error of 10.8%. Thus, we take $\mu_V = 19.9\%$ and $\sigma_V^2 = (10\%)^2$, which is slightly reduced to account for the estimation noise being included in the reported standard error.

Johnson et al. also noted that the average sample size in these experiments is over five million, which keeps the estimation noise low. However, the design of marketing experiments




Figure 2.7: The value gained by having some digital experimentation and measurement (DEM) capabilities $\mathcal{D}$ (x-axes, in per cent) under different capacity M (y-axes, in log scale) in the case study on 6,700 e-commerce experiments reported by Browne and Johnson [30] (see Section 2.7.1). In each plot, the dot represents the mean, and the error bar represents the 5th–95th percentile of the empirical value distribution. Each sub-caption denotes the estimation noise before and after the acquisition of DEM capabilities (i.e., σ_1^2, σ_2^2). We fix $\mu_{\mathcal{V}}, \mu_{\epsilon} = 0$, $\sigma_{\mathcal{V}}^2 = (0.6\%)^2$, and $N = 6700$.

often comes with additional sources of noise compared to standard A/B tests [105, 172]. Hence, we assume the same estimation noise as in the e-commerce case study above (i.e., $\sigma_2^2 = \{(0.8\%)^2, (0.6\%)^2, (0.4\%)^2\}$). The larger variance in the uplifts allows us to assume a larger estimation error without DEM capabilities, and we explore the scenarios where $\sigma_1^2 = \{(5\%)^2, (2\%)^2, (1\%)^2, (0.8\%)^2, (0.6\%)^2\}$. We set $N = 184$ (# experiments) and vary M between 10 and 100 for each combination of σ_1^2 and σ_2^2 .

Figure 2.8 shows the results. In the presence of a larger variability in the true uplift of the advertising campaigns (σ_V^2) and lower capacity (M), the level of estimation noise reduction that gave a statistically significant value gained in the e-commerce example is no longer sufficient. Therefore, one needs a larger noise reduction or to increase their capacity to effectively control the risk in investing in DEM capabilities. They may also be better off increasing their limited number of existing business propositions.

2.8 Empirical Extensions

We also provide two extensions, evaluated empirically, that open the door for future work in this area.

2.8.1 Valuation under independent t -distributed assumptions

So far, we have spent much of the work assuming that the business propositions' true value and the estimation noise are normally distributed. While possessing decent mathematical properties, more is needed to explain the heavy tail in the distribution of uplifts shown in [30] or [136].

In this section, we model the true value of the business propositions and the estimation noise as Generalised Student's t -distributions (see Equation (2.7)). It is difficult to derive the exact theoretical quantities under such model assumptions because Student's t -distributions do not have conjugate priors [230]. We instead simulate the empirical distribution of the value gained under different parameter combinations to understand if this model is a better




Figure 2.8: The value gained by having some digital experimentation and measurement (DEM) capabilities $\mathcal{D}$ (x-axes, in per cent) under different capacity M (y-axes) in the case study on 184 marketing experiments reported by Johnson et al. [136] (see Section 2.7.2). In each plot, the dot represents the mean, and the error bar represents the 5th–95th percentile of the empirical value distribution. Each sub-caption denotes the estimation noise before and after the acquisition of DEM capabilities (i.e., σ_1^2, σ_2^2). Here, we fix $\mu_V = 19.9\%$, $\mu_\epsilon = 0$, $\sigma_V^2 = (10\%)^2$, and $N = 184$.

alternative to that under normal assumptions. The sampling procedure is similar to that described in Section 2.6, with steps modified to generate samples using standard t -distributions, which are then scaled and located as specified by Equation (2.7).

We compare the value gain estimates obtained under t -distributed and normal assumptions. For each comparison, we randomly sample values for N , M , μ_V , μ_ϵ , σ_V^2 , σ_1^2 , σ_2^2 , and perform 1,000 simulation runs of the four-step sampling procedure in Section 2.6 to obtain samples of D using both the t_3 and normal distributions.¹⁸ We then compare the expected values, the 5th and 95th percentile of the value gained under the two distributions.

We observed from 840 comparisons that overall, the value gained under the t -distributed assumptions has a higher mean (7% higher mean) and variance (7% higher in the 95% percentile on average) than that under normal assumptions. The result arises despite us setting the mean/variance of the true value and estimation noise under the t -distributed assumptions to that under the normal assumptions. This suggests that the model under t -distributed assumptions can capture the “higher risk, higher reward” concept.

Individual comparisons paint a more nuanced picture, perhaps best illustrated by revisiting the case study in Section 2.7 under t -distributed assumptions. We select a few scenarios featured in the previous section and overlay the value gained by having DEM capabilities under t -distributed assumptions over that under normal assumptions in Figure 2.9. One can see that while t -distributed assumptions generally yield a higher value gained, this is not always the case – for the e-commerce case, as M increases, the value gained decreases quicker under t -distributed assumptions than under normal assumptions. This shows that the valuation of DEM capabilities is sensitive to model assumptions.

2.8.2 Partial estimation/measurement noise reduction

There are many situations when not all business propositions are immediately measurable upon acquiring DEM capabilities. This may be due to the extra work required to integrate additional capabilities in certain legacy systems or the limited ability to run experiments

¹⁸ t_3 (t -distribution with three degrees of freedom (d.f.)) is used as it is the distribution with the longest tail under the t family with a natural number of d.f. while retaining a finite variance.




Figure 2.9: The value gained by having some digital experimentation and measurement capabilities (x-axes, in per cent) under normal assumptions (blue error bars with crosses) and t -distributed assumptions (orange error bars with triangles). Both assumptions are evaluated under different but matching capacity M (y-axes, with the error bars displaced for clarity). See Section 2.7.1/Figure 2.7 and Section 2.7.2/Figure 2.8 for details of the parameters for e-commerce and marketing experiments, respectively.

on online but not offline activities. In the case where there is a single backlog, we ask the question, will an organisation still benefit from a partial noise reduction when some business propositions' values are obtained under reduced uncertainty while others are subject to the original noise level?

We address this by attempting to establish the relationship between the expected improvement in the mean true value of the selected business propositions and the proportion of business propositions that benefited from a reduced estimation noise (denoted $p \in [0, 1]$).¹⁹ The sampling procedure is similar to that described in Section 2.6, with Step 3 modified: instead of generating all samples from $\mathcal{N}(\mu_\epsilon, \sigma_2^2)$, we generate p of the samples from $\mathcal{N}(\mu_\epsilon, \sigma_2^2)$ (the lowered estimation noise) and $1 - p$ of the samples from $\mathcal{N}(\mu_\epsilon, \sigma_1^2)$ (the original estimation noise).

We run the procedure above under various scenarios, including a large/small N , a large/small ratio between an organisation's capacity and backlog (M/N), and a large/small magnitude of noise reduction upon acquisition of DEM capabilities ($\sigma_1^2 - \sigma_2^2$). Figure 2.10 shows the result. We can see that under most scenarios, the expected value gained increases with p linearly, while there are a few scenarios where the expected improvement in mean true value of the selected business propositions curve upward for increasing p . This shows that while organisations are incentivised to acquire DEM capabilities that cover the majority of their work, in many scenarios, a partial acquisition yields proportional benefits. Potential experimenters need not consider the acquisition a zero-one decision or worry about any steep initial investment required to unlock returns.

2.9 A Brief Recap

In this chapter, we have addressed the problem of valuing DEM capabilities, which enables one to justify acquiring the said capabilities. Such capabilities deliver three forms of value to organisations. These are

1. Improved recognition of the value of business propositions,

¹⁹ We can model the estimation noise using a two-component mixture distribution parameterised by p .




Figure 2.10: The near-linear relationship between p (proportion of business propositions with their value estimated under a lower estimation noise, x -axes) and the improvement in the mean true value of the selected business propositions (y -axes) under normal assumptions. In each plot, the dot represents the sample mean, and the error bar represents the 5%–95% percentile of the sample value gained. All figures assume $\sigma_v^2 = 1$, while the top four figures assume $\sigma_1^2 = 0.5^2$ and $\sigma_2^2 = 0.4^2$ (corresponding to a small reduction in estimation noise), and the bottom four figures assume $\sigma_1^2 = 0.8^2$ and $\sigma_2^2 = 0.2^2$ (corresponding to a large reduction in estimation noise).

2. Enhanced capability to prioritise, and
3. The ability to optimise individual business propositions.

Of these, improving prioritisation is the most challenging to address while being the most applicable for organisations seeking to build DEM capabilities from scratch.

We have established a methodology to value better prioritisation through reduced estimation error using the framework of ranking under uncertainty. The key insight is that DEM capabilities reduce the estimation error in the value of individual business propositions, allowing prioritisation to follow the optimal order of projects more closely had the true values of business propositions been observable. In addition, we have provided simple formulas that give the value of prioritising under lower estimation error brought by DEM capabilities and the Sharpe ratio governing investment decisions. Finally, alongside two case studies that illustrate how we can apply the methodology in specific situations, we have provided general guidelines for conditions when such investments are inappropriate.

Chapter 3

Statistical Testing

This chapter contains background material that appeared in multiple publications:

- “*An Evaluation Framework for Personalization Strategy Experiment Designs*”, presented at and awarded Best Student Paper of *AdKDD 2020 Workshop (in conjunction with SIGKDD '20)* [176];
- “*Datasets for Online Controlled Experiments*”, presented at the *35th Conference on Neural Information Processing Systems (NeurIPS 2021)* [173]; and
- “*Measuring e-Commerce Metric Changes in Online Experiments*”, presented at *ACM Web Conference 2023 (WWW '23)* [177].

3.1 Motivation

After establishing the business case of engaging in digital experimentation and measurement, we turn our focus to two necessary ingredients for running digital experiments. They are data, which record what happens during an experiment, and statistical tests, which indicate whether groups in the experiment give different responses beyond randomness.

We first discuss statistical testing in this chapter, which will complement our approach to datasets in Chapter 4. We will also build upon the concepts introduced in this chapter in Chapter 5, when we formally describe an experiment and its underlying causal reasoning,

and in Chapter 6, when we compare different experimental designs for digital experiments. Generally (and loosely) speaking, a statistical test is a statistical inference procedure that informs us whether the data collected is compatible with our hypothesis(es). In the context of digital experimentation, the data collected is often responses from two or more groups exposed to different treatments, and the hypothesis is usually whether or how much the responses from one group are different from another group.

Statistical testing may appear to be a simple and widely understood topic – at the end of the day, it is taught in many upper secondary education curricula around the globe¹ and forms part of the standard toolkit in scientific research. It is not. As we will show in this chapter, its development is riddled with assumptions that do not suit practical needs and its application is permeated with careless interpretations. Multiple generations of academics and practitioners have derided its misuse (and sometimes its use) [26, 47, 166, 200, 277, 297], with the latest round of controversy surrounding “*p*-hacking” necessitating an official statement from the American Statistical Association [276].

Despite its many associated problems, statistical testing remains useful for decision-making if applied correctly. This chapter aims to provide a comprehensive introduction to the area, with a good balance between the theoretical foundation and its practical use in digital experiments. We seek to enable practitioners in digital experimentation, having arrived from diverse backgrounds, to quickly acquire the knowledge required to perform a statistical test while being mindful of the nuances and pitfalls.

¹ A non-exhaustive list of examples includes (in alphabetical order):

- China: (National College Entrance Examination) Mathematics - Extended Topic 3: Probability and Statistics (Item 3(3)2), Elective Group A: Probability and Statistics (Topic 4), and Mathematics Elective Group B: Applied Statistics (Topic 4) [197, pp. 48, 57, 62]
- Germany (Berlin & Brandenburg): (Abitur) Mathematics (Semester 4 / Q4, Advanced-level only, Theme 5 / L5) [196, p. 31]
- Japan: Mathematics B (Topic (2)A(D)) [198, pp. 104, 108–109]
- United Kingdom (AQA): GCE A Level in Mathematics (Subject Content O) [7, p. 23]
- United Kingdom (Edexcel): GCE A Level in Statistics (Topics 7–10, 15–17, and 19–20) [219, pp. 11–13, 15–19]
- USA: Advanced Placement Statistics (Units 6–9) [48, pp. 125–205]
- Worldwide - International Baccalaureate Mathematics: applications and interpretation (Topics SL 4.11 & AHL 4.18) [132, pp. 57–58, 62]

Target audience The chapter is written with a mathematically-inclined audience in mind, i.e., those exposed to university-level introductory mathematics. As we place equal importance on the applications, we will not cover formal proofs and skip through many advanced theoretical concepts, relegating them to footnotes with pointers to relevant works. However, the chapter will involve a fair bit of mathematical notations and algebraic manipulations to help us properly navigate the many concepts in statistical testing. For brevity, we also assume readers are already familiar with concepts covered in introductory probability and statistics textbooks, including conditional probability, population vs sample, and probability distributions.² Readers may also find expositions aimed at a broader audience in digital experimentation, e.g., early chapters of [101], Chapter 17 of [155], and [194] plus its sibling online articles, as well as works on various aspects of statistical testing written for researchers and practitioners in other fields [36, 76, 87, 96, 166, 178, 275] helpful.

A note of caution It is worth emphasising that everything in this chapter only helps us establish whether the groups being compared give different responses and, at best, whether the difference in responses is *correlated* to the different treatments provided to the groups. As the old saying goes, “Correlation does not imply causation” [3]. We require further theoretical foundations to properly establish the causal relationship between the treatment(s) and the change in responses. This will be explored in Chapter 5 when we discuss experimental design in digital experimentation.

Chapter organisation The rest of the section is organised as follows. We first define the basic terminologies and outline the considerations that drive the development of many different statistical tests in Section 3.2, followed by some non-examples but important concepts related to statistical testing in Section 3.3. We then discuss tests and quantities commonly used in digital experimentation from Section 3.4 onward, categorised by the considerations outlined at the end of Section 3.2.

² One can refer to, e.g., Chapters 4–7 of [190] (more illustrative), Chapters 1–5 of [36] (more formal), Chapters 1–2 of [275] (more compact), or any relevant massive open online courses (MOOCs) delivered by a reputable higher education institution.

3.2 Pre-matter and Considerations

A statistical test starts with the experimenter stating one or more statistical hypotheses.³

Usually, one specifies a *null hypothesis* H_0 , the preferred/assumed state of matter,⁴ and one or more *alternative hypothesis(es)* $H_1, H_2, \dots$. Together, the hypotheses usually cover all possible scenarios. In digital experiments, we may have two populations, X and Y , and are interested in inferring the difference in their population mean, $\Delta = \mu_Y - \mu_X$, from their respective samples. Depending on the experimenter's claim, the hypotheses can be in the form of:

- In a *comparison* test, a.k.a. a *two-sided* test in some cases –

$$H_0 : \Delta = 0, H_1 : \Delta \neq 0.$$

- In a *superiority* test, a variant of a *one-sided* test [49, 272] –

$$H_0 : \Delta \leq 0, H_1 : \Delta > 0.$$

- In a *non-inferiority* test, another variant of a one-sided test commonly used to show a population is not substantially worse than another –

$$H_0 : \Delta \leq -\theta, H_1 : \Delta > -\theta,$$

where θ is a small margin that an experimenter is prepared to take to declare that μ_Y is not substantially lower than μ_X .

³ Not to be confused with a *scientific/research hypothesis*. A scientific hypothesis describes what the experimenter intends to discover/validate about the relationship between variables in the real world. A statistical hypothesis is a much narrower statement on how a statistical model/procedure (usually a simplified view of the real-world phenomenon) behaves. See [28, 160, 166, 189] for further discussions.

⁴ Not to be confused with a *nil hypothesis*, one that hypothesises the parameter of interest (e.g., Δ as defined in the same paragraph) is zero. Our comparison test example shows a nil hypothesis as the null, though other hypotheses, such as $H_0 : \Delta = -3$, are also valid. See [47, 104] for further discussions.

- In an *equivalence* test (commonly carried out using two one-sided tests, or TOSTs) –

$$\begin{array}{ll} H_0 : \Delta \leq -\theta_L, H_1 : \Delta > -\theta_L & \text{in one and} \\ H_0 : \Delta \geq \theta_U, H_1 : \Delta < \theta_U & \text{in the other.} \end{array}$$

Rejecting both null hypotheses would lead to the conclusion that $-\theta_L < \Delta < \theta_U$, where θ_L and θ_U are the negative and positive margins, respectively, that an experimenter is prepared to take to declare that μ_Y and μ_X are equal.

In practice, within and outside digital experimentation, it is common to treat the null hypothesis as one to be nullified, one set up to enable an experimenter to demonstrate the contrary despite its status as the assumed state of matter [47].⁵ One major exception is A/A tests – randomised controlled trials featuring two supposedly identical treatments. These tests contrast A/B tests that feature two different treatments (see Chapter 1 / Figure 1.1) and are often employed in digital experimentation to ensure one implements the systems and models correctly. Given identical treatments, one would expect results consistent with a nil null hypothesis (no difference in the population means) in a comparison test – the alternative hypothesis is the unexpected artefact.⁶

After specifying the statistical hypotheses, one then has to decide how to carry out the statistical test procedure. Many statistical tests have been proposed to date, each driven by a combination of different considerations:

Philosophical alignment The philosophical debates on the interpretation of probability and the role of models in statistical inference have led to different approaches to running a statistical test. The *frequentist* versus *Bayesian* debate dominates the former, and the latter manifests in whether a statistical test should solely inform one's degree of belief in a particular hypothesis or act as a decision rule between competing hypotheses [162, 165]. To

⁵ One can draw a parallel between the reasoning behind nullifying the null hypothesis in a statistical test and that behind proof by contradiction in pure mathematics. That said, the introduction of probabilistic statements has complicated matters. See [47], particularly the “The Permanent Illusion” section, for further discussions.

⁶ Instead of merely assuming the nil null hypothesis, some may instead put the hypothesis to test by running an equivalence test with $-\theta_L$ and θ_U very close to zero. However, we are back to the “nullifying the null hypothesis” regime in that case.

make things more confusing, statistical testing in modern times, known as *null hypothesis significance testing* (NHST), is a hybrid of the two procedures (and underlying schools of thought) [166]. We will introduce examples of NHST and Bayesian tests common in digital experiments in Sections 3.4/3.5 and 3.9, respectively.

Model/distributional assumptions As an inference procedure, statistical tests almost always utilise a statistical model – a set of assumptions related to the sample data generation process. The model can be *parametric* (i.e., assuming distributions with a finite number of parameters), *non-parametric* (i.e., assuming distributions with an infinite number of parameters or even no distributions), or a mixture of the two [102]. All tests that we will introduce in Section 3.4 are parametric. We will also discuss non-parametric tests in Section 3.7.

Number of samples A test can involve *one, two, or many* sample(s) (groups of responses in business speak). Some tests also utilise models that enable one to collapse a two-sample test into a one-sample test (which, one can argue, reduces the operational complexity). Examples of the latter include the *z*-test and *t*-test, which we will discuss in depth in Section 3.4.

Decision-making process Some tests are developed with a *fixed-sample/fixed-horizon* experiment in mind. Other tests, which include those that arise from *sequential* analysis and *Bayesian* statistics, support adaptive stopping (see Section 4.1 for a further exposition). It is important to separate this consideration from that on philosophical alignment, e.g., there is both frequentist and Bayesian interpretation to sequential tests. We will introduce only fixed-horizon tests in Section 3.4, followed by sequential tests in Section 3.8 and Bayesian tests in Section 3.9.

3.3 Alternative Approaches to Formal Tests

Before we proceed with the definitions of relevant statistical tests, we discuss some related concepts that are not formal statistical tests yet are often mentioned alongside statistical testing. They are the inter-ocular trauma test, effect size (of a treatment), and confidence intervals.

3.3.1 The inter-ocular trauma test

Given the number of considerations listed above, it is perhaps unsurprising that many works that collate and organise statistical tests already exist [184]. Indeed, many guides for experimenters from different fields in need of statistical tests feature a flowchart of various sophistication – see [224], [71, 107], and [58] for one intended for general use, clinical trials, and digital experimentation, respectively. However, many of these works may have neglected to ask an important question:

“Do we require a statistical test at all?”

Many statistical tests are unnecessary from the outset. Some results are obvious and way beyond any possible random fluctuations. Other results come from such a noisy environment that it is statistically unfeasible to extract any signal. The inter-ocular trauma (IOT) test deals with the former:

“You know what the data mean when the conclusion hits you between the eyes.” [84]

An IOT test is not a statistical test per se – it is often performed visually via plots [31, 204] or tables without involving formal statistical inference. Nonetheless, it reminds experimenters that statistical testing is not just a rite of passage or a recipe to follow to obtain a “significant result”, whatever “significant” might be.

We are not asserting that all statistical tests are unnecessary. Determining what constitutes inter-ocular trauma is more an art than science and requires the backing of substantial domain expertise. Indeed, Edwards et al. [84] added the following immediately after publicly coining the IOT test,

“But the enthusiast’s interocular trauma may be the skeptic’s random error. A little arithmetic to verify the extent of the trauma can yield great peace of mind for little cost.” [84]

3.3.2 Effect size

We move on to effect size, which we touched upon without explicitly naming when we set up our running example in Section 3.2. Recall that we have two populations, X and Y , and

are interested in the difference in their means

$$\Delta = \mu_Y - \mu_X . \quad (3.1)$$

When X and Y represent responses from two groups receiving different treatments in a randomised controlled trial, Δ is known as the *effect size* (of the mean difference) or the *average treatment effect*. It enables us to tell whether a treatment is creating an impact that matters in practice (i.e., *practically significant*). Effect size is just a measure,⁷ not a statistical test (as no inference is made). It is, though, a measure that many statistical tests are set up to infer and is one of the two quantities of interest when we compare digital experiment designs in Chapter 6.

A standardised effect size enables us to compare the difference across many (related) experiments, thus useful in meta-analyses. One commonly used effect size is Cohen's d , defined as the difference in sample means divided by the pooled sample standard deviation [46]. More formally, let $X_1, \dots, X_n$ and $Y_1, \dots, Y_m$ be independent and identically distributed (i.i.d.) samples from the two populations in our running example, with sample means and variances $(\bar{X}, \bar{Y})$ and (s_X^2, s_Y^2) , respectively. Cohen's d is defined as

$$d = \frac{\bar{Y} - \bar{X}}{\sqrt{\frac{(n-1)s_X^2 + (m-1)s_Y^2}{n+m-2}}} . \quad (3.2)$$

Other commonly used effect sizes, including Glass' Δ and Hedge's g [120], also involve dividing the difference in sample means by a standalone or pooled standard deviation, with the latter also involving a correction factor dependent on the sample sizes n and m .

3.3.3 Confidence Intervals

The final non-example is confidence intervals. Confidence intervals are range estimates for a parameter of interest θ , which is unknown but fixed. Examples of such parameters include the population mean of a sample and the effect size introduced above.

We define confidence intervals as follows. Consider a sample X generated from a statistical

⁷ In a layman's sense, i.e., "a basis or standard of comparison." [191]

model with a parameter of interest θ . A random interval constructed from X is an interval delimited by two transformations of X , themselves r.v.'s: $(u(X), v(X))$. It is said to be a $(1 - \alpha)$ confidence interval of θ if the random interval covers θ in at least $(1 - \alpha)$ of instances, i.e.,

$$\mathbb{P}(u(X) < \theta < v(X)) = 1 - \alpha, \quad (3.3)$$

where $(1 - \alpha)$ is known as the *confidence level* and is usually chosen to be a value close to but less than one (e.g., 95%). The probability statement should hold for all possible values of θ , as well as that for all other parameters in the model. This is achieved by choosing appropriate $u(\cdot)$ and $v(\cdot)$.

For example, consider an i.i.d. sample $X_1, \dots, X_n$ generated from a normal distribution with mean μ_X and variance σ_X^2 , estimated by the sample mean $\bar{X}$ and sample variance s_X^2 , respectively. A commonly used $(1 - \alpha)$ confidence interval of μ_X is

$$\left(\bar{X} - t_{n-1, 1-\alpha/2} \sqrt{\frac{s_X^2}{n}}, \bar{X} + t_{n-1, 1-\alpha/2} \sqrt{\frac{s_X^2}{n}} \right), \quad (3.4)$$

where $t_{n-1, 1-\alpha/2}$ is the $(1 - \alpha/2)$ quantile of the Student's t -distribution with $n - 1$ degrees of freedom.⁸ Upon observing the empirical responses, they are substituted into the formula in Expression (3.4) to obtain the realised confidence interval specific to the experiment.

Confidence intervals are not statistical tests as they do not involve any statistical hypotheses. That said, a “duality” exists between a $(1 - \alpha)$ confidence interval and the non-rejection region of a statistical test that utilises the same statistical model at significance level α , which leads to confidence intervals frequently appearing in test results. The “duality” enables one to construct a confidence interval for θ by taking values that a corresponding statistical test on its point estimate (i.e., $H_0 : \theta = \hat{\theta}$) fails to reject [52]. It also enables one to reject the null hypothesis $H_0 : \theta = \theta_0$ in a statistical test if the corresponding confidence interval does not cover θ_0 .

Some consider confidence intervals superior to a statistical test or p -values (see Section 3.6.1). They see confidence intervals rebalancing the influence from the null and alternate hypothe-

⁸ For brevity, we will skip how such an interval is constructed and why such construction is valid. For a more extensive exposition, readers can refer to, e.g., Sections 7.2 & 7.3 of [52].

ses – instead of giving the null a preferred status as one does in a (frequentist) test, confidence intervals make one look at the full range of effect sizes compatible with the data [109]. Others consider confidence intervals a more effective heuristic on the interaction between effect size and sampling uncertainty than any two numbers from a statistical test.

On the other hand, confidence intervals suffer from the same problem as p -values – there is no guarantee that they are always useful in practice. It is possible for confidence interval procedures, even the “good” ones, to produce intervals that are overly wide or narrow (or even empty) while maintaining a $(1 - \alpha)$ confidence level overall [199]. Moreover, its root in frequentist, parametric statistics also prevents meaningful comparison with non-parametric or Bayesian tests on a like-for-like basis. Any attempt at such will likely revert to debates on the underlying philosophy.

In addition, many interpretation pitfalls prevent one from fully reaping the benefits outlined above. Some stem from losing sight of θ , the parameter of interest, during the estimation procedure. Recall that a confidence interval is a random interval that estimates θ , a fixed quantity. It does not infer the samples themselves and thus does not provide any information on the proportion of the samples it may cover.⁹ Readers may have also encountered texts claiming “there is a $(1 - \alpha)$ chance for θ to fall within the $(1 - \alpha)$ confidence interval”. This is misleading. One may consider the phrases “the interval covering θ ” and “ θ falling within the interval” to describe the same event from a different perspective. However, things get complicated once they form part of a probability statement. The description above, which has θ as the subject, incorrectly implies θ is random.

Other misconceptions arise from confusing the procedure, which produces random intervals from random samples, with a particular interval constructed from realised response values. The confidence level is a probability statement on the estimation procedure itself – its value gives the proportion of confidence intervals that cover θ in the long run.¹⁰ It does not matter what value θ takes or what the parameter is – all $(1 - \alpha)$ confidence intervals that were and

⁹ A random interval that covers $(1 - \alpha)$ of the samples is known as a $(1 - \alpha)$ *prediction interval*. Generally, these two intervals are not comparable as they estimate different quantities. When one produces confidence intervals for the population mean of a normal distribution, a prediction interval usually shares the same mid-point with, yet is much wider than, a confidence interval with the same “rating”.

¹⁰ The notion of long-run proportion underpins the frequentist interpretation of probability.

will be constructed should cover the corresponding θ in $(1-\alpha)$ of the instances.¹¹ On the other hand, once an interval is realised from the response values, it is no longer random. While θ remains unknown, this particular interval either covers the fixed θ or does not – there is no probability statement to be made. Thus, the notion of confidence level is meaningless for realised intervals [207].

3.4 Null Hypothesis Significance Tests

We continue to discuss statistical tests by introducing null hypothesis significance tests (NHSTs), the most popular class of statistical tests featured in classrooms and most application areas, including digital experimentation. In this section, we describe the general procedure for running an NHST. Three examples follow in Section 3.5, and a practical guide to concepts and quantities featured in NHSTs in Section 3.6.

The general procedure for running an NHST is as follows:

1. Determine the statistical hypotheses, namely the null and alternative(s), as outlined in Section 3.2.
2. Determine an appropriate test based on underlying considerations (see Section 3.2) and state the associated *test statistic* T – see below for typical examples in digital experimentation.
3. Determine the appropriate *significance level* (α), and as best practice, minimum *power* ($\pi_{\min}$) and *sample size* (n) via an *a priori* power analysis – see Section 3.6 for the definition of the mentioned concepts and description of a power analysis.
4. Derive/obtain the distribution of T under the null hypothesis. We call a hypothesis *simple* if it completely determines the distribution of T and *composite* otherwise.
5. Calculate the test statistic based on the observed data, t_{obs} .
6. Determine the *critical region* $\mathbb{T}_{\text{crit}}$ – values of T that lead to us rejecting the null hypothesis. The critical region should be chosen so that the probability of T falling within the region is

¹¹ This assumes all intervals involved are correctly constructed and all other model assumptions are met.

at most α , as determined in Step 2 (see Section 3.6). For named tests, both the distribution and the optimal critical region¹² are standard results.

7. Reject the null hypothesis if the test statistic falls within the critical region, i.e., $t_{\text{obs}} \in \mathbb{T}_{\text{crit}}$. This indicates it is implausible for the null to generate the observed data, and one of the alternate hypotheses is now favoured.¹³ Otherwise, do not reject the null hypothesis.¹⁴

Another approach to the decision steps (Steps 6 and 7) is as follows:

- 6A. Calculate the *p-value*, the highest probability that, given the null hypothesis, a sampled value from T (itself an r.v.) will be at least or more extreme than observed test statistic t_{obs} (see Section 3.6).
- 7A. Reject the null hypothesis if the *p-value* is less than the significance level, i.e. $p < \alpha$. Otherwise, do not reject the null hypothesis.

The existence of the two approaches highlights the difference in philosophical alignment discussed in Section 3.2. The critical region was developed to provide an objective decision criterion. In contrast, the *p-value* was developed to measure the strength of evidence against the null hypothesis without making a formal decision. The latter approach is a result of mashing the two mathematical concepts together, perhaps driven by the desire to “have your cake and eat it”, i.e., to have a measure of the strength of evidence while making a formal decision, and further aided by the happy coincidence that there is a one-to-one mapping between these two concepts. In practice, both approaches are commonly used. Some experimenters also engage in both approaches when running the same NHST.

¹² For one-sided parametric tests, this is given by the uniformly most powerful (UMP) test under the same statistical assumptions. See Section 8.3.2 of [36] or Chapter 3 of [164] for further discussions.

¹³ It might sound innocuous to say that rejecting the null hypothesis indicates the null hypothesis is implausible given the observed data (or, in mathematical terms, a low $\mathbb{P}(H_0|\text{Data})$). This is not true – it indicates the observed data is implausible given the null hypothesis (or a low $\mathbb{P}(\text{Data}|H_0)$). We will explore this subtle difference in greater detail in Section 3.6.1.

¹⁴ Strictly speaking, there is no such notion of accepting the null hypothesis – failing to reject the null simply means it remains the preferred hypothesis as determined at the beginning of a statistical test.

3.5 Examples of NHSTs

We now focus on individual examples, i.e., some named tests commonly used in digital experimentation. These include the z -test, the (Welch's) t -test, and the binomial test of proportion. The examples are all parametric, fixed-horizon NHSTs.¹⁵ Needless to say, there are many other examples of parametric, fixed-horizon NHSTs (e.g., F -test on population variances, one-way analysis of variance test on population means, Pearson's correlation coefficient test on the correlation between two sets of data) that we have to omit for brevity.

We will follow a similar pattern when introducing each test. In the order laid out by the general procedure in Section 3.4, we first state the statistical assumptions, definition of the associated test statistic T , distribution of T , the optimal critical region $\mathbb{T}_{\text{crit}}$, and p -value. We then provide examples of their use in practice.

3.5.1 z -test

The z -test (or test for normal population mean(s)) is one of the most recognised statistical tests, perhaps because of its early coverage in introductory statistical textbooks.

Assumptions & Definitions We first consider a one-sample test: let $X_1, \dots, X_n$ be i.i.d. samples from a population with unknown mean μ_X and known variance σ_X^2 . We are interested in inferring μ_X , estimated by the sample mean $\bar{X}$. The z -test assumes all other parameters (in this case σ_X^2) are already known and defines the test statistic as

$$Z \triangleq \frac{\bar{X} - \mu_X}{\sqrt{\sigma_X^2/n}}. \quad (3.5)$$

Test statistic distribution The test also assumes the test statistic, also known as the standard score, follows the standard normal distribution, denoted $\mathcal{N}(0, 1)$. Theoretically, having normally distributed data ensures that the test statistic is normally distributed. This can eas-

¹⁵ Some regard these tests as *classical* NHSTs, though there is no formal definition or consensus on what constitutes a classical NHST.

ily be shown by observing

$$X_1, \dots, X_n \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(\mu_X, \sigma_X^2) \implies \bar{X} \sim \mathcal{N}(\mu_X, \sigma_X^2/n) \implies Z \sim \mathcal{N}(0, 1). \quad (3.6)$$

It leads to some textbooks (incorrectly) mandating normality assumptions in the *data* for a z -test. In practice, one would be hard-pressed to find data that is strictly normally distributed. Instead, many appeal to the central limit theorem (CLT), which implies that the sample mean of any underlying distribution approximately follows the normal distribution given sufficient samples,¹⁶ i.e., $\bar{X} \stackrel{\text{approx.}}{\sim} \mathcal{N}(\mu_X, \sigma_X^2/n)$. This enables one to run an approximate z -test using the same test statistic.

The implication above also means we can run an approximate z -test on two or more samples to infer a linear combination of the means, e.g., the effect size between two groups. We further let $Y_1, \dots, Y_m$ be i.i.d. samples from a population with unknown mean μ_Y (estimated by the sample mean $\bar{Y}$) and known variance σ_Y^2 . Using the same implication above, we can assume $\bar{Y} \stackrel{\text{approx.}}{\sim} \mathcal{N}(\mu_Y, \sigma_Y^2/m)$ and thus

$$\bar{Y} - \bar{X} \stackrel{\text{approx.}}{\sim} \mathcal{N}\left(\mu_Y - \mu_X, \frac{\sigma_X^2}{n} + \frac{\sigma_Y^2}{m}\right). \quad (3.7)$$

Noting the effect size as $\Delta = \mu_Y - \mu_X$ (see Equation (3.1)), the test statistic in this two-sample setting is

$$Z \triangleq \frac{(\bar{Y} - \bar{X}) - \Delta}{\sqrt{\frac{\sigma_X^2}{n} + \frac{\sigma_Y^2}{m}}} \stackrel{\text{approx.}}{\sim} \mathcal{N}(0, 1). \quad (3.8)$$

One may notice that Equations (3.5) and (3.8) place the parameters of interest (μ_X and Δ) in the definition of Z rather than the distribution.¹⁷ Such formulation enables one to use the same distribution for different statistical hypotheses on μ_X or Δ – parameters' values only affect the value of the test statistic calculated using the observed data. This is appealing for old clinical trials and pedagogical purposes as one only needs to refer to a single z -table.

¹⁶ Formally, the Lindeberg-Lévy CLT states that $\sqrt{n}(\bar{X} - \mu_X)$ converges in distribution to $\mathcal{N}(0, \sigma_X^2)$. The practical but naïve interpretation of the statement above is that the distribution of $\bar{X}$ is getting closer to a normal distribution with mean μ_X and variance σ_X^2/n as n increases. It makes no claim on what constitutes a sufficient number of samples.

¹⁷ With σ_X^2 and σ_Y^2 known, this makes Z a *pivotal quantity*.

Such appeal is less pronounced in digital experiments with modern computers.

Critical region & p -value While the distribution of a z -test statistic does not depend on the statistical hypotheses, it is not the case for the critical region $\mathbb{T}_{\text{crit}}$ and p -value – the direction of the hypotheses matters. Consider the one-sample setting, where we use Z defined in Equation (3.5) to infer on μ_X :

$$\mathbb{T}_{\text{crit}} = \begin{cases} (z_{1-\alpha}, \infty) & \text{under a greater/right-tailed test } (H_0: \mu_X \leq \theta_0, H_1: \mu_X > \theta_0) \\ (-\infty, z_\alpha) & \text{under a less/left-tailed test } (H_0: \mu_X \geq \theta_0, H_1: \mu_X < \theta_0) \\ (-\infty, z_{\alpha/2}) \cup (z_{1-\alpha/2}, \infty) & \text{under a two-sided test } (H_0: \mu_X = \theta_0, H_1: \mu_X \neq \theta_0), \end{cases} \quad (3.9)$$

$$p = \begin{cases} 1 - \Phi(z_{\text{obs}}) & \text{under a greater/right-tailed test} \\ \Phi(z_{\text{obs}}) & \text{under a less/left-tailed test} \\ 2 \cdot (1 - \Phi(|z_{\text{obs}}|)) & \text{under a two-sided test,} \end{cases} \quad (3.10)$$

where z_{obs} denotes the observed test statistic, z_q is the q^{th} quantile of a standard normal, and $\Phi(\cdot)$ is the CDF of a standard normal. The exposition for a two-sample test is identical, with Δ and the definition of Z in Equation (3.8) replacing μ_X and the definition in Equation (3.5), respectively.

Practical usage Despite its popularity in introductory statistical textbooks, a pure z -test is uncommon in digital experiments as one needs to know/assume the population variances (σ_X^2 and σ_Y^2) in advance. Nonetheless, experimenters have combined elements of the approximate z -test and the Student's/Welch's t -test to test whether the mean response of a group differs from that of another group. We will describe how such a combination works in practice below.

3.5.2 t -tests

We now introduce t -tests, the most popular class of statistical tests in digital experimentation [155]. It is a collection of closely-related variants – in addition to the one- or two-sample

distinction (similar to the z -test discussed above), there are also slightly different t -tests for independent or paired samples [129], equal or unequal response variance [106, 279], and full or partial range of data points [295]. In the rest of the section, we focus on Welch's t -test [279], which deals with independent samples with potentially unequal variance across groups and takes all available data.

Assumptions & Definitions Let $X_1, \dots, X_n$ and $Y_1, \dots, Y_m$ be i.i.d. samples from two populations with unknown means μ_X and μ_Y , respectively, and *unknown* variances σ_X^2 and σ_Y^2 , respectively. The population means and variances are estimated by the sample means $(\bar{X}, \bar{Y})$ and variances (s_X^2, s_Y^2) , respectively. Similar to the z -test, we are interested in inferring μ_X and $\Delta = \mu_Y - \mu_X$ in a one-sample and a two-sample t -test, respectively. However, unlike the z -test, which only assumes the test statistic is normally distributed, a t -test assumes the *data* (more precisely, the underlying r.v.'s generating the data) are normally distributed.¹⁸

The test statistics are

$$T \triangleq \frac{\bar{X} - \mu_X}{\sqrt{s_X^2/n}} \quad \text{for a one-sample test on } \mu_X \text{ and} \quad (3.11)$$

$$T \triangleq \frac{\bar{Y} - \bar{X} - \Delta}{\sqrt{\frac{s_X^2}{n} + \frac{s_Y^2}{m}}} \quad \text{for a two-sample test on } \Delta. \quad (3.12)$$

They are similar to that in the z -test (cf. Equations (3.5) and (3.8)), with the only difference being that the t -test statistics use the sample variances (s_X^2 and s_Y^2) instead of the population variance. The difference is critical – as the sample variances estimate the population variances (and hence can be formulated as estimators/r.v.'s), they introduce another source of uncertainty to the parameters we are inferring. Other variants of t -tests may pool the variance, resulting in slightly different denominators in the fraction within the square root.

¹⁸ Strictly speaking, neither the original t -test from Gosset/Student nor Welch's t -test explicitly states the r.v.'s generating the data has to be normally distributed. Instead, they assume (1) $\bar{X}$ (or its two-sample counterpart) follows the normal distribution, (2) $s^2(n-1)/\sigma^2$ follows the χ^2 distribution, and (3) the quantities in (1) and (2) are independent. However, the assumptions only hold when the underlying r.v.'s are normally distributed.

Test statistic distribution Both test statistics follow the Student's t -distribution with degrees of freedom ν determined by the Welch-Satterthwaite equation [238]:

$$T \sim t_\nu, \text{ where } \nu = \frac{\left(\frac{s_X^2}{n} + \frac{s_Y^2}{m}\right)^2}{\frac{(s_X^2/n)^2}{n-1} + \frac{(s_Y^2/m)^2}{m-1}} \text{ for a two-sample test.} \quad (3.13)$$

In cases where $s_X^2 \approx s_Y^2$ and $n \approx m$, the degrees of freedom are around $n + m - 2$. In a one-sample test on μ_X , the degrees of freedom are simply $n - 1$.

Critical region & p -value Similar to the z -test, the distribution of a t -test statistic does not depend on the statistical hypotheses. In contrast, the critical region and p -value do. For a one-sample test inferring on μ_X , they are

$$\mathbb{T}_{\text{crit}} = \begin{cases} (t_{\nu, 1-\alpha}, \infty) & \text{under a greater/right-tailed test } (H_0: \mu_X \leq \theta_0, H_1: \mu_X > \theta_0) \\ (-\infty, t_{\nu, \alpha}) & \text{under a less/left-tailed test } (H_0: \mu_X \geq \theta_0, H_1: \mu_X < \theta_0) \\ (-\infty, t_{\nu, \alpha/2}) \cup (t_{\nu, 1-\alpha/2}, \infty) & \text{under a two-sided test } (H_0: \mu_X = \theta_0, H_1: \mu_X \neq \theta_0), \end{cases} \quad (3.14)$$

$$p = \begin{cases} 1 - T_\nu(t_{\text{obs}}) & \text{under a greater/right-tailed test} \\ T_\nu(t_{\text{obs}}) & \text{under a less/left-tailed test} \\ 2 \cdot (1 - T_\nu(|t_{\text{obs}}|)) & \text{under a two-sided test,} \end{cases} \quad (3.15)$$

where t_{obs} is the observed test statistic, $t_{\nu, q}$ is the q^{th} quantile of a Student's t -distribution with ν degrees of freedom, and $T_\nu(\cdot)$ is the CDF of a Student's t -distribution with ν degrees of freedom.

Practical usage As mentioned at the beginning of the section, t -tests are popular in digital experiments. However, what digital experimenters call a t -test is usually not the pure Student's/Welch's t -test defined above. The Student's/Welch's t -test was initially developed for when the number of available samples is small. The test also impose normality assumptions on the data. Digital experimentation, on the other hand, often features a large number of responses that are unlikely to be well modelled by normal distributions.

Readers may recall that an approximate z -test, introduced in Section 3.5.1, is suited to experiments with many arbitrarily distributed responses. However, it requires the knowledge of the population variances (σ_X^2, σ_Y^2) , which we will never have in practice.¹⁹ To obtain the t -test digital experimenters use in practice, or the “practical t -test” in short, we need to extend the approximate z -test to take sample variances (s_X^2, s_Y^2) while keeping its other features. Fortunately, given enough samples, this is as simple as using the sample variance estimates as the population variances, or in other words, “plugging in” the sample variances.²⁰ This results in

$$T \stackrel{\text{approx.}}{\sim} \mathcal{N}(0, 1), \quad (3.16)$$

which enables experimenters to work with the normal distributions, instead of the more complicated Student’s t -distributions, during experiment design and analysis (see Section 3.6).

By delving into the underlying statistical theory, we can understand how the shift from Welch’s t -test to the practical t -test (i.e., an approximate z -test with plug-in variance estimates) happens. Readers interested only in practical details can skip to the next titled paragraph. The shift can happen in two ways. The first maintains the normality assumptions imposed by Welch’s t -test. Thus, the test statistic follows the Student’s t -distribution with ν degrees of freedom. As the sample size increases, so would ν , resulting in the test statistic distribution converging to a standard normal distribution:

$$T |_{X,Y \sim \mathcal{N}(\cdot, \cdot)} \sim t_\nu \rightarrow \mathcal{N}(0, 1) \text{ as } \nu \rightarrow \infty. \quad (3.17)$$

The second way removes the normality assumptions altogether. Using the same one-sample test statistic as Welch’s t -test (see Equation (3.11)), we observe that the numerator, including $\sqrt{n}$, converges in distribution to a normal distribution with a zero mean by the central

¹⁹ This is perhaps one of the reasons why the name “ t -test” sticks in the digital experimentation community: It serves as a reminder that we are dealing with the sample variances instead of the population variances.

²⁰ This is known as the *plug-in* principle. Why the plug-in principle works is a topic in its own right – see, e.g., Chapters 4 & 5 of [86] for further discussions.

limit theorem:

$$\sqrt{n}(\bar{X} - \mu_X) \xrightarrow{D} \mathcal{N}(0, \sigma_X^2). \quad (3.18)$$

At the same time, the denominator converges in probability to the population standard deviation by the law of large numbers and the continuous mapping theorem:

$$s_X^2 \xrightarrow{\mathbb{P}} \sigma_X^2 \implies \sqrt{s_X^2} \xrightarrow{\mathbb{P}} \sigma_X. \quad (3.19)$$

Thus, by Slutsky's theorem, the test statistic converges in distribution to the quotient of the respective limits – in somewhat sloppy mathematical notation:

$$T = \frac{\sqrt{n}(\bar{X} - \mu_X)}{\sqrt{s_X^2}} \xrightarrow{D} \frac{\mathcal{N}(0, \sigma_X^2)}{\sigma_X} \equiv \mathcal{N}(0, 1). \quad (3.20)$$

This formulation is a special case (square root version of the one-parameter test statistic) of the Wald test [275]. The exposition for the two-sample case follows in an identical fashion.

What constitutes a large enough sample size? Equations (3.17) and (3.20) show that the test statistic distribution converges (in other words, eventually gets arbitrarily close) to the standard normal, provided the sample size is large. However, they make no claim on how quickly the convergence is happening or at what sample size the test statistic distribution will be practically indistinguishable from the standard normal.

Readers may have come across many resources suggesting $n = 30$ as the “magic number” where beyond that, the test statistic is effectively normal [38]. Some might go further and suggest one should switch from a t -test to a z -test beyond that point, regardless of the underlying distribution of the data [258]. These suggestions may be acceptable for pedagogical purposes but are misleading or even harmful in digital experimentation. Firstly, they indicate a clear boundary exists where things are normal on one side and not on the other, whereas in reality, no such boundary exists – the t_{29} (from $n = 30$) and t_{30} (from $n = 31$) distributions are not much different from each other.²¹

²¹ The issue of having arbitrary thresholds for decision-making is pervasive in statistical testing. An obvious example would be the convention for one to reject the null hypothesis when the p -value drops below a certain threshold (see Section 3.6.1), which has led to years of bitter debates between statisticians.

Secondly, $n = 30$ is usually insufficient for the test statistic distribution to practically resemble the standard normal. This is the case even when the normality assumption for the data holds. Consider the t_{30} distribution, the distribution that the one-sample Welch's t -test statistic follows when $n = 31$. While its PDF may look similar to the standard normal distribution on a plot, the tails of the two distributions are still substantially different. We can see this via the 95% z - and t -quantiles:

$$z_{0.95} = t_{30, 0.94478...} \text{ and } z_{0.9517...} = t_{30, 0.95} . \quad (3.21)$$

The left equation above tells us that a t_{30} -distributed test statistic has a 5.52% chance to exceed the 95% z -quantile (the lower bound of the critical region for a greater/right-tailed z -test, see Equation (3.9)). Thus, if we assume the t_{30} -distributed test statistic is sufficiently normal and decide whether to reject based on the normal (instead of t_{30}) quantiles, we will inflate the Type I error from 5% to 5.52%, or 10% relatively. Such bias is further amplified at lower significance levels – for the 99% z - and t -quantiles,

$$z_{0.99} = t_{30, 0.98652...} \text{ and } z_{0.99299...} = t_{30, 0.99} . \quad (3.22)$$

In this case, rejecting at a 1% significance level using the normal (instead of t_{30}) quantiles will inflate the Type I error from 1% to 1.35%, or 35% relatively. One will require hundreds (if not thousands) of samples to effectively limit such bias,²² which, fortunately, is easily achievable in digital experiments.

It is worth reiterating that the bias above occurs even when the normality assumption in the data holds. A separate source of bias comes from when we deviate from the normality assumption. In this case, [150] provided a rule of thumb on the minimum sample size required per group for the test statistic to be sufficiently normal-like to run a practical t -test. It is based on the underlying distribution's moment coefficient of skewness:

$$355 \times s^2, \text{ where } s = \frac{\mathbb{E}[(X - \mu_X)^3]}{\sigma_X^3} , \quad (3.23)$$

²² For example, to limit such bias to under 1% relative to the significance level, we require at least 315 samples (per group) for a test at a 5% significance level and 1,000 samples for a test at a 1% significance level.

In practice, one usually plugs the empirical moment coefficient of skewness into Equation (3.23) to estimate the required sample size. The estimated sample size is most useful when $|s| > 1$. Needless to say, experimenters should ensure they have a sufficient sample size that can address both sources of bias. This is in addition to having enough samples to achieve sufficient test power (see Section 3.6.5).

3.5.3 The binomial test of proportions

Another popular statistical test deals with binary responses. By statistical convention, the response of interest is referred to as success and the other as failure. The successes and failures are often modelled by a Bernoulli r.v. The test decides whether the proportion of successes from the observed responses equals a pre-specified probability.

Definitions Let $X_1, \dots, X_n$ be independent, identically, and Bernoulli distributed responses with parameter π_X . We also let B be the test statistic formed by summing the responses. We observe that the test statistic follows the binomial distribution:

$$B = \sum_{i=1}^n X_i \sim \text{Bin}(n, \pi_X). \quad (3.24)$$

Unlike the z -test or t -test, the test statistic distribution depends on the parameter of interest π_X . Thus, one should define their statistical hypotheses if they have not yet done so. The direction of the null hypothesis does not matter at this stage – the test statistic follows the same distribution under the two-sided ($H_0 : \pi_X = \theta_0$) and one-sided null hypotheses ($H_0 : \pi_X \leq \theta_0$ or $H_0 : \pi_X \geq \theta_0$):

$$B \sim \text{Bin}(n, \theta_0). \quad (3.25)$$

p -value & critical region Suppose one then proceeds with the test statistic calculation and obtains k (i.e., they have observed k successes amongst n responses). Its associated p -value,

the probability that B attains k or more extreme values (see Section 3.6.1), is

$$p = \begin{cases} \sum_{i=k}^n \mathbb{P}(B = i) & \text{under a greater/right-tailed test } (H_0: \mu_X \leq \theta_0, H_1: \mu_X > \theta_0) \\ \sum_{i=0}^k \mathbb{P}(B = i) & \text{under a less/left-tailed test } (H_0: \mu_X \geq \theta_0, H_1: \mu_X < \theta_0) \\ \sum_{i \text{ s.t. } \mathbb{P}(B=i) \leq \mathbb{P}(B=k)} \mathbb{P}(B = i) & \text{under a two-sided test } (H_0: \mu_X = \theta_0, H_1: \mu_X \neq \theta_0), \end{cases} \quad (3.26)$$

where $\mathbb{P}(B = i) = \binom{n}{i} (\pi_X)^i (1 - \pi_X)^{n-i}$. Note how the summation indices change depending on the statistical hypothesis. For two-sided p -values, we obtain the summation indices by taking values that are equally or less likely to occur (as “more extreme”) than k . Such an elaborate definition is required as most binomial distributions are asymmetric. This renders the approach used by z - and t -tests to obtain the p -value, i.e., taking the probability of the relevant tail and multiplying by two, infeasible.

We also note that a binomial r.v. is discrete. Thus, we define a critical set instead:

$$\mathbb{T}_{\text{crit}} = \begin{cases} \left\{ i \mid i \in \mathbb{N}^0, F_{\text{Bin}(n, \theta_0)}^{-1}(1 - \alpha) < i \leq n \right\} & \text{under a greater/right-tailed test} \\ \left\{ i \mid i \in \mathbb{N}^0, 0 \leq i < F_{\text{Bin}(n, \theta_0)}^{-1}(\alpha) \right\} & \text{under a less/left-tailed test,} \end{cases} \quad (3.27)$$

where $F_{\text{Bin}(n, \theta_0)}^{-1}(\cdot)$ is the quantile function of a binomial. Two observations arise: Firstly, we exclude the binomial quantiles from the critical sets as it ensures the probability of the test statistic attaining a value in the critical set would not exceed the significance level α . Secondly, excluding the binomial quantiles means the critical region can be an empty set, usually when n is small. In such a scenario, the test is considered poorly designed as it is impossible to reject the null hypothesis.

The critical set for a two-sided binomial test is more complicated than that for a one-sided test: one can use the $\alpha/2$ and $1 - \alpha/2$ binomial quantiles, but the resultant set rarely agrees with what one would obtain using the p -values, especially when n is small. This brings unnecessary confusion; thus, we omit it from Equation (3.27). In practice, when n is small, it is easier for one to decide whether to reject using just the p -values. For large n , one generally uses the critical region provided by a practical t -test, as shown below.

Practical usage The test described above is *exact* – it gives exact probabilities provided all the assumptions are met. It generally works well for small n , though it is inadequate for the sample size featured in digital experiments due to the need to calculate the probability mass for each possible value B can take.

Instead, digital experimenters employ approximation methods. They appeal to the normal approximation to the binomial distribution (a special case of the CLT), which states that if the sample size n is large and the success rate π_X is not too close to zero or one,²³

$$B \stackrel{\text{approx.}}{\sim} \mathcal{N}(n\pi_X, n\pi_X(1 - \pi_X)) . \quad (3.28)$$

Another way to represent the same approximation is to divide the test statistic by n . The resultant test statistic is the mean of n Bernoulli r.v.'s, representing the proportion of successes:

$$B' \triangleq \frac{B}{n} \stackrel{\text{approx.}}{\sim} \mathcal{N}\left(\pi_X, \frac{\pi_X(1 - \pi_X)}{n}\right) \iff \frac{B' - \pi_X}{\sqrt{\frac{\pi_X(1 - \pi_X)}{n}}} \stackrel{\text{approx.}}{\sim} \mathcal{N}(0, 1) . \quad (3.29)$$

This enables one to use the practical t -test described in Section 3.5.2. Since there is no straightforward equivalent for the exact binomial test for two-sample comparison, digital experimenters almost always use the practical t -test with the test statistic defined in Equation (3.12).

3.6 Key Quantities in NHST

We introduce some quantities in NHST that are key to the design and analysis of digital experiments.

3.6.1 p -value

We start with p -values, one of the most commonly reported measures⁷ in NHSTs. Its construction is straightforward, with the resultant formulas/routines for different tests readily

²³ The literature usually states the requirements as $n\pi_X > 5$ and $n(1 - \pi_X) > 5$ [241]. Some texts recommend that they should be greater than 10 to be safe.

implemented in statistical software packages. However, perhaps due to the widespread availability and ease of access, it is also one of the most misconstrued quantities in NHSTs. The section aims to highlight the pitfalls surrounding its usage.

As alluded to in Section 3.4, the p -value of a test is often descriptively defined as

$$p \triangleq \mathbb{P}(T \text{ attaining } t_{\text{obs}} \text{ or more extreme values} \mid H_0 \text{ is true}). \quad (3.30)$$

A more mathematically precise definition of p -value depends on the hypotheses and test, as those affect what constitutes extreme. See Equations (3.10), (3.15), and (3.26) in the previous section for examples. One can also find a more abstract definition of p -value in statistical theory literature, e.g., Definition 8.3.26 in [36].

The definition above applies to a simple null hypothesis, e.g., a nil hypothesis within a comparison test. Under a composite null hypothesis, e.g., a one-sided test,

$$p = \sup_{h \in H_0} \mathbb{P}(T \text{ attaining } t_{\text{obs}} \text{ or more extreme values} \mid h \text{ is true}). \quad (3.31)$$

Each possible null hypothesis leads to a different test statistic distribution and, hence, a different p -value candidate for the supremum. At first glance, it seems like much work to derive all the test statistic distributions and calculate the p -value under every possible null hypothesis. In reality, virtually all named one-sided tests used in digital experimentation attain the supremum at the hypothesis boundary (e.g., parametric tests with $H_0 : \Delta \leq \theta_0$, $H_1 : \Delta > \theta_0$ attain the supremum at $\Delta = \theta_0$). Thus, computation of the test statistic and p -value always focuses on this single hypothesis as a simple hypothesis [130].

The widespread usage of p -values may give the impression that they give extensive information about the experiment. They do not. They have a fairly narrow scope: To the fullest extent of the interpretation, p -values can be characterised as a measure of the degree of surprise²⁴ for the data in hand under the assumptions imposed by the specific null hypothesis. Writing such characterisation as $\mathbb{P}(\text{Data} \mid H_0)$, we can see that changing the null hypothesis, which includes the underlying statistical model and experiment assumptions, will lead to a different, incomparable measure.

²⁴ Some prefer a degree of *discrepancy* [21] or *incompatibility* [276] instead.

Crucially, p -values, by themselves and as characterised above, tell us nothing about the null hypothesis and the data-generating process. Ideally, one would decide whether to reject the null hypothesis based on $\mathbb{P}(H_0)$ or $\mathbb{P}(H_0 \mid \text{Data})$, i.e., the chance that the null hypothesis is actually true, with or without observing the data, respectively. In reality, both probabilities are generally unavailable under the NHST framework. Using p -values as a substitute for the same ideal is where one can get tripped up. Clearly, $\mathbb{P}(\text{Data} \mid H_0)$ and $\mathbb{P}(H_0 \mid \text{Data})$ are different. The former – p -values – do not tell us the probability that the null hypothesis is true, as the entire p -value calculation assumes that the null hypothesis is the preferred state of affairs, the ground truth.

It is possible to circumvent the restriction above by making assumptions about $\mathbb{P}(H_0)$ or $\mathbb{P}(H_0 \mid \text{Data})$. Indeed, [248] showed that under some prior belief on how the parameters involved in H_0 may behave, $\mathbb{P}(\text{Data} \mid H_0)$ and $\mathbb{P}(H_0 \mid \text{Data})$ can be equivalent for certain statistical models and hypotheses.²⁵ That said, such assumptions would bring us into Bayesian territory (see Section 3.9), and we recommend readers who are new to p -values and their role in the frequentist vs Bayesian debate to revert to such equivalence only after familiarising themselves with the two perspectives.

In addition, p -values also tell us very little about the effect size of treatment and whether it is important in practice. Again, p -values merely represent how (in)compatible the data is with the specific null hypothesis assumptions. With carefully crafted hypotheses that seek to demonstrate specific effect sizes, e.g., those in non-inferiority or equivalence tests, one can indeed infer a non-negative or zero effect size solely from a small p -value. However, most digital experiments do not seek to demonstrate specific effect sizes (beyond, e.g., “not zero”). Thus, their null hypothesis specification and p -value calculations are generally ignorant to the underlying effect size despite the latter usually being encoded in the data. To appreciate the extent of decoupling between p -values and effect sizes in these experiments, we note that for any given effect size, it is possible to construct statistical tests that yield an arbitrarily chosen p -value.²⁶ One usually needs to report quantities discussed in Section 3.3.2 to communicate

²⁵ More specifically, [248] showed that under non-informative priors, the p -value of a one-sided NHST and the posterior probability of a corresponding Bayesian test are (asymptotically) equivalent. Such equivalence applies for binary data with Jeffrey’s prior and normal data with an improper flat prior. For two-sided tests, they demonstrated equivalence between the two concepts by reformulating a two-sided test as two one-sided tests and constructing a “two-sided posterior probability” that matches the p -value.

²⁶ One can do so by changing, e.g., the sample size or the null hypothesis parameter(s).

effect sizes. In all cases, they would also require domain expertise to appreciate the effect size in the relevant context and, more importantly, to decide the next steps based on the results.

Despite what p -values do not tell us, it remains useful in statistical testing to identify and potentially sift out specific statistical hypotheses and assumptions that the data is incompatible with. It is just not the one quantity to rule them all. In fact, no quantities in statistical testing assume such a role. Experiment analysis is nuanced, and we require multiple sources of signal, including but not limited to the concepts introduced in Section 3.3, to understand the big picture and make well-informed decisions. This cannot happen if we over-emphasise reporting a single quantity that does not represent the experiment results. The quantity could also be prone to manipulation, both intentionally and unintentionally,²⁷ making results less interpretable and thus less trustworthy.

To summarise the above, it is perhaps best to use the (slightly rearranged) words of the American Statistical Association [276]:

“ p -values can indicate how incompatible the data are with a specified statistical model [(one used by the null hypothesis)].”

“[They] do not measure the probability that the [scientific/statistical] hypothesis is true, or the probability that the data were produced by random chance alone.”

“[They also] do not measure the size of an effect or the importance of a result.”

“Proper inference requires full reporting and transparency.”

“By itself, a p -value does not provide a good measure of evidence regarding a model or hypothesis.”

“[One should not base] scientific conclusions and business or policy decisions [...] only on whether a p -value passes a specific threshold.”

3.6.2 Type I and II errors

The rest of the quantities stem from decisions made in an NHST. Here, we consider what happens when we decide between two competing statistical hypotheses.

²⁷ For p -value, it is known as data dredging [252] or “ p -hacking” [276].

Treating NHST as a binary decision problem enables us to analyse the different kinds of errors we commit. Of course, we do not know which of the null hypothesis (H_0) or the alternate hypothesis (H_1) is actually true – otherwise, we do not need a statistical test at all. We can, however, consider what would have happened had either of the hypotheses been true. This gives us four possible scenarios, as shown below:

	H_0 is true	H_1 is true
Test does not reject H_0	Correct decision / True negative	Type II error / False negative
Test rejects H_0	Type I error / False positive	Correct decision / True positive

Knowing how often each of the four scenarios happens enables us to quantify the performance of a decision rule. We are usually interested in the Type I and Type II error rates in statistical testing. The former is defined as the chance of committing a Type I error when H_0 is true (see Section 3.6.3), and the latter is defined as the chance of committing a Type II error when H_1 is true (see Section 3.6.4). A good experiment design always controls for these two types of error, though we will show in Section 3.6.4 that it is more of a trade-off.

It is worth noting that many other quantities exist to evaluate different aspects of a decision rule. To add to the complication, different communities may refer to the same quantity differently. For example, readers with a clinical background may recognise *specificity* (one minus the Type I error rate) and *sensitivity* (one minus the Type II error rate), while those with a computer science or machine learning background may be used to *precision* (one minus the false positive risk, see Section 3.6.3) and *recall* (another name for sensitivity). For consistency, we will stick to the terms presented in the table above for the remainder of the section.

3.6.3 Significance level

The significance level, as commonly referred to in digital experimentation,²⁸ of an NHST is defined as

$$\alpha \triangleq \mathbb{P}(\text{test rejects } H_0 \mid H_0 \text{ is true}). \quad (3.32)$$

²⁸ Strictly speaking, in statistical theory, Equation (3.32) describes the *size* of a statistical test, with the significance level being a property of a class of tests [183]. Having that said, in nearly all practical cases in digital experimentation, the size and level of a test are equal, and there is little practical impact in confusing the terms in our context.

It is usually specified as a parameter (as opposed to a calculated property) of a test and determines the specific form of the rejection rules.

Similar to the p -value, the definition of the significance level is slightly different when a composite null hypothesis (e.g., in a one-sided test) is involved:

$$\alpha = \sup_{h \in H_0} \mathbb{P}(\text{test rejects } H_0 \mid h \text{ is true}). \quad (3.33)$$

Under the same arguments for p -values, computation upon the specification of α always focuses on the boundary hypothesis as if it is a simple hypothesis. That said, given that a decision is now involved, experimenters should also take particular care when interpreting a test that failed to reject the null – clearly, only a particular (but not the full range of the) null hypothesis can be true.

The parameter signifies the *false positive ratio/rate* experimenters are willing to take. While we are probably more interested in controlling the confusingly similarly named *false positive risk* (i.e., $\mathbb{P}(H_0 \text{ is true} \mid \text{test rejects } H_0)$, with the events conditioned the other way round), it is simply not a feature of NHSTs.

There are no rules on what α should be. Common sense dictates that it should be small – a test is no good if it rejects the null hypothesis more often than not when the null hypothesis is true. Digital experimenters generally take $\alpha = 10\%$, 5% or 1% , which stemmed from the values used in clinical trials decades ago, themselves chosen more as a matter of convenience (e.g., access to statistical tables presented with arbitrarily decided thresholds) rather than principle [162].

It is also worth noting that some, including a large digital A/B testing platform provider [211], refer to the complement of the p -value (i.e., $1 - p$) as “significance”. Fortunately, the two usages usually yield values on different extremes and thus create little confusion once the experiment context is clear.

3.6.4 Power

The power of an NHST is defined as

$$1 - \beta \triangleq \mathbb{P}(\text{test rejects } H_0 \mid H_1 \text{ is true}), \quad (3.34)$$

where β is the probability of committing a Type II error / false negative.

In the case where H_1 is a composite hypothesis (which is the case for all one-sided and two-sided tests), such probability is difficult to obtain. To illustrate, in parametric tests utilising a model with parameters $\theta \in \Theta$, we have to consider the test rejection probability under every parameter combination to obtain $1 - \beta$, i.e., we have to calculate

$$\int_{\theta \in \Theta} \underbrace{\mathbb{P}(\text{test rejects } H_0 \mid \theta)}_{\triangleq 1 - \beta_\theta} f(\theta \mid H_1 \text{ is true}) d\theta. \quad (3.35)$$

The second term of the integrand, which describes the distribution of the parameters under H_1 , is usually unknown in an NHST.²⁹ Thus, we generally refer to the first term of the integrand, i.e., the probability of the test rejecting H_0 *given a specific alternate hypothesis*, when we talk about the power.

Unlike the significance level (α), one can specify the power as a parameter *a priori* or calculate the *a posteriori/post hoc* power. The *a priori* power,³⁰ commonly denoted $\pi_{\min}$, signifies the minimum power one desires for the upcoming test. Equivalently, via β , it represents the false negative ratio/rate they are willing to take. Clearly, we want the power to be close to one, and it is common to set $\pi_{\min} = 80\%$ or 90% in digital experiments. As shown below, specifying the power will affect what is possible with other test parameters, such as sample size (see Section 3.6.5) and effect size (see Section 3.6.6).

One can calculate the *a posteriori/post hoc* power before or after running the test. Both approaches answer, “Given all the other test parameters, what is the chance that the test detects a certain effect size if it actually exists?” The difference comes from the timing: in pre-test

²⁹ Imposing any assumptions on the parameter distribution will bring us into Bayesian territory (see Section 3.9). However, most Bayesian testing approaches do not feature the concept of power as they do not involve a decision.

³⁰ Also known as *prospective* power [87].

calculations, one often explores a range of potential test parameters based on their experience. In contrast, in post-test calculations, one often uses the observed effect size, response variance, and sample size as if they are the true values.

Calculating the power post-test using the observed effect size, known as a *retrospective* or *observed* power analysis, is controversial [94, 125, 151].³¹ This is because it relies on the unlikely assumption that the observed effect size, derived from a sample, is identical to the underlying population effect size. Its use is nonetheless common in other applications and, thus, not unheard of in digital experiments as experimenters move fields [151].

Example: z -test Consider a two-sample, one-sided z -test with $H_0 : \Delta = \theta_0$ and $H_1 : \Delta > \theta_0$. The test statistic under H_0 is $z = ((\bar{y} - \bar{x}) - \theta_0) / \sqrt{\frac{\sigma_X^2}{n} + \frac{\sigma_Y^2}{m}}$, and we would reject the null hypothesis if $z > z_{1-\alpha}$. The power for a specific $H_1 : \Delta = \theta$ is thus

$$\begin{aligned}
 1 - \beta_\theta &= \mathbb{P}(Z > z_{1-\alpha} \mid \Delta = \theta) = 1 - \mathbb{P}(Z < z_{1-\alpha} \mid \Delta = \theta) \\
 &= 1 - \mathbb{P}\left(\frac{\bar{Y} - \bar{X} - \theta_0}{\sqrt{\frac{\sigma_X^2}{n} + \frac{\sigma_Y^2}{m}}} < z_{1-\alpha} \mid \Delta = \theta\right) && \text{(substituting def. of } Z) \\
 &= 1 - \mathbb{P}\left(\frac{\bar{Y} - \bar{X} - \theta + \theta - \theta_0}{\sqrt{\frac{\sigma_X^2}{n} + \frac{\sigma_Y^2}{m}}} < z_{1-\alpha} \mid \Delta = \theta\right) && \text{(inserting self-cancelling } \theta \text{ pair)} \\
 &= 1 - \mathbb{P}\left(\frac{\bar{Y} - \bar{X} - \theta}{\sqrt{\frac{\sigma_X^2}{n} + \frac{\sigma_Y^2}{m}}} < z_{1-\alpha} - \frac{\theta - \theta_0}{\sqrt{\frac{\sigma_X^2}{n} + \frac{\sigma_Y^2}{m}}} \mid \Delta = \theta\right) && \text{(moving const. terms to RHS).}
 \end{aligned} \tag{3.36}$$

We note that under this specific H_1 , the r.v. on the RHS of the inequality follows the standard normal. Thus, the power evaluates to

$$1 - \beta_\theta = 1 - \Phi\left(z_{1-\alpha} - \frac{\theta - \theta_0}{\sqrt{\frac{\sigma_X^2}{n} + \frac{\sigma_Y^2}{m}}}\right). \tag{3.37}$$

³¹ The terminology surrounding power analysis is neither accurate nor precise. Some uses have semantics that deviate from the dictionary definition, and there is no consensus on the names of different approaches. Some consider an *a posteriori*, *post hoc*, and retrospective/observed analysis synonymous [5, 243], others differentiate between an *a posteriori*/*post hoc* analysis and a retrospective/observed analysis [94, 285].

In a two-sample, *two-sided* z -test, we can obtain the power using the same procedure as Equations (3.36) and (3.37), noting that we would reject the null hypothesis if $|z| > z_{1-\alpha/2}$ instead and taking particular care when expanding the absolute function:

$$1 - \beta_\theta = \underbrace{\Phi \left(z_{\alpha/2} - \frac{\theta - \theta_0}{\sqrt{\frac{\sigma_X^2}{n} + \frac{\sigma_Y^2}{m}}} \right)}_{\text{Left-tail contribution}} + \underbrace{1 - \Phi \left(z_{1-\alpha/2} - \frac{\theta - \theta_0}{\sqrt{\frac{\sigma_X^2}{n} + \frac{\sigma_Y^2}{m}}} \right)}_{\text{Right-tail contribution}}, \quad (3.38)$$

where the right-tail contribution is of the same form as that in Equation (3.37) (albeit with a different z -quantile).

There are also two alternate formulations for Equation (3.38). The first, which we use in Section 5.5, assumes one uses a nil hypothesis as the null. In this case, $\theta_0 = 0$ and

$$1 - \beta_\theta = \underbrace{\Phi \left(z_{\alpha/2} - \frac{\theta}{\sqrt{\frac{\sigma_X^2}{n} + \frac{\sigma_Y^2}{m}}} \right)}_{\text{Left-tail contribution}} + \underbrace{1 - \Phi \left(z_{1-\alpha/2} - \frac{\theta}{\sqrt{\frac{\sigma_X^2}{n} + \frac{\sigma_Y^2}{m}}} \right)}_{\text{Right-tail contribution}}. \quad (3.39)$$

Another formulation exploits that in Equation (3.39), the left-tail contribution is negligible when the ratio between θ and the standard error $\sqrt{\sigma_X^2/n + \sigma_Y^2/m}$ is large (i.e., $\theta/\sqrt{\dots} \gg 0$), the right-tail contribution is negligible when the ratio is negatively large (i.e., $\theta/\sqrt{\dots} \ll 0$), and the functions representing the two contributions are the horizontal reflection of each other along $\theta = 0$.³² This enables one to approximate the power using only the right-tail contribution term (similar to that in Equation (3.37)):

$$1 - \beta_\theta \approx 1 - \Phi \left(z_{1-\alpha/2} - \frac{|\theta|}{\sqrt{\frac{\sigma_X^2}{n} + \frac{\sigma_Y^2}{m}}} \right). \quad (3.40)$$

This form is attractive as it eases downstream algebraic manipulation – one only needs to deal with one normal CDF when making other parameters the subject of the equation (see Sections 3.6.5 and 3.6.6). However, it does require the ratio between the absolute value of θ

³² What matters is the magnitude of the entire ratio instead of its components (θ or the standard error). In digital experiments, one usually deals with small effect sizes and hence θ (see Section 3.6.6 and the “Improving sensitivity” paragraphs in Section 5.4.1 for further discussions). However, the same experiments often involve a large sample size (n and m), leading to a workable ratio.

and the standard error to be large (i.e., $|\theta|/\sqrt{\cdots} \gg 0$).³³ This formulation is used in the following sections and Chapter 6, where the ratio is assumed to be large enough to benefit from this approximation.

Determining factors In general, the power depends on many factors:

1. The significance level (α) - power increases when one increases α as β_θ , the Type II error rate, decreases;
2. The sampling uncertainty (σ^2) - power increases when σ^2 decreases as one is more likely to observe the same treatment effect under a lower noise level;
3. The sample size (n/m) - power increases when n/m increases due to the same reason above; and
4. The effect size (e.g., θ , assuming $\theta_0 = 0$ as in common practice) - power increases when θ increases as one is more likely to observe a larger effect size under the same noise level.

Returning to our z -test examples, we observe from Equations (3.37) and (3.38) that power depends entirely on the above factors. In other words, assuming we can specify the value of $1 - \beta_\theta$ alongside the other parameters above, specifying any four out of the five (groups of) parameters will fix the final one. This gives experimenters multiple options during the power analysis stage that were alluded to in Section 3.4.

3.6.5 Required sample size

One common approach to conducting a power analysis seeks to obtain the sample size required for the statistical test and experiment. This is done by either solving for the sample size (n or m) in a known power formula given all other parameters. Alternatively, one finds the sample size that leads to a desired power value while holding other parameters constant in a power simulation. The approach is popular in medical experiments as experimenters have a higher influence on sample size (compared to sampling uncertainty/response vari-

³³ In practice, we can obtain good approximations when the ratio is greater than $z_{1-\alpha/2} + 1$. At $\alpha = 5\%$, the quantity is approximately equal to three.

ance or the effect size). It also appeals to digital experimenters as the number informs how long it takes to collect the samples, a sizeable portion of an experiment's overall duration.³⁴

Example: z -test We can obtain the required sample size for a z -test by specifying the desired power as $\pi_{\min} < 1 - \beta_\theta$, substituting Equation (3.37) into the inequality, and rearranging the terms to make n or m the subject. Usually, one assumes $n = m$, i.e., the two groups have an equal sample size, as that generally yields the highest power and, thus, the lowest sample size requirement (assuming the two samples have similar variance).³⁵ This yields

$$\pi_{\min} < 1 - \Phi \left(z_{1-\alpha} - \frac{\theta - \theta_0}{\sqrt{\frac{\sigma_X^2}{n} + \frac{\sigma_Y^2}{n}}} \right) \iff n > \frac{(z_{1-\alpha} - z_{1-\pi_{\min}})^2 (\sigma_X^2 + \sigma_Y^2)}{(\theta - \theta_0)^2}. \quad (3.41)$$

We use the power definition in Equation (3.40) instead for a two-sided test instead; otherwise, the procedure is identical.

A helpful rule of thumb [23, 155, 194] is to express the minimum required sample size as

$$n = \frac{16\sigma^2}{\theta^2}, \quad (3.42)$$

where σ^2 is the population variance of *one* group, and θ is the effect size as defined and used throughout this section (assuming $\theta_0 = 0$). The factor of 16 comes from rounding up the product of $(z_{1-\alpha} - z_{1-\pi_{\min}})^2$ and the factor of two extracted from $\sigma_X^2 + \sigma_Y^2$ in Equation (3.41) (assuming $\alpha = 5\%$, $\pi_{\min} = 80\%$, and $\sigma_X^2 \approx \sigma_Y^2$ as per common practice).

One can also observe from the equations above that the required sample size n is inversely proportional to the square of the effect size θ . In other words, all other things being equal, one requires four times the number of samples to detect half the effect size. This highlights

³⁴ Translating the required sample size to the number of days required to collect the samples can be either straightforward or tricky, depending on the experimental units. If the responses are based on website visits, the translation is as simple as dividing the required sample size by the average daily number of visits. On the other hand, if the responses are based on visitors who can return if and when they please, we require more advanced methods to estimate the time taken to observe a certain number of visitors [228, 296].

³⁵ Let $\text{se}(n) = \sqrt{\frac{\sigma_X^2}{n} + \frac{\sigma_Y^2}{n_{\text{tot}} - n}}$ be a function of the standard error to the sample size of the first group, where σ_X^2 and σ_Y^2 are the population variances, and n_{tot} is the total number of experimental units available to be split into the two groups. Minimising $\text{se}(n)$ will maximise the power in Equations (3.37) and (3.40). One can show (via, e.g., a derivative test) that $\text{se}(n)$ is minimised when $n = \frac{n_{\text{tot}}}{\sigma_Y/\sigma_X + 1}$. When $\sigma_X^2 = \sigma_Y^2$, $\text{se}(n)$ is minimised when $n = n_{\text{tot}}/2$, i.e., when the two groups have equal sample size.

the importance of having a sensible effect size estimate. One can either draw on results from previous experiments or, without prior knowledge, use the minimum practically significant effect size – one that would matter in practice.

3.6.6 Minimum detectable effect

Another useful approach to power analysis is to compute the required effect size θ based on other experiment parameters and see whether it is practically achievable. As power increases when the effect size increases, we are interested in the minimum effect size that yields the desired power, in other words, can be detected (via rejecting the null hypothesis) reliably (at the specified power level). Similar to the required sample size, this can be done by solving a known power formula (see below) or by simulation.

The approach appeals to those with little influence on other experiment parameters (e.g., sample size and response variance) due to operational constraints. Given the fact that digital experiments often feature small effect sizes [150, 286], one may leverage a power analysis that requires an unrealistically large effect size to request better resources for their planned experiment, e.g., more participants or investment in methods that improve experiment sensitivity (see Section 5.4.1).

We will also discuss how the minimum detectable effect (MDE) of a test comes into play when we evaluate competing digital experiment designs in Chapter 6.

Example: z -test Similar to how we obtain the required sample size, we can obtain the MDE by specifying the desired power as $\pi_{\min} < 1 - \beta_{\theta}$, substituting Equation (3.37) into the inequality, and rearranging the terms to make θ the subject. The only difference is that we do not need to assume anything about the sample size. This yields

$$\pi_{\min} < 1 - \Phi \left(z_{1-\alpha} - \frac{\theta - \theta_0}{\sqrt{\frac{\sigma_X^2}{n} + \frac{\sigma_Y^2}{m}}} \right) \iff \theta - \theta_0 > (z_{1-\alpha} - z_{1-\pi_{\min}}) \sqrt{\frac{\sigma_X^2}{n} + \frac{\sigma_Y^2}{m}}. \quad (3.43)$$

The MDE, commonly denoted θ^* , is then the minimum $\theta - \theta_0$ that satisfies Inequality (3.43), i.e., that specified by the RHS of the inequality. Some variants of Inequality (3.43) may re-

place $\theta - \theta_0$ and $z_{1-\alpha}$ with, respectively, $|\theta - \theta_0|$ and $z_{1-\alpha/2}$ (by assuming a two-sided test and using Equation (3.40) instead), or leave out θ_0 (by assuming $\theta_0 = 0$), or both.

3.7 Non-parametric Tests

So far, we have discussed NHSTs comparing statistical hypotheses on the value of one or more statistical model parameters. In some cases, such hypotheses may not suit an experimenter's need – they may not want to assume the data is generated by a specific distribution, or they may feel the test assumptions are too restrictive.

Instead of inferring specific parameter(s) in a statistical model, a non-parametric test focuses on the overall distributions. In a one-sample setting, the statistical hypotheses often have the following form:

$$H_0 : F_X(\cdot) \equiv F(\cdot), \quad H_1 : F_X(\cdot) \not\equiv F(\cdot), \quad (3.44)$$

i.e., whether a sample X is generated from a specified distribution. Whereas in a two-sample setting, the hypotheses often have the following form:

$$H_0 : F_X(\cdot) \equiv F_Y(\cdot), \quad H_1 : F_X(\cdot) \not\equiv F_Y(\cdot), \quad (3.45)$$

i.e., whether two samples X and Y are generated by the same distribution. Similar to parametric tests, one-sided non-parametric tests exist and are often used, though different tests may have different definitions of what constitutes “greater” or “less”.

This section will introduce two non-parametric tests commonly used in digital experiments: the Mann-Whitney U test (Section 3.7.1) and the χ^2 goodness-of-fit test (Section 3.7.2). Other examples, which we omit for brevity, include the Wilcoxon signed-rank [281], Kruskal-Wallis [157], and Kolmogorov–Smirnov tests [75].

3.7.1 Mann-Whitney U test

A Mann-Whitney U test (also known as a Wilcoxon rank-sum test [163]) is a two-sample test determining whether the underlying distributions are equivalent [181]. Here, the null hypothesis is similar to that in Expression (3.45). In contrast, the greater (less than) alternative hypothesis is that one is more (less) likely to see a randomly chosen observation from a population Y exceeding in value than a randomly chosen observation from the other population X and, therefore, Y is greater (less) than X overall.³⁶

Let $X_1, \dots, X_n$ and $Y_1, \dots, Y_m$ be our samples. We make no assumptions on the distribution of X and Y except that they are at least ordinal (i.e., there is some notion of a response being greater/less than another), and each copy is i.i.d. The test statistic is defined as

$$U \triangleq \sum_{i=1}^n \sum_{j=1}^m S(X_i, Y_j), \quad \text{where } S(X, Y) = \begin{cases} 1 & \text{if } Y < X, \\ 1/2 & \text{if } Y = X, \\ 0 & \text{if } Y > X. \end{cases} \quad (3.46)$$

For large n and m , a rank-based method that generates the same test statistic while avoiding the inefficient double sum is available. Statistical software packages should handle all this automatically upon presenting the responses.

When the sample size is large, U approximately follows a normal distribution:

$$\frac{U - \mu_U}{\sigma_U} \underset{\text{approx.}}{\sim} \mathcal{N}(0, 1), \quad \text{where } \mu_U = \frac{n \cdot m}{2}, \quad \sigma_U = \sqrt{\frac{n \cdot m \cdot (n + m + 1)}{12}}. \quad (3.47)$$

This enables us to use the p -values and critical regions of an approximate z -test (see Section 3.5.1) to make a reject/not reject decision. In cases with many ties, one would adjust σ_U to account for the $1/2$ contributions to the test statistic [163].

³⁶ More formally, the greater (less than) alternative hypothesis says Y is *stochastically greater (less)* than X , defined as the probability that an observation from Y exceeds that from population X is greater (less) than the probability in the reverse case, i.e., observation from X exceeds that from Y . This is defined mathematically as $\mathbb{P}(Y > X) > \mathbb{P}(X > Y)$ ($\mathbb{P}(Y > X) < \mathbb{P}(X > Y)$). The two-sided alternative hypothesis states that the two probabilities are unequal.

Practical usage The Mann-Whitney test is often employed as an alternative to the (practical) t -test. This may be due to a severe deviation from the normality assumptions in the data or an insufficient sample size for one to rely on the CLT or both.

That said, one should not mindlessly use a Mann-Whitney test in lieu of a t -test. Strictly speaking, switching from a t -test to a Mann-Whitney test removes one's ability to produce statements on the effect size of a treatment described in Section 3.3.2. This is due to the statistical hypotheses involved – a t -test infers on the (difference in) population means while a Mann-Whitney test infers the overall distributions. At best, one can make claims like “overall engagement has improved” or “users are spending more in general” directly off the result of the test, but not “average engagement has improved by 0.2 units” or “mean spend per user has increased by 83 pence”.³⁷ Moreover, a Mann-Whitney test can be less efficient, or in other words, requires more samples to reject a nil-as-null hypothesis, than a t -test given the same effect size if the underlying distributions are sufficiently normal-like (see Section 3.5.2 on relevant heuristics).

Some literature refers to the Mann-Whitney test as a test for difference in medians. This is only true if one is willing to assume, in addition to that stated at the start of the section, that (1) the underlying probability distributions X and Y are continuous and (2) the distribution of Y is a location-shifted version of X , i.e. $F_X(x) = F_Y(x + \delta)$ for some value of δ . Failing these assumptions, it is possible to construct r.v.'s X and Y such that the median of Y is equal to / less than X , yet a Mann-Whitney test will reject the null hypothesis in favour of a greater alternative hypothesis (i.e., Y is stochastically greater than X). Generally, a Mann-Whitney test considers both the location and shape of the two distributions we are comparing [118].

3.7.2 χ^2 goodness-of-fit test

We also discuss the χ^2 goodness-of-fit test, which deals with categorical responses. In particular, it is a test to decide whether an observed frequency distribution over multiple categories is the same as the theoretical frequency.

Suppose there are $|\text{Cats.}|$ categories, with the expected/theoretical and observed frequency

³⁷ Many digital experimenters make the latter two claims off a Mann-Whitney test regardless.

for each category i being E_i and O_i , respectively. The test statistic is defined as

$$C^2 = \sum_{i=1}^{|\text{Cats.}|} \frac{(O_i - E_i)^2}{E_i}. \quad (3.48)$$

Under the null hypothesis, it follows the χ^2 distribution with ν degrees of freedom (d.f.):

$$C^2 \sim \chi_\nu^2. \quad (3.49)$$

For a goodness-of-fit test, ν equals the number of categories with counts minus the number of restrictions.³⁸ The restrictions can be in the form of fitted parameters required to produce the expected count (taking one d.f. away per parameter), or rules that mandate how different sets of expected counts should behave, or both. Usually, there is only one rule for one-dimensional counts: the expected counts should sum up to the observed count sum. However, one may encounter many copies of such a rule along the marginals for multi-dimensional counts, substantially reducing the d.f.

We observe that the test statistic C^2 evaluates to zero when the observed frequency distribution is identical to the theoretical frequency distribution and grows large as the discrepancy between the two distributions increases. Thus, the test only makes sense as a greater/right-tailed test, and we reject the null hypothesis if C^2 exceeds the $1 - \alpha$ quantile of the χ_ν^2 distribution.

Practical usage The χ^2 test is often employed to check for sample ratio mismatch (SRM) in digital experiments. SRM occurs when the ratio between the sample size of two or more groups in an experiment deviates from the designed split. It often indicates one or more problems in software code that trigger a user's entry to an experiment, perform experimental unit randomisation, and track events. These problems often introduce biases that can invalidate experiment results [74, 155]. Of course, not all deviations are necessarily problematic. Small deviations can occur purely due to randomness, but the same deviation becomes

³⁸ The same χ^2 test statistic can also be used to test for homogeneity and independence. They feature a different approach to obtaining the degrees of freedom in the test statistic distribution.

unlikely as the sample size grows.³⁹ The χ^2 goodness-of-fit test formalises this process by quantifying how likely the observed deviation would happen, assuming the split is correct.

The SRM test operates like other χ^2 goodness-of-fit tests, with the expected frequency of each group obtained by multiplying the total number of samples by the designed split proportion. The d.f. is the number of groups minus one, reflecting the only restriction that the expected group counts should sum up to the total number of samples. Digital experimenters also often set a lower significance level, e.g., at $\alpha = 1\%$ or even $\alpha = 0.1\%$. This is done on practical grounds: investigating why a test rejects the null hypothesis (i.e., flagging a potential SRM) is often costly [93], incentivising experimenters to limit the number of false positives.

3.8 Sequential Tests

As stipulated in Section 3.4, NHSTs require one to decide and commit to the number of samples required prior to the start of the experiment. This prevents experimenters from continuously monitoring and stopping an experiment early – also known as *peeking* [134] and often used to achieve a shorter time-to-insight – as it inflates the Type I / false positive error rate due to the look-elsewhere effect.⁴⁰

Some experimenters turn to sequential tests to circumvent the restriction on peeking [242]. Sequential tests (a.k.a. sequential analysis or *interim analysis*) concern statistical tests that do not fix their sample size in advance. Instead, these tests evaluate the responses as they are observed. In addition to having a decision rule for the test outcome (e.g., reject/not reject H_0), a sequential test also features a *stopping rule*, which dictates whether an experimenter should continue or stop collecting new data. The decision rule in a sequential test may also depend on the stopping time or other elements of the stopping rule.

³⁹ If we perform a random 50/50 split on our stream of incoming users, we often see 105 users in one group and 95 in the other from the first 200 users. However, one would be rightly suspicious if we see 1 050 000 users in one group and 950 000 in the other, which indicates systemic bias.

⁴⁰ The look-elsewhere effect is best described as “The harder you look for something, the more likely you will find it.” Peeking is closely related to the multiple comparison problem (see Section 5.4). The key difference is that the former concerns observing and making decisions based on successive (dependent) test statistics within the same statistical test. In contrast, the latter concerns making decisions based on test statistics across different tests that may or may not be dependent.

Many examples of sequential tests exist and are used in digital experimentation [195]. Below, we introduce two sub-classes, including (mixed) sequential probability ratio tests and group sequential tests. Readers can refer to, e.g., [242], for a critique on the pros and cons of different sequential tests.

(Mixed) sequential probability ratio test One of the earliest tests in the sequential paradigm is the sequential probability ratio test (SPRT) [271]. It is a one-sample test that assumes the i.i.d. data $X_1, \dots, X_n, \dots, X_{n'}$ follows a parametric distribution and compares two simple hypotheses on a parameter of interest θ , i.e., $H_0 : \theta = \theta_0$ and $H_1 : \theta = \theta_1$. The test statistic S_n is the cumulative sum of the log-likelihood ratio up until the latest observed data point X_n . More specifically, $S_0 \triangleq 0$ and

$$S_n = S_{n-1} + \log \left(\frac{f(X_n | \theta_1)}{f(X_n | \theta_0)} \right). \quad (3.50)$$

The stopping and decision rules are jointly specified as follows. Upon observing X_n and calculating S_n :

- If $S_n \geq \log \left(\frac{1-\beta}{\alpha} \right)$, accept H_1 ;
- Else if $S_n \leq \log \left(\frac{\beta}{1-\alpha} \right)$, accept H_0 ,⁴¹
- Else (i.e., when $\log \left(\frac{\beta}{1-\alpha} \right) < S_n < \log \left(\frac{1-\beta}{\alpha} \right)$), continue to take observations;

where α and $1 - \beta$ correspond to the significance level and power (see Section 3.6) specified by the experimenter during the design stage, respectively.⁴²

Many extensions of SPRT have been developed since. The mixture SPRT (or mSPRT [229]) supports the now more commonly used hypothesis pair $H_0 : \theta = \theta_0$ and $H_1 : \theta \neq \theta_0$ by using a mixture distribution H to specify likely values of θ among all possibilities. In other words, it specifies $H_1 : \theta \sim H$. One then integrates the likelihood ratio across the PDF of H ,

⁴¹ Readers may notice the test refers to actions such as accepting H_0 , something not supported in a classical NHST. This is not a typo. The SPRT stems from the (Neyman–Pearsonian) hypothesis test framework, an “ancestor” of the NHST framework. Instead of giving the null hypothesis a preferred status like that in an NHST, a hypothesis test decides between two competing hypotheses without any prior judgment and thus can accept either hypothesis. The test also does not have the notion of p -value, a feature of a (Fisherian) significance test [166].

⁴² We refer readers to the original work [271] for a full justification of *why* these specific decision/stopping boundaries will control the Type I and II errors at the specified level.

denoted $f_H(\theta)$, when calculating the test statistic. This is defined, upon observing the first n samples, as

$$\Lambda_n^{H,\theta_0} \triangleq \int \prod_{i=1}^n \frac{f(X_i|\theta)}{f(X_i|\theta_0)} f_H(\theta) d\theta. \quad (3.51)$$

The version of mSPRT commonly used in digital experimentation is specified in [134], which adapted the test into a two-sample setting. It assumes the samples $X_1, \dots, X_n, \dots, X_{n'}$ and $Y_1, \dots, Y_n, \dots, Y_{n'}$ are normally distributed with known variances σ_X^2 and σ_Y^2 , respectively. The normality assumptions mean one can reduce the two samples into a single sample $(Y_1 - X_1), \dots, (Y_n - X_n), \dots, (Y_{n'} - X_{n'})$, observe the resultant sample follow a normal distribution with mean θ (to be inferred) and variance $\sigma_X^2 + \sigma_Y^2$,⁴³ and apply the one-sample mSPRT on the resultant sample. The authors of [134] also noted that by taking H as a normal distribution centred at the null hypothesis (i.e., $H = \mathcal{N}(\theta_0, \tau^2)$, where τ^2 is a hyperparameter that we specify or learn from data), the test statistic would have the following closed-form upon observing the first n samples:

$$\tilde{\Lambda}_n^{H,\theta_0} \triangleq \sqrt{\frac{\sigma_X^2 + \sigma_Y^2}{\sigma_X^2 + \sigma_Y^2 + n\tau^2}} \exp\left(\frac{n^2\tau^2(\bar{Y}_n - \bar{X}_n - \theta_0)^2}{2(\sigma_X^2 + \sigma_Y^2)(\sigma_X^2 + \sigma_Y^2 + n\tau^2)}\right), \quad (3.52)$$

where $\bar{X}_n = n^{-1} \sum_{i=1}^n X_i$ and $\bar{Y}_n = n^{-1} \sum_{j=1}^n Y_j$ represent the sample mean of the two samples up to sample n . The test rejects the null hypothesis if $\tilde{\Lambda}_n^{H,\theta_0} \geq \frac{1}{\alpha}$. However, it can also seek new responses indefinitely if the inequality is not met.

Group sequential tests Another popular sub-class of sequential tests is group sequential tests. Instead of specifying a new test statistic, a group sequential test reuses the statistical test that one would have used if the experiment is fixed-horizon multiple times during the test duration. Instead, it adjusts each interim test's significance level (and thus critical value) to ensure the overall Type I error rate is equal to what has been specified.⁴⁴

⁴³ In [134], the text states that the resultant sample has a variance of $2\sigma^2$. This arises from assuming $\sigma_X^2 = \sigma_Y^2 = \sigma^2$.

⁴⁴ The procedure resembles family-wise error rate control procedures such as Bonferroni correction, which we will cover in Section 5.4. That said, there are fundamental differences between the two sets of procedures – see Footnote 40.

The test procedure is then as follows. Many steps featured below are similar to that in an NHST, as described in Section 3.4:

1. Determine the statistical hypotheses, the appropriate test *for each interim analysis*, overall significance level, power, *maximum* sample size, and test statistic distribution under the null hypothesis H_0 (cf. Steps 1–4 in Section 3.4);
2. Upon collecting the required number of responses for the subsequent interim analysis:
 - (a) Obtain the adjusted significance level (see below for examples),
 - (b) Calculate the test statistic based on responses observed so far (cf. Step 5),
 - (c) Obtain the adjusted critical region or calculate the p -value (cf. Steps 6/6A),
 - (d) Decide to reject/not reject H_0 *within the context of the interim analysis*, using the adjusted critical region or significance level (cf. Steps 7/7A),
 - (e) If the interim analysis rejects H_0 , decide the same for the overall test;
Otherwise, continue to take observations;
3. Upon reaching the maximum sample size specified in Step 1, stop taking further observations and do not reject H_0 .

There are many ways to adjust the significance level in each interim test. Early approaches set a fixed discount on its value based on the number of *groups* (i.e., the number of times the experimenter intends to peek) and apply the discounted value uniformly to every interim test [220]. While experimenters can expect to use the same adjusted significance level for each interim test, such approaches require them to specify and commit to the number of times they intend to peek.

More recent approaches feature an alpha spending function specified before starting an experiment. It describes how the overall significance level α , which can be characterised as a budget of “chance to commit Type I error” in this context, is spent across multiple interim analyses [158]. Any monotonically increasing function that maps the proportion of responses observed so far to significance levels bounded above by α , i.e., $[0, 1] \rightarrow [0, \alpha]$, can be used as an alpha spending function with various effectiveness. While these approaches require non-trivial calculations to obtain the values of adjusted significance level and critical

region for each interim analysis,⁴⁵ they do not require experimenters to specify the number of times or when they intend to peek during the experiment [60].

3.9 Bayesian Tests

We also discuss statistical testing from a Bayesian perspective. Recall that we introduced $\mathbb{P}(H \mid \text{Data})$, the probability of a statistical hypothesis H being true given the data, in Section 3.6.1.⁴⁶ We observed that this probability is often considered the more intuitively correct heuristic when making data-driven decisions than, e.g., p -values [47]. In Section 3.6.1, “Data” was shorthand for the specific event “test statistic attaining the observed or more extreme values”. However, the observation still holds if we take the term’s general definition of “the probability of the samples taking their observed value”.

The Bayesian perspective enables one to obtain $\mathbb{P}(H \mid \text{Data})$ by treating it as a posterior belief. Under such an approach, one first specifies their prior belief on the hypothesis as a probability, denoted as $\mathbb{P}(H)$. They then update the prior belief with the observations, summarised by the likelihood under the chosen statistical model (denoted as $f(\text{Data} \mid H)$), using Bayes’ theorem to obtain the posterior belief.

We outline two approaches in the Bayesian setting. The first focuses on the posterior probability of a single hypothesis being true. By applying Bayes’ theorem directly:

$$\mathbb{P}(H \mid \text{Data}) = \frac{f(\text{Data} \mid H)\mathbb{P}(H)}{f(\text{Data})}. \quad (3.53)$$

In cases where the chosen statistical model is parametric, the hypothesis is usually a statement of plausible values of the parameter of interest θ . The likelihood function has to account for all such parameter values:

$$f(\text{Data} \mid H) = \int f(\text{Data} \mid \theta, H) f(\theta \mid H) d\theta. \quad (3.54)$$

⁴⁵ That said, software packages that perform such calculations are readily available. See, e.g., [37].

⁴⁶ The actual hypothesis featured in Section 3.6.1 was, in fact, the null hypothesis H_0 . The switch to a general statistical hypothesis H is intentional, as all hypotheses are treated equally in the Bayesian regime.

The second approach arises from the need to compare and select statistical models. Of course, we already have those models in hand – that specified by the null (H_0) and alternative hypotheses (H_1). Unlike the NHST regime, where there is a preferred/assumed state of matter and hence an unequal status balance between the two hypotheses, the models specified by the two hypotheses are treated equally in the Bayesian regime. Both are treated as candidate models, with updates done in the same way using the same observations.⁴⁷

The comparison and selection process utilises the Bayes factor [143]. We first substitute H_0 and H_1 into Equation (3.53) separately to obtain two posterior probability formulas. Then, we observe that the quotient between the two posterior probabilities is

$$\underbrace{\frac{\mathbb{P}(H_1 | \text{Data})}{\mathbb{P}(H_0 | \text{Data})}}_{\text{Posterior odds}} = \frac{f(\text{Data} | H_1)\mathbb{P}(H_1)}{f(\text{Data})} \cdot \frac{f(\text{Data})}{f(\text{Data} | H_0)\mathbb{P}(H_0)} = \underbrace{\frac{f(\text{Data} | H_1)}{f(\text{Data} | H_0)}}_{\text{Bayes factor } (\triangleq \text{BF}_{10})} \cdot \underbrace{\frac{\mathbb{P}(H_1)}{\mathbb{P}(H_0)}}_{\text{Prior odds}}, \quad (3.55)$$

with the expansion in Equation (3.54) applied where appropriate. The Bayes factor BF_{10} (the second-from-right fraction in Equation (3.55)) represents the strength of evidence provided by the observations in favour of H_1 . There is no universal scale with descriptors, though [143] considers a Bayes factor of $\sqrt{10} \approx 3.162$, 10, and 100 to be the lower thresholds for “substantial”, “strong”, and “decisive” evidence.

An advantage of the Bayes factor over NHST is that the former can also represent the strength of evidence *in favour* of H_0 , the “null” hypothesis that one can only accumulate evidence *against* in an NHST (Sections 3.4–3.6). We can obtain that by taking the reciprocal of the Bayes factor BF_{10} (or the entire Equation (3.55)) to obtain BF_{01} . One can also recover the posterior probability of each of the two hypotheses from the posterior odds quickly if H_0 and H_1 cover the entire event space. This is usually the case if we use the standard one- and two-sided hypotheses featured in the previous sections.

Applications in digital experimentation Bayesian testing has gained traction in digital experiments in recent years [61, 66]. It is often employed in lieu of sequential testing (see section above) to address the continuous monitoring/peeking problem – see Section 4.1 for a more detailed account of the problem and [66] for a discussion on how Bayesian and se-

⁴⁷ One can still favour a particular hypothesis via the specification of the prior belief.

quential testing differs in practice.

Many Bayesian tests in digital experimentation work off the statistical models used in a Wald/practical t -test. Consider a two-sample setting, with $X_1, \dots, X_n, \dots, X_{n'}$ and $Y_1, \dots, Y_m, \dots, Y_{m'}$ being i.i.d. samples from the two groups. We recall from Section 3.5.2 that given a large sample, the Wald/practical t -test statistic approximately follows the standard normal distribution. The Bayesian perspective allows one to update their belief with incoming observations. Thus, it makes sense to define a test statistic upon observing the first n samples from X and m samples from Y :

$$\sqrt{W_{n,m}} \triangleq \frac{\bar{Y}_m - \bar{X}_n - \Delta}{\sqrt{\left(\frac{s_X^2}{n} + \frac{s_Y^2}{m}\right)}} \stackrel{\text{approx.}}{\sim} \mathcal{N}(0, 1), \quad (3.56)$$

where $\bar{X}_n = \frac{1}{n} \sum_{i=1}^n X_i$ and $\bar{Y}_m = \frac{1}{m} \sum_{j=1}^m Y_j$ are the running sample means. s_X^2 and s_Y^2 are the plug-in estimates for the population variances σ_X^2 and σ_Y^2 , respectively.

Unlike the practical t -test, where one usually states hypotheses on the unstandardised effect size Δ , we prefer inferring the effect size standardised by the pooled variance in a Bayesian test. This is because standardising the effect size will make it scale-independent, easing the specification or learning of the prior. To obtain the standardised effect size, we note $\sqrt{W_{n,m}}$ can be expanded as

$$\sqrt{W_{n,m}} = \left(\underbrace{\frac{\bar{Y}_m - \bar{X}_n}{\sqrt{\left(\frac{s_X^2}{n} + \frac{s_Y^2}{m}\right) / \left(\frac{1}{n} + \frac{1}{m}\right)}}}_{\triangleq \delta_{n,m}} - \underbrace{\frac{\Delta}{\sqrt{\left(\frac{s_X^2}{n} + \frac{s_Y^2}{m}\right) / \left(\frac{1}{n} + \frac{1}{m}\right)}}}_{\triangleq \Delta^\circ} \right) \underbrace{\frac{1}{\sqrt{\frac{1}{n} + \frac{1}{m}}}}_{\triangleq \sqrt{E_{n,m}}}, \quad (3.57)$$

where $\delta_{n,m}$, Δ° , and $E_{n,m}$ are the test statistic, standardised effect size and the effective sample size of the test, respectively. Given that $\sqrt{W_{n,m}}$ approximates the standard normal distribution, $\delta_{n,m}$, a linear transformation of $\sqrt{W_{n,m}}$, approximates a (non-standard) normal distribution:

$$\delta_{n,m} = \frac{\bar{Y}_m - \bar{X}_n}{\sqrt{\left(\frac{s_X^2}{n} + \frac{s_Y^2}{m}\right) / \left(\frac{1}{n} + \frac{1}{m}\right)}} = \frac{\sqrt{W_{n,m}}}{\sqrt{E_{n,m}}} + \Delta^\circ \stackrel{\text{approx.}}{\sim} \mathcal{N}\left(\Delta^\circ, \frac{1}{E_{n,m}}\right). \quad (3.58)$$

This enables us to assume a normal likelihood for the standardised effect size.

Next, we specify the statistical hypothesis to calculate the Bayes factor. Similar to NHSTs, many hypothesis combinations exist (see Section 3.2).⁴⁸ A commonly used hypothesis pair is $H_0 : \Delta^\circ = \theta_0$ and $H_1 : \Delta^\circ \sim \mathcal{N}(\theta_0, V^2)$, where V^2 is a hyperparameter that we specify or learn from data. We observe that H_0 is identical to that used in NHSTs. Meanwhile, the prior specified in H_1 is a reasonable approximation to the distribution of (standardised) effect sizes (see Section 2.3.1) that also has nice mathematical properties – the resultant model under H_1 is a hierarchical normal model. Given the above, the Bayes factor is

$$(BF_{10})_{n,m} = \frac{f(\delta_{n,m} | H_1)}{f(\delta_{n,m} | H_0)} = \frac{\phi\left(\delta_{n,m}; \theta_0, V^2 + \frac{1}{E_{n,m}}\right)}{\phi\left(\delta_{n,m}; \theta_0, \frac{1}{E_{n,m}}\right)}, \quad (3.59)$$

where $\phi(\cdot; \mu, \sigma^2)$ is the PDF of a normal distribution with mean μ and variance σ^2 . One can also obtain the posterior odds shown in Equation (3.55) if they specify $\mathbb{P}(H_0)$ and $\mathbb{P}(H_1)$, the prior belief on the probability that each hypothesis is true.

Like every other Bayesian analysis, determining a good prior distribution for Δ° is difficult. If one prioritises removing experimenters' biases from prior specification, they may consider leveraging past experiment data to learn the prior using an empirical Bayes approach [61].⁴⁹ In our running example, this means using the expectation-maximisation algorithm to learn the parameter V^2 and the prior probability $\mathbb{P}(H_0)$. In the case where the speed in obtaining a large Bayes factor is essential, they may also consider using a *non-local prior* for H_1 , which is perhaps best described as drilling hole(s) on the probability density near the parameter(s) covered by H_0 [138]. By unnesting H_0 from H_1 , we avoid evidence in favour of H_0 to also contribute to H_1 as well and thus enable a quicker convergence had H_0 been the ground truth.

⁴⁸ That said, one should take extra care when specifying the statistical hypotheses in a Bayesian test, particularly the alternative hypothesis. In NHSTs, the alternative hypothesis is generally not involved in test calculations and can be lazily specified as the complement of the null hypothesis. This is different for Bayes factor calculations, which involve both hypotheses heavily.

⁴⁹ In [61, 66], both prominent works on Bayesian testing in digital experimentation, the authors called the prior distribution obtained from the empirical Bayes "objective". We believe this is used in the general sense to contrast with a *subjective prior* elicited from a domain expert. In Bayesian analysis, *objective priors* usually mean *non-informative priors*, default priors constructed by some formal or structural rules that are largely ignorant to the context surrounding the experiment [50, 144], which is a different approach to empirical Bayes priors [265].

3.10 A Brief Recap

We provided an introduction to statistical testing in digital experimentation in this chapter. Starting from the beginning of a statistical test – the specification of one or more statistical hypotheses – we introduced four (overlapping) classes of statistical tests. They are NHSTs, non-parametric tests, sequential tests, and Bayesian tests. For each class of tests, we described one or more examples that are popular among digital experimenters in detail. We also outlined three alternatives to formal statistical tests: the inter-ocular trauma test, effect size, and confidence intervals.

We have aimed to balance the theoretical foundations and the practical applications. For the former, we covered, among many topics, the distributional results that underpin the practical t -tests and the derivation of the required sample size formulas from first principles using the concept of power. For the latter, we have outlined the situations where a specific test is applied in digital experiments and the pitfalls in interpreting quantities (e.g., p -values) featured in a test. In addition, we have carefully collated and linked the terminologies used in different communities. This is done with the intention for the guide to be accessible to a diverse yet mathematically-inclined audience.

That said, the expansive nature of statistical testing means we have inevitably omitted the details of many examples and (primarily theoretical) concepts. We provided further details and pointers for interested readers in the footnotes and references throughout the chapter.

Chapter 4

Datasets for Digital Experiments

The chapter is adapted from the research paper “*Datasets for Online Controlled Experiments*”, presented at the *35th Conference on Neural Information Processing Systems (NeurIPS '21)* [173].

4.1 Motivation

We move on and discuss data, another necessary ingredient for running digital experiments. Data plays a central role in developing the digital economy, especially in digital content, social networks, digital advertising and e-commerce. Its generation, processing, use, and storage have been the topic of interest for many software engineering, databases, and machine learning researchers and practitioners. That said, we feel data is less rigorously explored in the digital experimentation context than statistical tests (see Chapter 3). Its existence is often taken for granted by researchers or practitioners in digital experimentation or treated as merely one of the means to tackle the research/business problem (instead of the object of interest).

To illustrate, the ability to run experiments on the Web allows one to interact with many subjects within a short time frame and collect many responses. This, together with the scale of experimentation carried out by tech organisations, should lead to a wealth of datasets describing the result of an experiment. However, there are not many publicly available datasets describing digital experiments, and we believe they were never systematically reviewed nor

categorised. This is in contrast to the machine learning field, which also enjoyed its application boom in the past decade yet already has established archives and detailed categorisations for its datasets [78, 266].

We argue that the lack of relevant datasets arising from real experiments hinders the further development of digital experimentation and measurement methods (e.g., new statistical tests, bias correction, and variance reduction methods; see Chapter 5). Many statistical tests proposed relied on simulated data that impose restrictive distributional assumptions and thus may not represent the real-world scenario. Moreover, it may be difficult to understand how methods differ and assess their relative strengths and weaknesses without a common dataset to compare them on.

To address this problem, we present the first ever survey for online controlled experiments (OCE) datasets, together with a taxonomy that can be generalised to all digital experiments.¹ Our survey identified 13 datasets, including standalone experiment archives, accompanying datasets from scholarly works, and demo datasets from online courses on the design and analysis of experiments. We also categorise these datasets based on dimensions such as the number of experiments each dataset contains, how granular each data point is time-wise and subject-wise, and whether it includes results from real experiment(s).

The taxonomy enables us to discuss the data requirements for an experiment by systematically mapping out which data dimension is required for which statistical test and learning the hyperparameter(s) associated with the test. We also recognise that, in practice, data are often used for purposes beyond what it is originally collected for [146]. Hence, we posit the mapping is equally useful in allowing one to understand their options when choosing statistical tests given the format of data they possess. Together with the survey, the taxonomy helps us to identify what types of datasets are required for commonly used statistical tests yet are missing from the public domain.

One of the gaps identified by the survey and taxonomy is datasets that can support the design and running of experiments with *adaptive stopping* (a.k.a. continuous monitoring / optional stopping). We motivate their use below. Traditionally, experimenters analyse experiments

¹ Online controlled experiments, or online randomised controlled trials, are a subset of digital experiments. They are characterised by the experimenters' ability to perform random assignments and apply scientific control. See Section 5.3 for further details.

using null hypothesis significance tests (NHSTs), e.g., a Student's t -test. These tests require one to calculate and commit to a required sample size based on some expected treatment effect size, all prior to starting the experiment. Making extra decisions during the experiment, be it stopping the experiment early due to seeing favourable results or extending the experiment as it teeters "on the edge of statistical significance" [203], is discouraged as they risk one having more false discoveries than intended [109, 202].

Clearly, the restrictions above are incompatible with modern decision-making processes. Businesses operating online are incentivised to deploy beneficial features and roll back damaging changes as quickly as possible. In the "free delivery" banner example in Chapter 1, the business may have calculated that they require four weeks to observe enough users based on an expected 1% change in the decision metric. If the experiment shows, two weeks in, that the banner is leading to a 2% improvement, it will be unwise not to deploy the banner to all users simply due to the need to run the experiment for another two weeks. Likewise, if the banner is shown leading to a 2% loss, it makes every sense to immediately terminate the experiment and roll back the banner to stem further losses.

As a result, more experimenters are moving away from NHSTs and adopting adaptive stopping techniques. Experiments with adaptive stopping allow one to decide when to stop an experiment (i.e. stopping it earlier or prolonging it) based on the sample responses observed so far without compromising the statistical validity of false positive/discovery rate control. To encourage further development in this area, both in methods and data, we release the ASOS Digital Experiments Dataset, which contains daily checkpoints of decision metrics from multiple real OCEs run on the global online fashion retail platform.

The dataset design is guided by the requirements identified by the mapping between the taxonomy and statistical tests, and to the best of our knowledge, is the first public dataset that can support the end-to-end design and running of digital experiments with adaptive stopping. We demonstrate it can indeed do so by (1) running a sequential test and a Bayesian hypothesis test on all the experiments in the dataset and (2) estimating the value of hyperparameters associated with the tests. While the notion of ground truth does not exist in real digital experiments, we show that the dataset can also act as a quasi-benchmark for statistical tests by comparing results from the tests above with that of a t -test.

To summarise, our contributions are:

1. (Sections 4.2 & 4.3) We create, to the best of our knowledge, the first ever taxonomy on digital experiment datasets and apply it to publicly available online controlled experiment datasets;
2. (Section 4.4) We map the relationship between the taxonomy and statistical tests commonly used in experiments by identifying the minimally sufficient set of statistics and dimensions required in each test. The mapping, which also applies to offline experiments, enables experimenters to quickly identify the data collection requirements for their experiment design (and conversely, the test options available given the data availability); and
3. (Section 4.5) We make available, to the best of our knowledge, the first real, multi-experiment time series dataset, enabling the design and running of experimentation with adaptive stopping.²

4.2 A Taxonomy for Digital Experiment Datasets

We begin by presenting a taxonomy on digital experiment datasets, which is necessary to characterise and understand the results of a survey. To the best of our knowledge, there are no surveys nor taxonomies specifically on this topic prior to this work. While there is a large volume of work concerning the categorisation of datasets in machine learning [78, 266], research work in the online randomised controlled experiment methods [10, 12, 90, 233], and general experiment design [128, 176], our search on Google Scholar and Semantic Scholar using combinations of the keywords “online controlled experiment”/“A/B test”, “dataset”, and “taxonomy”/“categorization”/“categorisation” yields no relevant results.

The taxonomy focuses on the following four main dimensions:

Experiment count A dataset can contain the data collected from a *single* or *multiple* experiments. Results from a single experiment are useful for demonstrating how a test works, though any learning should ideally involve multiple experiments. Two closely related but relatively minor dimensions are the *variant count* (number of control/treatment groups in the

² Link to the dataset and accompanying datasheet: <https://osf.io/64jsb/>.

experiment) and the *metric count* (number of performance metrics the experiment is tracking). Having an experiment with multiple variants and metrics enables the demonstration of methods such as false discovery rate control procedures (e.g., [25], see Section 5.4.1 for more details) and learning the correlation structure within an experiment.

Response granularity Depending on the experiment analysis requirements and constraints imposed by the digital experimentation platform, the dataset may contain data aggregated to various levels. Consider the “free delivery” banner example in Chapter 1, where the website users are randomly allocated to the treatment (showing the banner) and control (not showing the banner) groups to understand whether the banner changes the proportion of users who bought something. Each individual user is considered a randomisation unit [155].

A dataset may contain, for each experiment, only summary statistics on the *group* level, e.g., the proportion of users who have bought something in the control and treatment groups, respectively. It can also record one response per *randomisation unit*, with each row containing the user ID and whether the user bought something. The more detailed activity logs at a *sub-randomisation unit* level can also have each row containing information about a particular page view from a particular user.

Time granularity An experiment can last between a week and many months [155], providing many possibilities for recording the result. A dataset can opt to record the *overall result only*, showing the end state of an experiment. If the dataset contains multiple randomisation units or experiments, it may or may not have a *timestamp* for each instance. It can also record *intermediate checkpoints* for the decision metrics, ideally at regular intervals such as *daily* or *hourly*. These checkpoints can either be a *snapshot* of the interval (recording activities between time t and $t + 1$, time $t + 1$ and $t + 2$, etc.) or *cumulative* from the start of the experiment (recording activities between time 0 and 1, time 0 and 2, etc.).

Syntheticity A dataset can record data generated from a *real* process. It can also be *synthetic* – generated via simulations with distributional assumptions applied. A dataset can also be *semi-synthetic* if generated from a real-life process and subsequently augmented with synthetic data.

Note that we can also describe datasets arising from any experiments (including offline and non-randomised controlled experiments) using these four dimensions. We will discuss in Section 4.4 how these dimensions map to common statistical tests used in digital experimentation.

In addition, we also record the *application domain*, *target demographics*, and the *temporal coverage* of the experiment(s) featured in a dataset. In an age when data are often reused, one must understand the underlying context and that learnings from a dataset created under a specific context may not translate to another context. We also see the surfacing of such context as a way to promote considerations in fairness and transparency for experimenters as more experiment datasets become available [133]. For example, having target demographics information on a meta-level helps experimenters identify who were involved, or perhaps more importantly, who were *not* involved in experiments and could be adversely impacted by an effectively untested treatment.

Finally, the datasets can also differ in their medium of *documentation* and the presence/absence of a *data management / long-term preservation plan*. The latter includes the hosting location, the presence/absence of a DOI, and the type of license. We record these attributes for the datasets surveyed below for completeness.

4.3 Public Online Controlled Experiment Datasets

Here, we discuss our approach to produce the first ever survey on OCE datasets and present its results. The survey is compiled via two search directions, which we describe below. For both directions, we conduct a first round search in May 2021, with follow up rounds in August and October 2021 to ensure we have the most updated results.

We first searched on the vanilla Google search engine using the keywords “Online controlled experiment “dataset””, “A/B test “dataset””, and “Multivariate test “dataset””. For each keyword, we inspect the first 10 pages of the search result (top 100 results) for scholarly articles, web pages, blog posts, and documents that may host or describe a publicly available OCE dataset (or both). The search term “dataset” is in double quotes to limit the search results to those explicitly mentioning dataset(s). We also searched on specialist data search en-

gines/hosts, namely Google Dataset Search (GDS) and Kaggle, using the keywords “Online controlled experiment(s)” and “A/B test(s)”. We inspect the metadata and description for all the results returned (except for GDS, where we inspect the first 100 results for “A/B test(s)”) for relevant datasets as defined below.³

A dataset must record the result arising from a randomised controlled experiment run online to be included in the survey. The criterion excludes experimental data collected from offline experiments, e.g. those in agriculture [284], medicine [77], and economics [81]. It also excludes datasets used to perform quasi-experiments and observational studies, e.g. the LaLonde dataset used in econometrics [59] and datasets constructed for uplift modelling tasks [72, 124].⁴

The result is presented in Table 4.1. We place the 13 OCE datasets identified in this exercise along the four taxonomy dimensions defined in Section 4.2 and record the additional features. These datasets include two standalone archives, one for online media experiments [185] and the other for experiments on an online education platform [244], plus two accompanying datasets for peer-reviewed research articles [262, 293]. There are also tens of Kaggle datasets, blog posts, and code repositories that describe or duplicate one of the six example datasets used in five massive open online courses on online controlled experiment design and analysis [15, 33, 111, 113, 264]. Finally, we identify three standalone datasets hosted on Kaggle with relatively light documentation [13, 88, 147].

The table shows several gaps in OCE dataset availability, the most obvious being the lack of datasets that record responses at a sub-randomisation unit level. In the sections below, we will identify more of these gaps and discuss their implications for digital experiment analysis.

³ Searching for the keyword “Online controlled experiment” on GDS and Kaggle returned 42 and 7 results, respectively, and that for “A/B test” returned “100+” and 286 results, respectively. Curiously, replacing “experiment” and “test” in the keywords with their plural form changes the number of results, with the former returning 6 and 10 results on GDS and Kaggle, respectively, and the latter returning “100+” and 303 results, respectively.

⁴ Uplift modelling (UM) tasks for online applications often start with an OCE [226], and thus we can consider UM datasets as OCE datasets with extra randomisation unit level features. The nature of the tasks is different, though: OCEs are concerned with validating the average treatment effect across the population using a statistical test, whereas UM is concerned with modelling the conditional average treatment effect for each individual user, making it more a general causal inference task that is outside the scope of this survey.

Dataset Name	Ref.	Experiment Count - Variant / Metric Count	Response Granularity	Time Granularity	Syntheticity	Application Domain	Target Demographic	Temporal Coverage	Documentation	Data management / long- term preservation plan
Upworthy Research Archive	[185]	Multiple (32 487) - 2-14 / 2 (C)	Group	Overall result only Timestamp per expt.	Real	Media & Ads Copy/Creative	Mostly English- speaking users in USA	Jan 2013 – Apr 2015	Peer reviewed data article	Host: Open Science Framework DOI: ✓ (see [185]) Licence: CC BY 4.0
ASSISTments Dataset from Multiple Randomized Controlled Experiments	[244]	Multiple (22) - 2 / 2 (BC)	Rand. Unit	Overall result only ✗ timestamp	Real	Education Teaching model	Mostly middle school students (age 11-14) in/near MA, USA	2013 – 2015	Peer reviewed data article	Host: Author website DOI: Unknown Licence: Unknown
A/B Testing Web Analytics Data (From [292])	[293]	Single - 5 / 2 (C)	Group	Overall result only ✓ timestamp	Real	Education UX/UI change	Mostly English- speaking university library users	May 2013 – Jun 2013	Accompanying dataset to peer-reviewed research article	Host: University library DOI: ✓ (see [293]) Licence: CC BY-SA 4.0
Dataset of two experiments of the application of gamified peer assessment model into online learning environment MeuTutor (From [261])	[262]	Multiple (2) - 3+2 / 3+3 (R+C)	Rand. Unit	Overall result only ✗ timestamp	Real	Education Teaching model	High school students in Brazil taking ENEM (age 17)	Jul 2015 – Aug 2015	Peer reviewed data article	Host: Journal website DOI: ✓ (see [262]) Licence: CC BY 4.0
Udacity Free Trial Screener Experiment (From Udacity A/B Testing Course - Final Project [111])	See e.g. [205, 247, 273]	Single - 2 / 4 (C)	Group	Daily checkpoint Snapshot	Real	Education UX/UI change	Mostly English- speaking users	Unknown	Blog posts & Kaggle notebooks e.g. [192, 236, 294]	Host: Kaggle / GitHub (multiple) DOI: Unknown Licence: Unknown
“Analyse A/B Test Results” Dataset (From Udacity Online Data Analyst Course - Project 3 [264])	See e.g. [4, 57, 215]	Single - 2 / 1 (B)	Rand. Unit	Overall result only Timestamp per RU	Unknown	E-commerce UX/UI change	Unknown	Unknown	Blog posts & Kaggle notebooks e.g. [42, 180, 227]	Host: Kaggle / GitHub (multiple) DOI: Unknown Licence: Unknown
Mobile Games A/B Testing with Cookie Cats (DataCamp project [15])	See e.g. [82, 257, 291]	Single - 2 / 3 (BC)	Rand. Unit	Overall result only ✗ timestamp	Real	Gaming Design change	Unknown (likely Facebook users)	Unknown	Kaggle notebook	Host: Kaggle / Course website DOI: Unknown Licence: Unknown
Experiment Dataset (From DataCamp A/B Testing in R Course [33])	[35]	Single - 2 / 1 (B)	Rand. Unit	Overall result only Timestamp per RU	Synthetic	Tech UX/UI Change	N/A	N/A	Course notes Blog posts	Host: Course website DOI: Unknown Licence: Unknown
Data Visualization Website - April 2018 (From DataCamp A/B Testing in R Course [33])	[34]	Single - 2 / 4 (BR)	Rand. Unit	Overall result only Timestamp per RU	Synthetic	Tech UX/UI Change	N/A	N/A	Course notes	Host: Course website DOI: Unknown Licence: Unknown
AB Testing Result (From Customer Analytics and A/B Testing in Python [113])	[112]	Single - 2 / 2 (CR)	Rand. Unit	Overall result only Timestamp per RU	Unknown	Tech UX/UI Change	Unknown	Jan 2014 – Jan 2018	Course notes	Host: Course website DOI: Unknown Licence: Unknown
Grocery Website Data for AB Test	[147]	Single - 2 / 1 (B)	Rand. Unit	Overall result only ✗ timestamp	Unknown	E-commerce UX/UI change	Unknown	Unknown	Kaggle notebook	Host: Kaggle DOI: Unknown Licence: Unknown
Ad A/B Testing (aka SmartAd AB Data)	[88]	Single - 2 / 2 (B)	Rand. Unit	Overall result only Timestamp per RU	Unknown	Media & Ads Display ads	Unknown	Jul 2020	Kaggle notebook	Host: Kaggle DOI: N/A Licence: CC BY-SA 3.0
Synthetical A/B-Tests	[13]	Multiple (25 856) - 2 / 1 (R)	Group.	Overall result only ✗ timestamp	Synthetic	N/A	N/A	N/A	Kaggle notebook	Host: Kaggle DOI: Unknown Licence: CDLA-Sharing-1.0

Table 4.1: Results from the first ever survey of OCE datasets. The 13 datasets identified are placed on the four taxonomy dimensions defined in Section 4.2, together with the additional attributes recorded. In the Experiment Count column, a second-line value (corresponding to Variant/Metric Count) of “ x / y (BCR)” means the dataset features x variants and y metrics in each experiment, with the metrics based on **B**inary, **C**ount, and **R**eal-valued responses. In the Response Granularity and Time Granularity columns, Randomisation Unit is abbreviated as RU or Rand. Unit. Many resources are accessed via links to non-scholarly articles – blog plots, Kaggle dataset pages, and GitHub repositories. These resources may not persist over time.

4.4 Matching Dataset Taxonomy with Statistical Tests

Specifying the data requirements (or structure) and performing statistical tests are perhaps two of the most common tasks carried out by data scientists. However, the link between the two processes is seldom mapped out explicitly. It is all too common to consider from scratch the question, “I need to run this statistical test, how should I format my dataset?” (or more controversially, “I have this dataset, what statistical tests can I run?” [146]) for every new project/application, despite the list of possible dataset dimensions and statistical tests remaining essentially the same.

We aim to speed up the process above by describing what summary statistics are required to perform common statistical tests in digital experiments and link the statistics back to the taxonomy dimensions defined in Section 4.2. The exercise is similar to identifying the sufficient statistic(s) for a statistical model [95]. However, the identification is done for the encapsulating statistical inference procedure, with a practical focus on data dimension requirements. We do so by stating the formula used to calculate the corresponding effect sizes and test statistics and observing the summary statistics required in common. The general approach also enables one to apply the resultant mapping to any experiments that involve a two-sample statistical test, including offline experiments and experiments without a randomised control. We will refrain from discussing the full model assumptions and their applicability for brevity. Instead, we point readers to the relevant work in the literature.

4.4.1 Effect size and Welch’s t -test

We first look at the requirements to calculate effect sizes (using Cohen’s d) and perform the most basic statistical test (using Welch’s t -statistic). Recall from Sections 3.3.2 and 3.5.2 that the formulas to calculate these two quantities are

$$d = \frac{\bar{Y} - \bar{X}}{\sqrt{\frac{(n-1)s_X^2 + (m-1)s_Y^2}{n+m-2}}}, \quad (*3.2)$$

$$T = \frac{\bar{Y} - \bar{X} - \Delta}{\sqrt{\frac{s_X^2}{n} + \frac{s_Y^2}{m}}}, \quad (*3.12)$$

where $(\bar{X}, \bar{Y})$, (s_X^2, s_Y^2) , and (n, m) are the means, variances, and counts of the samples X and Y , respectively.

We observe that we require only six *group-level* summary statistics to calculate the two quantities above: two means, two variances, and two counts. Besides the counts, the exact list of quantities required depends on the distributional assumptions used in a test, which dictates what the sufficient statistics are.⁵ We call these quantities *Dimension Zero (D0)* quantities, as they are the bare minimum required to run a statistical test. These quantities will be expanded along the taxonomy dimensions defined in Section 4.2.

Cluster randomisation / dependent data The sample variance estimates (s_X^2, s_Y^2) may be biased when cluster randomisation is involved. Using the “free delivery” banner example again, instead of randomly assigning each individual user to the control and treatment groups, the business may randomly assign postcodes to the two groups, with all users from the same postcode getting the same version of the website. In this case, user responses may become correlated, violating the independence assumptions in statistical tests. Common workarounds, including the use of bootstrap [19] and the Delta method [64] (see Sections 5.4.1 and 5.5), generally require access to *sub-randomisation unit* responses.

4.4.2 Experiments with adaptive stopping

As discussed in Section 4.1, experiments with adaptive stopping are getting increasingly popular among the digital experimentation community. Here, we motivate the data requirement for statistical tests in this domain by looking at the quantities required to calculate the test statistics for a mixture sequential probability ratio test (mSPRT, see Section 3.8 or [134]) and a Bayesian hypothesis test using Bayes factor (see Section 3.9 or [66]), two popular approaches in digital experimentation. Many other tests support adaptive stopping [194, 271], though the data requirements should be largely identical in terms of the dimensions defined in Section 4.2.

⁵ Welch’s t -test assumes normally distributed data and hence requires the sample mean and variance for both samples as they are a joint sufficient statistic for normal distributions with unknown variance. For decision metrics derived from binary responses, we only require the sample mean as it is a sufficient statistic for a Bernoulli distribution.

We first recall from Section 3.8 that running a mSPRT with a normal mixing distribution $H = \mathcal{N}(\theta_0, \tau^2)$ involves calculating the following test statistic upon observing the first n samples from X and Y :

$$\tilde{\Lambda}_n^{H, \theta_0} = \sqrt{\frac{\sigma_X^2 + \sigma_Y^2}{\sigma_X^2 + \sigma_Y^2 + n\tau^2}} \exp\left(\frac{n^2\tau^2(\bar{Y}_n - \bar{X}_n - \theta_0)^2}{2(\sigma_X^2 + \sigma_Y^2)(\sigma_X^2 + \sigma_Y^2 + n\tau^2)}\right), \quad (*3.52)$$

where $\bar{X}_n = n^{-1} \sum_{i=1}^n X_i$ and $\bar{Y}_n = n^{-1} \sum_{j=1}^n Y_j$ represent the sample mean of X and Y up to sample n , respectively.

For Bayesian hypothesis tests, we recall from Section 3.9 that the test statistic $\delta_{n,m}$ and effective sample size $E_{n,m}$ can be obtained by decomposing the Wald test statistic upon observing the first n samples from X and m samples from Y :

$$\sqrt{W_{n,m}} = \left(\underbrace{\frac{\bar{Y}_m - \bar{X}_n}{\sqrt{(\frac{s_X^2}{n} + \frac{s_Y^2}{m}) / (\frac{1}{n} + \frac{1}{m})}}}_{\triangleq \delta_{n,m}} - \underbrace{\frac{\Delta}{\sqrt{(\frac{s_X^2}{n} + \frac{s_Y^2}{m}) / (\frac{1}{n} + \frac{1}{m})}}}_{\triangleq \Delta^\circ} \right) \frac{1}{\underbrace{\sqrt{\frac{1}{n} + \frac{1}{m}}}_{\triangleq \sqrt{E_{n,m}}}}, \quad (*3.57)$$

where Δ° is the standardised effect size. $\delta_{n,m}$ and $E_{n,m}$ are then used to calculate the Bayes factor under the hypothesis pair $H_0 : \Delta^\circ = \theta_0$ and $H_1 : \Delta^\circ \sim \mathcal{N}(\theta_0, V^2)$:

$$(BF_{10})_{n,m} = \frac{f(\delta_{n,m} | H_1)}{f(\delta_{n,m} | H_0)} = \frac{\phi\left(\delta_{n,m}; \theta_0, V^2 + \frac{1}{E_{n,m}}\right)}{\phi\left(\delta_{n,m}; \theta_0, \frac{1}{E_{n,m}}\right)}, \quad (*3.59)$$

where $\phi(\cdot; \mu, \sigma^2)$ is the PDF of a normal distribution with mean μ and variance σ^2 .

During an experiment with adaptive stopping, we calculate the test statistics stated above many times for different n and m . This means a dataset can only support the running of such experiments if it contains *intermediate checkpoints* for the counts (n, m) and the means $(\bar{X}_n, \bar{Y}_m)$, ideally *cumulative* from the start of the experiment. One often also requires the variances at the same time points (see below). The only exception to the dimensional requirement above is where the dataset contains responses at a *randomisation unit* or finer level of granularity and, despite recording the *overall results only*, has a *timestamp per randomisation unit*. Under this special case, we can still construct the cumulative means $(\bar{X}_n, \bar{Y}_m)$ for all relevant values of n and m by ordering the randomisation units by their associated timestamps.

Learning the effect size distribution (hyper)parameters The two tests introduced above feature some hyperparameters (τ^2 and V^2) that must be specified or learned from data. These parameters characterise the prior belief of the effect size distribution, which will be the most effective if it “matches the distribution of true effects across the experiments a user runs” [134]. Common parameter estimation procedures [1, 14, 115] require results from *multiple related experiments*.

Estimating the response variance In the equations above, the response variance of the two samples, σ_X^2 and σ_Y^2 (or s_X^2 and s_Y^2), are assumed to be known before we observe all the samples. In practice, we often use the plug-in empirical estimates $(s_X^2)_n$ and $(s_Y^2)_m$ – the sample variances for the first n samples from X and m samples from Y , respectively. Thus, the data dimensional requirement is identical to that of the counts and means as discussed above. We will also require a *sub-randomisation unit* response granularity when the plug-in estimates are biased due to dependent data (see Section 4.4.1).

4.4.3 Non-parametric tests

We also briefly discuss the data requirements for non-parametric tests. Recall from Section 3.7 that one of the most commonly used non-parametric tests in digital experimentation, the Mann-Whitney U -test, calculates the following test statistic:

$$U = \sum_{i=1}^n \sum_{j=1}^m S(X_i, Y_j), \quad \text{where } S(X, Y) = \begin{cases} 1 & \text{if } Y < X, \\ 1/2 & \text{if } Y = X, \\ 0 & \text{if } Y > X. \end{cases} \quad (*3.46)$$

While a rank-based method is available for large n and m , both methods require knowledge of all the X_i and Y_j . This requirement applies to many other non-parametric tests (see Section 3.7). The observation suggests that a dataset can only support a non-parametric test if it at least provides responses at a *randomisation unit* level.

We conclude by showing how we can combine the individual data requirements above to obtain the requirement to design or run experiments for more complicated statistical tests.

This is possible due to the orthogonal design of the taxonomy dimensions. Consider an experiment with adaptive stopping using Bayesian non-parametric tests (e.g., with a Pólya Tree prior [43, 127]). It involves a non-parametric test, requiring responses at a *randomisation unit* level. It computes multiple Bayes factors for adaptive stopping and hence requires *intermediate checkpoints* for the responses (or a timestamp for each randomisation unit). Finally, it needs to learn the hyperparameters of the Pólya Tree prior, which requires *multiple related experiments*. The substantial data requirement along 3+ dimensions perhaps explains the lack of relevant digital experiment datasets and the tendency for experimenters to use simpler statistical tests for the day-to-day design or running of digital experiments (or both).

4.5 A Novel Dataset for Experiments with Adaptive Stopping

We finally introduce the ASOS Digital Experiments Dataset, which we believe is the first public dataset that supports the end-to-end design and running of digital experiments with adaptive stopping. We motivate why this is the case, provide a light description of the dataset (and a link to the more detailed accompanying datasheet),² and showcase the capabilities of the dataset via a series of experiments. We also discuss the ethical implications of releasing this dataset.

Recall from Section 4.4.2 that in order to support the end-to-end design and running of experiments with adaptive stopping, we require a dataset that

1. Includes *multiple* related experiments;
2. Is *real*, so that any parameters learned are reflective of the real-world scenario;

and either

- 3a. Contains *intermediate checkpoints* for the summary statistics during each experiment (i.e., time-granular), or
- 3b. Contains responses at a *randomisation unit* granularity with a timestamp for each randomisation unit (i.e. response-granular with timestamps).

None of the datasets surveyed in Section 4.3 meet all three criteria. While the Upworthy [185], ASSISTments [244], and MeuTutor [262] datasets meet the first two criteria, they all fail to

meet the third.⁶ The Udacity Free Trial Screener Experiment dataset meets the last two criteria by having results from a real experiment with daily snapshots of the decision metrics (and hence time-granular), which supports the running of an experiment with adaptive stopping. However, the dataset only contains a single experiment, which does not help in learning the effect size distribution (the design).

The ASOS Digital Experiments Dataset contains results from OCEs run by a business unit within ASOS.com, a global online fashion retail company. In terms of the taxonomy defined in Section 4.2, the dataset contains *multiple* (78) *real* experiments, with two to five variants in each experiment and four decision metrics based on binary, count, and real-valued responses. The results are aggregated on a *group* level, with *daily or 12-hourly checkpoints* of the metric values *cumulative* from the start of the experiment. The dataset design meets all the three criteria stated above and thus differentiates itself from other public datasets.

We provide readers with an accompanying datasheet (based on [99]) that provides further information about the dataset. We also host the dataset on Open Science Framework to ensure it is easily discoverable and can be preserved long-term.² It is worth noting that the dataset is released with the intent to support development in the statistical methods required to run digital experiments. The experiment results shown in the dataset are not representative of ASOS.com’s overall business operations, product development, or experimentation program operations, and no conclusion of such should be drawn from this dataset.

4.5.1 Potential use cases

Meta-analyses The multi-experiment nature of the dataset enables one to perform meta-analyses. A simple example is characterising the distribution of p -values (under Welch’s t -test) across all experiments (see Figure 4.1). We observe that roughly a quarter of experiments in this dataset attain $p < 0.05$ and attribute this to the fact that what we experiment in digital experiments is often guided by what domain experts think may have an impact. That said, we invite external validation on whether there is evidence for data dredging using,

⁶ All three report the overall results only and hence are not time-granular. The Upworthy dataset reports group-level statistics and hence is not response-granular. The ASSISTments and MeuTutor datasets are response-granular but lack the timestamp to order the samples.




Figure 4.1: Distribution of p -values attained by the 99 OCEs in the ASOS Digital Experiment Dataset using Welch’s t -tests, split by decision metrics. The leftmost bar in each histogram represents experiments with $p < 0.05$. Here, we treat OCEs with multiple variants as multiple independent OCEs.

e.g., [193].

Design and running of experiments with adaptive stopping We also demonstrate that the dataset can indeed support digital experiments with adaptive stopping by performing a mixed sequential probability ratio test (mSPRT) and a Bayesian hypothesis test via the Bayes factor for each experiment and metric. This requires learning the hyperparameters τ^2 and V^2 . We learn, for each metric, a naïve estimate for V^2 by collating the $\delta_{n,m}$ (see Equation (3.57)) at the end of each experiment and taking their sample variance. This yields the estimates $1.30\text{e-}05$, $1.07\text{e-}05$, $6.49\text{e-}06$, and $5.93\text{e-}06$ for the four metrics respectively.

For τ^2 , we learn near-identical naïve estimates by collating the value of Cohen’s d (see Equation (3.2)) instead. However, as τ^2 captures the spread of unstandardised effect sizes, we specify in each test $\tau^2 = d \cdot (s_X^2)_n$, where $(s_X^2)_n$ is the sample variance of all responses up to the n^{th} observation in that particular experiment. The Bayesian tests also require a prior belief in the null hypothesis being true ($\mathbb{P}(H_0)$) – we set it to 0.75 based on what we observed in the t -tests above.

We then calculate the p -value in mSPRT and the posterior belief in the null hypothesis ($\mathbb{P}(H_0|\text{data})$) in the Bayesian test for each experiment and metric at each daily/12-hourly checkpoint, following the procedures stated in [134] and [66, 143], respectively. We plot the results for five experiments selected at random in Figure 4.2. It shows that the p -value for an mSPRT is monotonically non-increasing. At the same time, the posterior belief for a Bayesian test can fluctuate depending on the effect size observed so far.




Figure 4.2: Change in (left) p -values in a mixed sequential probability ratio test ($\tau^2 = 5.92\text{e-}06$) and (right) posterior belief in the null hypothesis ($\mathbb{P}(H_0|\text{data})$) in a Bayesian hypothesis test ($V^2 = 5.93\text{e-}06$, $\mathbb{P}(H_0) = 0.75$) during the experiment for five experiments selected at random. The experiment duration (x-axis) is normalised by its overall runtime. Only results for Metric 4 are shown.

A quasi-benchmark for adaptive stopping methods Real digital experiments, unlike machine learning tasks, generally do not have a notion of ground truth. Nonetheless, the dataset can still be used as a quasi-ground truth to compare two hyperparameter settings of the same adaptive stopping method or between two adaptive stopping methods. For example, we can treat the reject / not reject verdict from a Welch’s t -test at the end of an experiment as the “ground truth” and compare it to the reject / not reject verdict of an mSPRT at different experiment stages. This yields many “confusion matrices” over different stages of individual experiments, where a “Type I error” corresponds to cases in which a Welch’s t -test fails to reject the null hypothesis and an mSPRT rejects the null hypothesis. A confusion matrix for the end of each experiment can be seen in Table 4.2. As the dataset was collected without early stopping, it allows us to perform sensitivity analysis and optimisation on the hyperparameters of mSPRT under what can be construed as a “precision-recall” trade-off of statistically significant treatments.

Other use cases The time series nature of this dataset enables one to detect bias (of the estimator) across time, e.g., those caused by concept drift or feedback loops. In the context of digital experiments, [41] described several methods to detect invalid experiments over time that may be run on this dataset. Moreover, with the time series spanning multiple experiments and decision metrics, one can also learn the correlation structure across experiments, decision metrics, and time [221, 274].

Table 4.2: Comparing the number of reject / not reject verdicts given by Welch’s t -test and an mSPRT at the end of an experiment for all four metrics.

t -test	Rejects H_0		Does not reject H_0	
mSPRT	Rejects H_0	Does not reject H_0	Rejects H_0	Does not reject H_0
Metric 1	19	7	2	71
Metric 2	20	7	18	49
Metric 3	16	9	6	63
Metric 4	16	11	4	63

4.5.2 Ethical considerations

We finally discuss the ethical implications of releasing the dataset, touching on data protection and anonymisation, potential misuse, and the ethical considerations for running digital experiments in general.

Data protection and anonymisation The dataset records aggregated activities of hundreds of thousands or millions of website users for business measurement purposes. Hence, it is impossible to identify a particular user. Moreover, to minimise the risk of disclosing business sensitive information, all experiment context is either removed or anonymised. One should not be able to tell who is in an experiment, when the experiment is run, what treatment the experiment involves, and what decision metrics are used. We refer readers to the accompanying datasheet² for further details in this area.

Potential misuses A digital experiment dataset, no matter how anonymised it is, reflects the behaviour of its participants under a specific application domain and time. We urge potential users of this dataset to exercise caution when attempting to generalise the learnings. It is important to emphasise that the learnings differ from the statistical methods and processes demonstrated on this dataset. We believe the latter are generalisable, i.e., they can be applied to other datasets with similar data dimensions regardless of the datasets’ application domain, target demographics, and temporal coverage, and appeal for potential dataset users to focus on such.

One example of generalising the learnings is using this dataset as a full performance benchmark. As discussed above, this dataset does not have a notion of ground truth, and any quasi-ground truths constructed are themselves a source of bias to estimators. Thus, experiment design comparisons need to be considered at a theoretical level [176]. Another example will be directly applying the value of hyperparameter(s) obtained while training a model on this dataset to another dataset. While this may work for similar application domains, the less similar they are, the less likely the hyperparameters learned will transfer. This may risk introducing bias, both in the estimator and in fairness.

Running digital experiments in general The dataset is released to support experiments with adaptive stopping, which will enable a faster experimentation cycle. The ethical concerns will naturally mount as we run more digital experiments, which are ultimately human subjects research. We reiterate the importance of the following three principles when we design and run experiments [206, 208]: respect for persons, beneficence (properly assess and balance the risks and benefits), and justice (ensure participants are not exploited), and refer readers to Chapter 9 of [155] and its references for further discussions in this area.

4.6 A Brief Recap

Digital experiments are a powerful tool for online organisations to assess the impact of their digital products and services. To safeguard future methodological development in the area, it is vital that we have access to and a systematic understanding of relevant datasets arising from real experiments. We described the result of the first ever survey on publicly available OCE datasets. We also provided a dimensional taxonomy that links the data collection and statistical test requirements for digital experiments in general. In addition, we released the first ever dataset that can support digital experiments with adaptive stopping, which design is grounded on a theoretical discussion between the taxonomy and statistical tests. Via extensive experiments, we also showed that the dataset is capable of addressing the identified gap in the literature.

Our work on surveying, categorising, and enriching the publicly available digital experiment datasets is just the beginning, and we invite the community to join in the effort. As discussed

above, we have yet to see a dataset that can support methods dealing with correlated data due to cluster randomisation, or the end-to-end design and running of experiments with adaptive stopping using Bayesian non-parametric tests. We also see ample opportunity to generalise the survey to cover datasets arising from uplift modelling tasks, quasi-experiments, and observational studies. Finally, we can further expand the taxonomy, which already supports datasets from all experiments, with extra dimensions (e.g., number of features to support stratification, control variate, and uplift modelling methods) as the area matures.

Chapter 5

Recent Advances in Digital Experimentation Methods

Part of the chapter (Section 5.5) is adapted from the research paper “*Measuring e-Commerce Metric Changes in Online Experiments*”, presented at *ACM Web Conference 2023 (WWW '23)* [177].

5.1 Motivation

After establishing the statistical tests and datasets required to run a digital experiment in Chapters 3 and 4, we can now describe what a digital experiment entails. To begin, we outline the causal reasoning behind experiments. Establishing causation (i.e., making conclusions such as “this treatment causes that business metric to change”) is a step further than establishing correlation¹ that a statistical test may do. We thus require additional theoretical grounding to represent the concept of causation. This chapter provides an initial view of the potential outcomes framework / Rubin causal model and its associated language.²

Using the potential outcomes framework, we can describe different classes of experiments. The classes, primarily characterised by the level of control an experimenter has, are randomised controlled trials (RCTs), quasi-experiments, and natural experiments. Having a

¹ Also known as *association* within the causal inference community [122].

² Other tools and frameworks exist, e.g., directed acyclic graphs and do-calculus, which are arguably more generalisable. Here, we prioritise the ease of understanding.

classification of digital experiments enables us to highlight the different methodological and implementation challenges faced by digital experimenters on each front.

Given the availability of excellent and up-to-date review articles that deal with established and emerging approaches in each area, it is perhaps not the most effective for us to cover yet another systemic review of digital experimentation methods. Thus, we provide a more lightweight, narrative overview beyond introducing the review articles. We will describe the major challenges and highlight several prominent works for each challenge as entry points for readers interested in further exploration.

As we continue to place equal emphasis on the theory and practice of digital experimentation methods, we also enrich the overview with case studies. These case studies provide spotlights on digital randomised controlled trials with dependent responses and quasi-experiments that assign treatments based on geographical regions. In addition to providing evidence on the extent of these challenges, the case studies also detail the practical considerations when implementing the methods to address them.

The rest of the chapter is organised as follows. We introduce the basics of the potential outcomes framework in Section 5.2. Using its notation and rules, we outline the three classes of experiments in Section 5.3. We then interleave the reviews on advances in each class of experiments (RCTs in Section 5.4 and quasi-experiments/natural experiments in Section 5.6) with spotlights on methods that deal with dependent responses (Section 5.5) and geo-experiments (Section 5.6.3).

5.2 From Correlation to Causation – The Potential Outcomes Framework

Readers may recall that in Chapter 3, we noted that statistical tests alone are insufficient to establish causal relationships between treatments and changes in responses. In this section, we bridge the gap between correlation/association thinking (as featured in statistical tests) and causal thinking (as featured in experiments and causal inference) by introducing the most basic concepts of the potential outcomes framework. In subsequent sections, we will refer to and extend these concepts.

This introduction extends the notations in Chapter 3 and takes a more leisurely pace than other introductions to background materials in the thesis. This is deliberate – in the author’s experience, it is a big jump to fully appreciate the framework’s premises, rules, and notations, especially for those who are used to dealing only with randomised control trials. Readers familiar with the potential outcomes framework and the $Y_i(0)$ -style notation in causal inference may consider skimming through the section. They may also find other expositions on the framework [62, 122, 131] useful.

Observed responses In Chapter 3, we often deal with two sets of responses: $X_1, \dots, X_n$ and $Y_1, \dots, Y_m$. In the general context of statistical testing, the responses can arise from two arbitrary sources, e.g., the height of males vs females in a survey or ice cream sales vs the number of drowning incidents [201]. In the context of an experiment, the two sets of responses would have come from experimental units that are assigned and subjected to two different treatments.³ Formally, by modelling the treatment assignment using a binary r.v. $\mathcal{A}$ and denoting the two possible assignments $\mathcal{A} = 0$ and $\mathcal{A} = 1$, we can represent the responses under the additional context as

$$X_i | \mathcal{A} = 0 \quad \text{and} \quad Y_j | \mathcal{A} = 1. \quad (5.1)$$

Assuming all X_i and Y_j are identically distributed copies of their respective populations, by taking the sample mean of the two sets of responses, we obtain an unbiased estimator of the expected response of the overall populations X and Y given the corresponding treatments:

$$\mathbb{E}(X | \mathcal{A} = 0) \quad \text{and} \quad \mathbb{E}(Y | \mathcal{A} = 1). \quad (5.2)$$

Counterfactual responses Note X_i and Y_j in Expression (5.1) are merely the *observable* responses from their associated experimental units. We do not observe how the experimental unit responding X_i under $\mathcal{A} = 0$ would have responded had $\mathcal{A} = 1$, nor do we observe how the unit responding Y_j under $\mathcal{A} = 1$ would have responded had $\mathcal{A} = 0$. These “would haves” are known as the *counterfactual* responses and are essential to causal reasoning – understand-

³ We focus on the two-treatment case for the rest of the chapter. The numerical designation of the treatment assignments, i.e., 0, 1, $\dots$, eases the extension of the framework to cover multiple treatments.

ing the impact of a treatment requires understanding both the outcomes when the treatment is and is not prescribed.

The potential outcome framework unifies the observed and counterfactual responses from the same experimental unit by treating all equally as potential responses (or *outcomes*) and giving each potential response a separate designation. In the two-treatment case, they are denoted as

$$X_i \triangleq \begin{cases} X_i(0) & \text{if } \mathcal{A} = 0 \\ X_i(1) & \text{if } \mathcal{A} = 1, \end{cases} \quad Y_j \triangleq \begin{cases} Y_j(0) & \text{if } \mathcal{A} = 0 \\ Y_j(1) & \text{if } \mathcal{A} = 1. \end{cases} \quad (5.3)$$

More generally, we can write $X_i = X_i(a)$ if $\mathcal{A} = a$ and $Y_j = Y_j(a)$ if $\mathcal{A} = a$, where $a \in \{0, 1\}$.⁴ The framework features a missing data problem – some potential responses (generally one max per experimental unit) are observable, and others must be estimated.

Using the framework, we can define the causal impact of a treatment over another more formally. In a two-treatment setting, the treatment effect (of treatment 1 over 0) on a single experimental unit that responds X_i or Y_j is

$$\Delta_{X_i} = X_i(1) - X_i(0), \quad \Delta_{Y_j} = Y_j(1) - Y_j(0). \quad (5.4)$$

Assuming again that all X_i and Y_j are identically distributed copies of their respective populations, we can also define the average treatment effect (ATE) across the overall populations X and Y as

$$\Delta_X = \mathbb{E}(X(1) - X(0)), \quad \Delta_Y = \mathbb{E}(Y(1) - Y(0)). \quad (5.5)$$

The framework also enables us to formalise the adage “correlation does not imply causation”. So far, we have encountered the expected observed response given a treatment $\mathbb{E}(X \mid \mathcal{A} = a)$ (see Expression (5.2)) and the expected potential response under the said treatment $\mathbb{E}(X(a))$. They may appear to be descriptions of the same quantity but are not the same. This is because a third variable (e.g., a confounder) may affect both X and $\mathcal{A}$.

⁴ In [122], the same quantities are denoted X_i^a and Y_j^a , sometimes with the numeric value of a explicitly specified, e.g., $Y_j^{a=0}$ for disambiguation.

To illustrate, we use the “ice cream consumption causes drowning” fallacy [201] – let X be the number of (observed) drowning incidents and $\mathcal{A}$ be a binary r.v., representing high ($\mathcal{A} = 1$) and low ($\mathcal{A} = 0$) ice cream sales. It is established that the number of drowning incidents increases along with ice cream sales. In other words, $\mathbb{E}(X | \mathcal{A} = 1)$ is high, $\mathbb{E}(X | \mathcal{A} = 0)$ is low, and the difference $\mathbb{E}(X | \mathcal{A} = 1) - \mathbb{E}(X | \mathcal{A} = 0)$ is positive. However, it is also established that ice cream consumption does not cause drowning, nor does drowning incidents cause increased ice cream sales.⁵ Instead, the correlation/association in the expected observed responses arises due to a third factor: weather. Hot and sunny spells generally raise the demand for ice cream and, unfortunately, the number of drowning incidents as more people visit various water bodies and thus are exposed to the risk of drowning.

The expected potential responses better reflect the lack of a causal relationship between ice cream sales and drowning incidents. Given the potential number of drowning incidents under high ice cream sales $\mathbb{E}(X(1))$, one would expect the counterfactual response $\mathbb{E}(X(0))$, i.e., the number of drowning incidents had the ice cream sales been low *and all else equal*, will be similar if not the same. This leads to a negligible causal impact $\mathbb{E}(X(1)) - \mathbb{E}(X(0))$.

5.3 Classes and Examples of Digital Experiments

To aid the discussion in the rest of the chapter, we provide an overview of different classes of experiments defined in the traditional statistics literature and further examples within digital experimentation. Figure 5.1 shows the classes of experiments and what differentiates a class from another.

5.3.1 Randomised controlled trials

The class of experiments where experimenters have the most control is randomised controlled trials (RCTs). They are also known as online controlled experiments (OCEs) in the digital experimentation community [155]. The free delivery banner experiment described at the beginning of the thesis (see Figure 1.1) belongs to this class. These experiments are

⁵ It is conceivable for drowning incidents to cause *decreased* ice cream sales due to emergency responses.




Figure 5.1: The onion model for different types of experiments, in decreasing degree of control from the experimenter (inner to outer): Randomised controlled trials, quasi-experiments, and natural experiments. Quasi-experiments and randomised controlled trials are differentiated by the experimenter’s ability to assign experiment participants randomly into different groups for different treatments. Natural experiments and quasi-experiments are differentiated by the experimenter’s ability to vary the treatment (or lack thereof) prescribed to an experiment participant.

characterised by the experimenter’s ability to apply a scientific control and randomly assign experimental units to different groups.⁶ For the free delivery banner example, this corresponds to the ability to show a version of the website without the “free delivery” banner and randomly assign users to see different website versions, respectively.

Most of the value of an RCT comes from the ability to perform random assignments, which has two major implications. Here, we describe the implications using the potential outcomes framework. Firstly, randomisation generally implies exchangeability.⁷ This means the potential responses and the actual treatment are independent. We denote such exchangeability as $X(0), X(1) \perp\!\!\!\perp \mathcal{A}$ and $Y(0), Y(1) \perp\!\!\!\perp \mathcal{A}$ and observe that if the potential responses are ex-

⁶ Unlike in medical trials, the issue of *blinding*, where information that may influence the result is withheld from the experimenter or experiment participants (or both) to minimise biases, is rarely discussed in digital experimentation. It is often addressed by having the Internet between the experimenter and participants: the overwhelming majority of participants are unaware of what treatment(s) are in place when they visit a website. Moreover, experimenters would only know who the participants are during post-experiment analysis as they delegate the task of assigning and tracking participants to automated computer systems.

⁷ Also known as *unconfoundedness* or *ignorability*. While there are randomisation schemes that lead to non-exchangeable potential responses, in practice, almost all online controlled experiments rely on randomisation schemes (e.g., simple randomisation) that yield some form of exchangeability [240].

changeable,

$$\begin{aligned}\mathbb{E}(X(a)) &= \mathbb{E}(X(a) \mid \mathcal{A} = 0) = \mathbb{E}(X(a) \mid \mathcal{A} = 1) \text{ and} \\ \mathbb{E}(Y(a)) &= \mathbb{E}(Y(a) \mid \mathcal{A} = 0) = \mathbb{E}(Y(a) \mid \mathcal{A} = 1), \quad a \in \{0, 1\}.\end{aligned}\quad (5.6)$$

Secondly, as experimental units in both groups can be viewed as a simple sample of the same underlying population under randomisation, the r.v.'s representing the potential responses of *both* groups under the same treatment, i.e.,

$$X_1(a), \dots, X_n(a), Y_1(a), \dots, Y_m(a) \text{ for } a \in \{0, 1\}, \quad (5.7)$$

are interchangeable.⁸ This means they can be swapped within and across the groups without consequences.

The two implications mean we can obtain an unbiased estimate/estimator on the average treatment effect using only the observed responses under randomisation. To illustrate, we can express the ATE defined in Equation (5.5) as

$$\begin{aligned}\Delta_X &= \mathbb{E}(X(1) - X(0)) && \text{(By definition – Eq. (5.5))} \\ &= \mathbb{E}(X(1)) - \mathbb{E}(X(0)) && \text{(Linearity of expectation)} \\ &= \mathbb{E}(Y(1)) - \mathbb{E}(X(0)) && \text{(Interchangeability between } X(1) \text{ and } Y(1)) \\ &= \mathbb{E}(Y(1) \mid \mathcal{A} = 1) - \mathbb{E}(X(0) \mid \mathcal{A} = 0) && \text{(Exchangeability of } Y(1) \text{ and } X(0) \text{ – Eq. (5.6))} \\ &= \mathbb{E}(Y \mid \mathcal{A} = 1) - \mathbb{E}(X \mid \mathcal{A} = 0) && \text{(By definition – Eq. (5.3)).} \\ & && (5.8)\end{aligned}$$

We can also show that Δ_Y is equal to the last line of Equation (5.8) using the same arguments (plus the interchangeability between $X(0)$ and $Y(0)$). This should not come as a surprise as X and Y are based on experimental units from the same population and thus should share the same ATE. As stated in Section 5.2, estimating the last line of Equation (5.8) is simply a matter of taking two sample means over the two groups of observed responses. This has the effect of reducing the ATE estimation problem into a much more straightforward two-sample

⁸ Also known as an *exchangeable sequence of r.v.'s* [44]. There are a few subtle differences between the two exchangeability definitions. Thus, we use a different name to disambiguate these two concepts.

testing problem. Experimenters can use the many readily available methods in statistical testing (see Chapter 3) and avoid estimating expected counterfactual responses, which may attract considerable bias.

To sum up, the results above show that randomisation generally implies that any difference in the responses can be attributed to the causal effect of the treatment and not to other measured or unmeasured confounding. Such a desirable property perhaps explains why RCTs are considered the “gold standard” in medical trials [139] and are the most common type of digital experiment. We will outline further statistical challenges to running RCTs online in Section 5.4.

5.3.2 Quasi-experiments

Taking away the ability to randomly assign experiment units in randomised controlled trials leaves us with a quasi-experiment. Consider the example in [287], where a digital organisation developed a mobile application downloadable from app stores such as Apple App Store and Google Play. The organisation is interested in whether a new app feature shipped within a new app version will improve a business metric relative to the existing app version. For the new feature to reach app users, developers usually upload the new version onto the app stores. The app stores then have to review the new app version, and once approved, app users have to update their app.

In this scenario, the ability to assign who sees the new or existing version of the app (let alone randomly) is taken away from the experimenters and put in the hands of the users, who decide when to update the app, and the app stores, who decide when to allow the new version of the app to be downloaded.

The inability to randomise treatment assignments generally marks the beginning of causal inference’s remit. Experimenters can no longer expect the exchangeability mentioned in the previous subsection to automatically hold and have to employ other methods to estimate the counterfactual responses from one or more observed responses. In the language of the potential outcomes framework, we want to estimate $X_i(1)$ and $Y_j(0)$ using $X_i | \mathcal{A} = 0$ and $Y_j | \mathcal{A} = 1$, where X_i represents the responses from users who have yet to download the new app version, Y_j represents the responses from users who have seen the new version

(and presumably interacted with the new feature), and $\mathcal{A}$ be the binary treatment assignment indicator. We will introduce several causal inference approaches and how they apply to quasi-experiments in Section 5.6.

5.3.3 Natural experiments

If we further remove the ability to apply scientific control,⁹ we end up with natural experiments. While experimenters cannot apply different treatments (in this case, different versions of a digital product or experience), participants may still encounter scenarios that include one or more of the said treatments by the act of nature. This means it remains possible to draw inferences on the causal effect of the treatments from observational data.

An experimenter may opt not to apply a scientific control and run a natural experiment instead due to various operational, reputational, legal and ethical considerations. For example, experiments aiming to measure the value added by existing call centre support may be prevented from actively blocking some users' access to the call centre as it negatively impacts user experience. Other experiments seeking to understand the price elasticity of demand¹⁰ may be keen to avoid engaging in the controversial practice of price discrimination¹¹ to avoid negative press and risk of legal actions [73, 282]. Lastly, testing different news headlines simultaneously in front of readers may also damage editorial credibility or even “fan the flames” of partisan political beliefs in the long run [133]. Experimenters may instead leverage responses from users who have never contacted the call centre, sales on similar digital products/services priced differently, and reactions on past news headlines on similar topics to understand the respective causal impact.

Most natural experiments share the same causal inference approaches as quasi-experiments. We will cover a few applications in Section 5.6.

⁹ Also known as *manipulating* the treatment or independent variables [231].

¹⁰ The relationship between the price of a product/service sold by an organisation and its demand.

¹¹ The act of offering different users a different price for the same digital product/service.

5.3.4 Special case – Experiments for personalisation strategies

It is worth noting that the boundaries between each experiment class introduced above are not clear-cut. Depending on the situation, experimenters may have various degrees of control over random assignments.

Consider a follow-up experiment to our free delivery banner experiment introduced above. Suppose the organisation is now interested in whether varying the *timing* of showing the “free delivery” banner will lead to further improvement. One can run an experiment as simple as comparing “showing the banner after a user has viewed *three* shop products” vs “showing the banner after a user has viewed *four* shop products” and see which timing rule leads to a better outcome.

In the experiment, the treatment (timing rule) a user is subjected to is *jointly* determined by the assignment from experimenters and the user’s actions. While the experimenters can randomly assign which timing rule to put the user through, the user must also qualify under the rules by seeing three or four shop products. Only by doing so, the user would see the “free delivery” banner, without which the notion of treatment effect would not make sense. This means experimenters have some, but not complete, control in random assignments – more than that in the new app feature experiment described in Section 5.3.2, but certainly less than the vanilla free delivery banner experiment discussed in Section 5.3.1.¹²

The banner display timing experiment is one of many experiments for *personalisation strategies* – complex sets of targeted user interactions common in e-commerce and digital marketing.¹³ We will focus on the design of these experiments and further elaborate on their unique challenges in Chapter 6.

¹² Some literature describes this scenario as *one-sided non-compliance* [270].

¹³ The example provided is grossly simplified to facilitate discussion. In practice, display timing strategies, in true personalisation spirit, usually specify how different users are shown their “free delivery” banner at a different stage of their website visit, depending on their browsing behaviour. Thus, an experiment that compares two display timing strategies is essentially comparing the execution of two *sets* of display timing rules generated from two meta-principles.

5.4 Advances in Online Controlled Experiments

Here, we provide a brief review of the advances in OCEs. Since its inception in the late 1990s [153], the field has seen tremendous growth in research output and its economic and societal impact [263]. More recently, we have also seen some comprehensive and up-to-date reviews in the area:

- [155] provided a narrative, general-purpose guide to the field while collating many relevant works.¹⁴
- [12] and [225] provided systematic literature reviews. The former built a taxonomy based on the process-infrastructure distinction, type of methodological approach, and the challenges/benefits involved. The latter categorised the works by the subject of an experiment,¹⁵ choices made in the experiment design/execution processes, and the role of stakeholders.
- [116] and [159] described the challenges. The former provided a general view by including data, engineering, and organisational culture topics. The latter focused on the statistical aspect and provided a more technical treatment to each topic presented.

This review aims to be a gateway to the multitude of research works in the area by unifying the many nomenclatures, providing a concise problem statement for each challenge, and suggesting works accessible to those new to a particular line of work. For a fuller exposition of the challenges and research in this area, readers should refer to the five high-quality reviews listed above.

Most of the section is dedicated to advanced statistical methods, building on concepts and notations covered in Chapter 3 (on statistical testing) and Sections 5.2 and 5.3.1 (on the potential outcomes framework). We will also briefly outline works addressing other topics in running successful OCEs, as mentioned in the second research question in Section 1.2. These include choosing decision metrics / evaluation criteria, scalable computer systems, and fostering a growth mindset organisational culture.

¹⁴ The author reviewed part of the manuscript before the book's publication in 2020.

¹⁵ Defined in [225] as the application domain and the element subjected to experimental manipulation.

5.4.1 Advanced statistical methods

We begin with the statistical challenges in OCEs that necessitate the development of advanced methods. These challenges included in the section are the multiple testing/comparison problem, post-selection inference, heterogeneous treatment effects, dependent responses, stable unit treatment value assumption (SUTVA) violations, and improving experiment sensitivity.

Multiple testing/comparisons Recall from Section 3.6.3 that the probability of rejecting the null H_0 in an NHST when H_0 is true (i.e. the false positive rate) is α . The false positive rate is usually small so that when a test rejects H_0 , we have reasonable conviction to hail the experiment result as a discovery. That said, the probability of committing a false positive accumulates as we perform additional comparisons. For N_{test} independent NHSTs, we expect $(\alpha \cdot N_{\text{test}})$ tests to reject H_0 when H_0 is true. This means that for every 20 NHSTs run under $\alpha = 0.05$, we expect the tests to reject H_0 once on average, even when all the H_0 are true [202]. Hailing all test rejections as discoveries without accounting for the inflated false positive rate will lead to biased decisions.¹⁶

The problem is increasingly prevalent due to the increasing number of experiments, treatment pairs, and user segments to compare (see paragraph on heterogeneous treatment effect below). Usually, experimenters seek to control the family-wise error rate (FWER) or the false discovery rate (FDR), defined over many experiments as

$$\text{FWER} \triangleq \mathbb{P}(\# \text{ Type I error} \geq 1) = 1 - \mathbb{P}(\# \text{ Type I error} = 0), \quad (5.9)$$

$$\text{FDR} \triangleq \mathbb{E} \left(\frac{\# \text{ Type I error}}{\# \text{ Test rejecting } H_0} \right). \quad (5.10)$$

The most common techniques that control for FWER include Bonferonni correction [80] and Šidák correction [249], whereas those for FDR include the Benjamini-Hochberg [25] and knockoff procedures [97].

¹⁶ This is closely related to the continuous monitoring/peeking problem discussed in Section 3.8 and Section 4.1. That said, comparisons in the continuous monitoring/peeking problem are often correlated, leading to the risk of committing a false positive accumulating differently.

Post-selection inference Having multiple comparisons over experiments, treatment pairs and sub-populations will remain a potential source of bias even when we move away from binary decisions and on to effect size estimation. Consider a two-sample, right-tailed t -test (and all its associated notation, see Section 3.5.2). In addition to deciding to reject/not reject $H_0 : \Delta = \theta_0$, we also estimate the treatment effect Δ . We know $\bar{Y} - \bar{X}$ is an unbiased estimator for Δ . However, with many statistical tests, we generally ignore all tests that fail to reject H_0 – assuming $\Delta = \theta_0$ as per the null hypothesis – and focus on tests that reject H_0 . The treatment effect estimate after such test selection now overestimates Δ :

$$\mathbb{E}(\bar{Y} - \bar{X} \mid T > t_{\nu, 1-\alpha}) = \mathbb{E}\left(\bar{Y} - \bar{X} \mid \bar{Y} - \bar{X} > \theta_0 + t_{\nu, 1-\alpha} \sqrt{\frac{s_X^2}{n} + \frac{s_Y^2}{m}}\right) > \mathbb{E}(\bar{Y} - \bar{X}) = \Delta. \quad (5.11)$$

Another way to put this in the context of multiple comparisons is that the test selection process above favours “lucky” OCEs, some of which are merely false discoveries. Inferring quantities with the test selection process is an example of *post-selection inference* [65], and the estimation bias it introduces is also known as the *winner’s curse* [161].

The best way to detect the existence of the winner’s curse and correct the bias is via replication [65]. In practice, it is often operationally challenging. It also seems unnecessary if one only seeks the aggregated treatment effect from multiple successful OCEs; in this case one can employ a maximum likelihood approach to estimate the expected bias from the test selection process [161]. To estimate the treatment effects of individual OCEs, one can perform pseudo-replication by splitting an experiment into two sub-experiments with the same treatments and build a regression model on the expected treatment effect of the second sub-experiment (without any selection) using that of the first sub-experiment (which goes through test selection) [45]. They can also use empirical Bayes methods [14] such as Tweedie’s formula and the James–Stein estimator [85]. Finally, [65] combined elements of the approaches above to improve the precision of individual treatment effect estimates.

Heterogeneous treatment effects Many experiments seek to measure the average treatment effect across the population. That said, it is valuable for experimenters to understand whether there are subgroups whose treatment effect may deviate from that of the entire population [79]. To detect such heterogeneous treatment effects, experimenters often estimate

the conditional average treatment effect (CATE) instead, i.e.

$$\mathbb{E}(Y_j(1) - Y_j(0) | \mathcal{C}_j), \quad (5.12)$$

where $\mathcal{C}_j$ represents the covariates for experimental unit j . Many works assume conditional interchangeability and exchangeability, such that they can estimate CATE using the conditional observed responses $\mathbb{E}(X_i | \mathcal{C}_i = c)$ and $\mathbb{E}(Y_j | \mathcal{C}_j = c)$ from an OCE. See Section 5.6.1 for further discussions on CATE estimation and Section 3 of [159] for a summary of works in the area.

Dependent responses Many OCEs feature statistical tests (see Section 3.5) that assume i.i.d. responses. While many OCEs do have observations that are approximately i.i.d. [67], such assumptions are not always justified. We have (implicitly) covered deviations from the identically distributed assumptions when we discussed heterogeneous treatment effects above. Below, we look at deviations from the independence assumptions.

One can find dependent responses in many experiment designs. A significant portion features dependent treatment assignments. This could be explicitly imposed by randomly assigning pre-existing groups of experimental units (instead of individuals) to treatment groups, also known as *cluster randomisation*, as briefly discussed in Section 4.4.1. It could also implicitly arise as one randomises individual experimental units independently yet observes and infers based on more granular responses, effectively forming a cluster per experimental unit.¹⁷ In Section 5.5, we will present an extensive case study on dependent responses in e-commerce OCEs, which arises from experimenters randomising by users and measuring changes to business metrics based on more granular transaction- or item-based responses.

Having dependent responses can affect treatment effect estimates and decisions from statistical tests. Consider a one-sample t -test (see Section 3.5.2 for all its associated notation). We recall from Equation (3.11) that its test statistic is $(\bar{X} - \mu_X) / \sqrt{s_X^2/n}$. Moreover, the test

¹⁷ Digital experimenters often differentiate an experimental unit, a *randomisation unit* (one subjected to independent randomisation), and an analysis unit (one that yields a response). These units do not necessarily coincide [70, 67].

statistic's denominator is an estimator of the standard error of the mean (SE):

$$\sqrt{\text{Var}(\bar{X})} = \sqrt{\frac{1}{n^2} \left(\sum_{i=1}^n \text{Var}(X_i) + 2 \sum_{i < j} \text{Cov}(X_i, X_j) \right)}. \quad (5.13)$$

Under i.i.d. assumptions, $\text{Cov}(X_i, X_j) = 0 \ \forall i \neq j$ and thus $\text{Var}(\bar{X}) = \sum_i \text{Var}(X_i)/n^2 = \sigma_X^2/n$. This implies $\sqrt{s_X^2/n}$ is an unbiased estimator of the SE. However, if X_i are dependent, then $\text{Cov}(X_i, X_j) \neq 0$, which makes $\sqrt{s_X^2/n}$ a biased estimator. In the case study below, we will demonstrate how the bias affects the power and confidence interval calculations and, thus, the decisions made in statistical tests and experiments.

Digital experimenters use the delta method [64] or one-way/block bootstrap [19] to mitigate the bias. We will provide a critique of these methods in the case study below. Another interesting work in the area is [67], which explores scenarios where an experimenter can treat responses as i.i.d. before providing a variance formula that exploits the relationship between the experimental/randomisation/analysis units and accommodates a wide range of randomisation mechanisms in practical settings.

Stable unit treatment value assumption (SUTVA) violations Responses can also become dependent by violating the *stable unit treatment value assumption* (SUTVA), an assumption we made implicitly when we defined the potential outcomes in Section 5.2 / Equation (5.3).¹⁸ More precisely, SUTVA is an assumption in two parts – in the words of [131]:

1. (No interference [51, 234]) – “The potential outcomes for any [experimental] unit do not vary with the treatments assigned to other units, and,”
2. (No hidden variations of treatments) – “[F]or each [experimental] unit, there are no different forms or versions of each treatment level, which lead to different potential outcomes.”

In OCEs, interference mainly arises due to the network effects. As experimental units interact with each other in the network, their potential response no longer depends on what treatment they have received but also on what treatments their connections have received. In other

¹⁸ Also known as having *spillover effects* [250, 253], in the sense that a treatment spills over to experimental units in another treatment group (due to, e.g., interactions between experimental units in different groups) and induces unwanted treatment effects.

words, a treatment may spill over to an experimental unit's connections. The effect can be found directly in social networks (where a user interacts with their connections) [142] or indirectly via marketplaces (where suppliers/buyers compete for shared resources) [167] and recommender systems (where the system recommend multiple items/users to the same user/item) [221].

While dependent treatment assignment, which breaks the independence assumption but maintains SUTVA, generally affects the variance estimate, a SUTVA violation can bias the treatment effect estimates. [19] have shown that interaction amongst users and, separately, the items recommended can result in an increased false positive rate in experiments when applied to Facebook data, likely due to bias in both the treatment effect and variance estimates.

A common way to mitigate the spillover effect in social networks is to randomise treatment group assignments over clusters¹⁹ of experimental units (e.g. communities of users) instead of the units themselves [18]. Clustering by non-overlapping communities means one can hold SUTVA at a cluster level, as the units generally interact heavily within the same cluster but sparsely between clusters [169]. However, the choice of clusters is non-trivial; this means a cluster-based treatment assignment may only be justified when the network effect is apparent. [239] developed a framework that detects whether network effects exist for a particular OCE by simultaneously running the same experiment under individual-based and cluster-based assignment and testing if the treatment effect differs (and hence if SUTVA holds up reasonably to justify individual-based assignment). Section 6 of [159] describes many more relevant works, including those focusing on indirect network effects.

The second part of SUTVA (no hidden variations of treatments) is seldom explicitly discussed in OCEs. That said, it is conceivable that many “imperfections [...] hidden throughout the engineering stack or in the design process” [283] are capable of breaking this part of SUTVA and thus bias the (potential) responses. Many examples are presented as pitfalls in OCE operations, which we will cover in Section 5.4.2 / Table 5.1.

¹⁹ In network speak, sub-graphs that are highly connected within.

Improving sensitivity We look beyond challenges that lead to biases in our treatment effect estimator. Recall from Section 3.6.4 that the power of a test depends on many factors, including the significance level, sampling uncertainty, sample size and effect size. While we have discussed the role of sample and effect sizes in Sections 3.6.5 and 3.6.6, respectively, little regarding the sampling uncertainty has been said so far in the thesis. Some works look into increasing the sensitivity of OCEs by reducing sampling uncertainty, which enables an experimenter to reduce the sample size required for an experiment or the minimum detectable effect (or both). The works are motivated by observations where [150] noted that ordinarily, successful experiments in technology companies only improve metrics by a fraction of a per cent; [286] also described the need to detect small effects as huge customer bases often translate them into substantial gains in revenue and profit.

The methods described below rely on decomposing the variance in a metric and attempting to eliminate unnecessary components. Controlled-experiment Using Pre-Existing Data (CUPED) is a control variates method that measures a metric modified by a linear function of the covariates [69].²⁰ The idea is to eliminate the effect of correlated covariates on the metric under measurement. [63] later apply the idea to reformulate dilution²¹ as a variance reduction problem, resulting in a more easy to implement and less error-prone formula. A similar idea is used by [222], which subtracted the predicted value of a metric using boosted decision trees as predictors. In the same spirit, stratified sampling is used to eliminate variance between strata. Stratified sampling can be applied pre- or post-experiment, and while the theoretical bounds are better for pre-experiment stratification, post-stratification is often preferable for large-scale OCEs [286].

5.4.2 Other topics

To complete our review of OCEs, we also review topics that we did not cover in the statistical tests and data chapters, as well as the advanced methods section above.

²⁰ Also known as *regression adjustment* in econometrics [168].

²¹ The act of translating the treatment effect from users who have been exposed to a treatment (“triggered” users) to the general user base. The metric is diluted as we incorporate untriggered users (who contribute zero).

Anatomy of an OCE A mature OCE system/platform often features complex processes and infrastructure, and many seek to provide an overview or describe specific parts of the ecosystem. [154] provided a definitive guide on setting up, running, reporting, and scaling OCEs. Several papers described the scalable, concurrent testing infrastructure and its associated challenges at major technology companies [149, 260, 288]. Others described the OCE lifecycle [91] and characterised the elements defining an experiment [11].

State and progress of OCEs' usage A class of works aim to capture the state and track the progress of the use of OCEs (or DEM capabilities in general) across many organisations. Some propose models that represent the maturity/growth of DEM capabilities [92, 280]. Others focus on experiment outcomes and perform meta-analyses on results from multiple OCEs – see [29, 30, 136] for results across industries, in e-commerce, and in display advertising, respectively. [103] bridged the two topics above by linking experiment performance to organisational structure.

Pitfalls in OCE operations Despite the increasing adoption and maturity of the use of OCE, running a successful OCE after another is, by the admission of the biggest players in the industry, difficult [155]. This is due to many non-obvious challenges and pitfalls experimenters face while running an OCE end-to-end. In addition to that related to statistical testing (see Chapter 3) and requiring advanced statistical methods (see Section 5.4.1), many accounts [53, 74, 89, 152, 148] describe pitfalls related to experimentation infrastructure, decision metric selection and interpretation [68], and external effects [126]. We summarise these challenges and pitfalls in Table 5.1. To mitigate the potential problems, [179] proposed techniques that separate problematic user responses affected by some of the external effects (namely dilution, carry-over effect, and novelty effect) from valid ones. [41, 93] also developed a series of diagnostic tests to identify the root cause(s) of a sample ratio mismatch, a common symptom for tests that are potentially invalidated due to encountering one or more of the pitfalls (see “practical usage” paragraph of Section 3.7.2).

5.5 A Spotlight on Dependent Responses in e-Commerce Experiments

To complement our review of advances in OCE, we present a case study on dependent responses in e-commerce OCEs. We intend for the case study to be a short but crucial piece of evidence that drives awareness of the pitfall among experimenters in e-commerce and hence encourages the adoption of established mitigation approaches.

Like many technology organisations, e-commerce organisations aim to provide a consistent user experience; hence, randomisation in OCEs is generally done on a per-user basis. On the other hand, e-commerce organisations are unique as they feature physical inventories and thus track a set of business metrics that are transaction- or item-based. These metrics include *average basket value* (ABV),²² *average basket size* (ABS),²³ and *average selling price* (ASP).²⁴

Experiments that measure changes to ABV, ABS, and ASP often feature dependent responses. In these experiments, the responses (or *analysis units*) are at a more granular transaction or item level than the *randomisation units*, which are at a user level. Given that a user can make many transactions and purchase many items during an experiment, the value of these transactions and items will likely correlate based on the user's preference. This creates a local dependence structure between users and transactions/items, which violates the usual i.i.d. assumptions in common statistical tests that OCEs employ (e.g., a Student's *t*-test). If left unmitigated, it risks experimenters having wrong estimates of sampling uncertainty and making wrong conclusions from a statistical test.

Having dependent responses in experiments due to differences in granularity between randomisation and analysis units is not a recent revelation [155]. In some sense, one may regard such a setup as a cluster-randomised controlled experiment, where an experimenter randomly assigns clusters of transactions or items belonging to the same user to the same treatment groups.²⁵ In terms of obtaining an unbiased estimate for the sampling uncer-

²² Average (mean) amount spent in each transaction/basket by a user.

²³ Average (mean) number of items purchased in each transaction/basket by a user.

²⁴ Average (mean) price per item sold.

²⁵ In such a setup, the size of a cluster can be zero and would only be determined after the experiment ends. This is because experimenters cannot control how much a user buys on the website.

tainty, many viable approaches already exist: experimenters in medicine, economics, and social sciences often employ a crossed random effects model [16] or cluster robust standard errors [32]. Meanwhile, those in digital technology prefer using bootstrap [19] or the delta method [64] to estimate the sample variance and the standard error due to their ability to scale to large datasets and relative lack of model assumptions.

Despite the existing literature dedicated to dependent responses in digital experiments, there is little published evidence specifically in the context of OCEs in e-commerce. Most published research is based on the context of OCEs in digital advertising and content, where experimenters randomise by users and analyse by sessions or page views [19, 64]; or social networks, where experimenters both randomise and analyse by users whose responses may correlate via their social connections [83, 114]. We believe the lack of evidence contributes to insufficient awareness of the issue from experimenters in this area.

This case study's contribution is evidence from real-life e-commerce experimentation operations. In particular:

1. We show, using three e-commerce transaction datasets (with two publicly available), a positive correlation between the value/size of transactions and items from the same user. This leads to inflation in measurement uncertainty (Section 5.5.1);
2. We quantify the extent of the said inflation, which can be anywhere between 1x to >100x the usual estimate and is dependent on the business metric, experiment duration, and organisation (Section 5.5.2);
3. We highlight the impact of the said inflation on test power and confidence interval coverage in null hypothesis significance tests (Section 5.5.3); and
4. We share lessons learned while incorporating some of the established mitigation approaches into an e-commerce company's internal experiment analysis platform, including a critique of some popular approaches and practical tradeoffs during implementation (Section 5.5.4).

Table 5.2: Summary of the three online retail/e-commerce transaction datasets featured in this paper (post data cleaning).

Dataset	# Users/ Customers	# Transactions/ Orders	# Items/ Units	# Products/ SKUs	Time span
ASOS (Proprietary)	~4 180 000	~9 680 000	~32 900 000	~330 000	62 days (2 months)
UCI Online Retail II [78, 40]	5 852	36 594	10 690 447	4 621	739 days (2 years)
Olist Brazilian e-Commerce [209]	94 983	98 199	112 101	32 729	729 days (2 years)

5.5.1 Beware of Hidden Inflation in Measurement Uncertainty

We first motivate why we may obtain inaccurate estimates of the sampling uncertainty when measuring changes to e-commerce metrics based on transactions or items in OCEs. Recall from the previous section (see Equation (5.13)) that $\sqrt{s_X^2/n}$, the vanilla sample standard error (SE), is an unbiased estimator for the SE under i.i.d. assumptions yet becomes biased when the responses are dependent.

The i.i.d. assumption is the prevailing assumption when experimenters randomise by users and analyse by user [67]. However, it is unlikely to hold when experimenters randomise by users and analyse by transactions or items. Many users tend to make a transaction that is similarly sized and valued to their previous transaction(s). They also tend to purchase items that are at a similar price point. To demonstrate the claims above, we utilise three online retail/e-commerce transaction datasets. The UCI Online Retail II [40, 78] and Olist Brazilian e-Commerce [209] datasets are publicly available.²⁶ We also use a proprietary dataset from ASOS.com, a global online fashion retail company, that records transactions within a specific two-month period in 2022 on a particular mobile platform.²⁷ We summarise the three datasets in Table 5.2.

²⁶ The experiment code and results on the two publicly available datasets are available on GitHub: <https://github.com/liuchbryan/oce-ecomm-abv-calculation>.

²⁷ The data described are not representative of ASOS.com's overall business operations, and one should not draw any such conclusion from the dataset.




Figure 5.2: Kernel density estimation plots showing a user's transaction (Top row) value / (Bottom row) size against that of their next transaction – if it exists – in the (Left) ASOS, (Middle) UCI Online Retail II, and (Right) Olist Brazilian e-Commerce datasets. We normalise the transaction values/sizes by the ABV/ABS of respective datasets.

For each dataset, we plot the value and size of a user's transaction against that of their next transaction (if it exists) in Figure 5.2. We observe from the figures a high empirical density along the $x = y$ diagonal and strong directionality in the kernel density estimates. This suggests a positive correlation between transaction values/sizes from the same user. For brevity, we do not show the plots for item values, yet contend that a perfect correlation arises when a customer purchases multiple units of the same product.

As responses from the same user become positively correlated, we have $\text{Cov}(X_i, X_j) > 0$ for some i and j in Equation (5.13) and thus a higher SE than the vanilla sample SE.

5.5.2 Re-estimating the SE using bootstrap

We now explore how much the SE can inflate. We apply the bootstrap procedure described in Section 2.2 of [19] to the datasets described above to re-estimate the SE. Bootstrapping




Figure 5.3: The ratio between the one-way bootstrap standard error (SE) estimates [19] and the vanilla sample SE estimates for (Left) ABV, (Middle) ABS, and (Right) ASP. The error bars represent the 95% confidence interval for the one-way bootstrap SE estimates.

generates a resample by sampling the original set of responses with replacement (or applying a random weight in this case). This yields a different sample mean. By repeating the process many times and taking many bootstrap means, we can estimate the SE by calculating the standard deviation of the bootstrap.

We use a one-way bootstrap (a.k.a. block bootstrap) to account for the dependency between transactions/items and users. Instead of generating the resample by sampling each transaction/item individually, we sample clusters of transactions/items belonging to the same user, mirroring our randomisation process.²⁸ Here, we use random weights generated from a Poisson(1) distribution to reweight our samples before calculating a bootstrap mean. The random weights emulate the number of times a transaction/item was sampled with replacement (along with other transactions/items associated with the same user). Finally, we estimate the SE from 500 bootstrap means.

Figure 5.3 shows how the one-way bootstrap SE differs from the vanilla sample SE in successive expanding windows, corresponding to different potential OCE durations. We observe that the bootstrap SE is significantly higher than the vanilla sample SE, with the exact

²⁸ For ASP, in addition to the dependency between items and users, there may be additional dependence between items and products/stock-keeping units (SKUs). One may account for both using a two-way bootstrap described in Section 5.5.4 or Section 2.3 of [19].




Figure 5.4: Trajectories of the one-way bootstrap standard error (SE) estimates (solid colored lines) and the vanilla SE estimates (dashed black lines) under different dataset-metric combinations. Estimates are cumulative from day one of the datasets. All x- and y-axes are on different scales.

ratio between the one-way bootstrap SE and vanilla sample SE heavily dependent on the dataset/audience, business metric, and duration. For example, the bootstrap SE for ASP in the UCI Online Retail II dataset is $>100\times$ the vanilla SE as many of their users buy tens or hundreds of the same product. Meanwhile, the ratios for ABV and ABS in the Olist Brazilian e-Commerce dataset are practically one. This is due to only 3.2% of transactions coming from returning users. Given such differences, experimenters should strive to obtain a reasonably accurate estimate.

We also observe from Figure 5.4 that the bootstrap SE may no longer drop as more transactions are involved over time and may go up again in some cases. This is likely due to returning users making further transactions, thus creating more and larger clusters of transaction/item values and sizes. It leads to the response variance increasing, sometimes more quickly than the increase in sample size. The observation suggests that lengthening an experiment solely to collect more responses (hence lowering the SE) may backfire. In addition to established practices on sample size estimation [228], experimenters should also consider how the variance of their metrics evolves when designing OCEs with dependent data.




Figure 5.5: (Left) The actual test power of and (Right) the coverage of the centred confidence interval (CI) from a two-tailed z -test against different standard errors (SE), expressed as a multiple of the vanilla sample SE $\sqrt{s^2/n}$. We choose the test parameters such that the test power = 80% and the CI coverage = the nominal values when the SE multiples are one.

5.5.3 Impact on OCE decisions

We finally discuss how an inflated SE due to dependent data can affect decisions made from a NHST. Firstly, it reduces the power of the test, making any potential treatment effect harder to detect. Consider a two-tailed Student's t -test with significance level α and ν d.f. Its power (adapted from Equation (3.39)) is

$$1 - T_{\nu}(t_{\nu,1-\alpha/2} - \theta/SE) + T_{\nu}(-t_{\nu,1-\alpha/2} - \theta/SE), \quad (5.14)$$

where θ is the effect size, $T_{\nu}(\cdot)$ is the CDF, and $t_{\nu,q}$ is the q^{th} quantile of a t -distributed r.v. If the SE increases, then both the standardised effect size θ/SE and the test power in Expression (5.14) decrease.

Secondly, having an inflated SE without knowing so will lead to tests producing confidence intervals that are too narrow, risking more false positives. In the test above, the $(1 - \alpha)$ confidence interval (CI) is $(\bar{x} \pm t_{\nu,1-\alpha/2} \cdot SE)$ – see Equation (3.4). Fixing the CI bounds, we observe that $t_{\nu,1-\alpha/2}$ must decrease (and cease to be a $(1 - \alpha/2)$ quantile) when the SE increases. This means that the said CI will no longer have a $(1 - \alpha)$ but a lower coverage, i.e., there is now a lower chance the interval will contain the actual effect size.

To show the full extent of the issues above, we plot the test power and CI coverage of a two-tailed z -test, essentially a Student's t -test with large degrees of freedom, against different SEs in Figure 5.5. We observe that the power of a z -test with a 5% significance level (dashed line) tumbles from 80% to around 29% when we merely double the vanilla sample SE. Moreover, the centred 95% CI calculated using the vanilla sample SE would only have roughly 67% coverage. These results reinforce the importance of having an accurate estimate of the SE.

5.5.4 Lessons Learned During Deployment

We finally share some lessons learned when incorporating the calculations of e-commerce metrics into ASOS.com's experiment analysis platform, which has been in production since Sep 2019. The analysis platform extends the functionality of a third-party experimentation platform, namely splitting, logging, and alerting. It adds value to the business by computing and performing statistical tests on bespoke business metrics that use internal datasets.

As mentioned at the beginning of the section, experimenters in digital technology prefer the use of bootstrap (see Section 2.2 of [19]) and the delta method (see Section 3 of [64]) to re-estimate the SE, which acts as a plug-in estimate for subsequent statistical tests. We described and implemented the bootstrap approach in Section 5.5.2 when we quantified the extent of the SE inflation. The delta method approach expresses a transaction-/item-based metric as the quotient of two user-based metrics. For example, we can express ABV as the quotient of the business metrics "mean spend per user (across all baskets)" and "mean number of baskets per user." Given that both metrics in the quotient are asymptotically normal and based on randomisation unit-based responses, which are commonly assumed to be i.i.d. [67], we can use the delta method to estimate the variance of the quotient. The resultant formula has an explicit term for the dependency between users and transactions/items.

Delta method - a straightforward formula with hidden complexity and inflexibility The delta method approach is more compute efficient than the bootstrap approach. It requires passing only group-level statistics around instead of analysis/sub-randomisation unit-level

responses required by bootstrap methods.²⁹ However, it is unclear how the results in [64], developed with a one-way dependency between randomisation and analysis units in mind, would apply to business metrics that feature responses dependent on other types of units.²⁸ The approach also requires considerable statistical mastery to understand and implement, risking decreased engagement from engineering teams and non-technical stakeholders due to general apprehensiveness towards “black box” algorithms.

Multi-way bootstrap may be overkill by being overly conservative The multi-way bootstrap described in Section 2.3 of [19] aims to address dependencies between different types of units. One may apply such an approach to item-based metrics (e.g., ASP) calculations, where dependence between items and products exists alongside dependence between items and users.

Recall from Section 5.5.2 that in a one-way reweighting bootstrap, we apply a random weight drawn from $\text{Poisson}(1)$ to each user and all their associated items before calculating a single bootstrap mean. In a two-way bootstrap, we apply a random weight to each user and a separate random weight to each product. The final weight for each item is the product of the user weight and the product weight.³⁰ Similar to the one-way bootstrap, the final weight emulates the number of times the item will be selected if we separately sample the users and products with replacement. The calculations for both the bootstrap means and SE remain identical post-reweighting.

As a two-way bootstrap preserves the correlation structure between items, users, and products during sampling, one would expect the method to generate less biased (if not unbiased) SE estimates. However, the SE estimates from a two-way bootstrap are impractically large (see Table 5.3). This appears to match the results reported in [19], where two-way bootstrap SE estimates generally yield CIs that are overly conservative in A/A tests.³¹ While the true

²⁹ That said, to calculate the group-level statistics, which include the covariance between users and transactions/items, one still requires *access* to sub-randomisation unit responses initially (see Section 4.4.1).

³⁰ In the general case where there are dependencies between multiple types of units, we apply a random weight to each unit within a single unit type, repeat for other unit types with a separate set of random weights, and take the final weight of each link/interaction/dependency between the unit types as the product of the weight of their associated units.

³¹ A/A tests are tests featuring two supposedly identical treatments and, thus, produce results that are consistent with a nil null hypothesis. See Section 3.2 for further discussions.

Table 5.3: The ratio between the bootstrap standard error (SE) estimates and the vanilla sample SE estimates for ASP. Each entry represents the mean and standard deviation of the ratios across multiple experiment runs.

Dataset	Duration	One-way bootstrap	Two-way bootstrap
ASOS	First 30d	3.51 ± 0.13	17.77 ± 0.56
	First 10%	64.45 ± 1.94	72.83 ± 2.20
UCI	First 50%	87.95 ± 2.82	105.15 ± 3.33
	All	69.82 ± 1.37	94.33 ± 2.28
Olist	First 20%	1.104 ± 0.035	2.083 ± 0.075
	First 50%	1.058 ± 0.036	2.531 ± 0.066
	All	1.076 ± 0.037	3.364 ± 0.109

SE will always be unknown, we believe overestimating it is just as bad as underestimating it, as experimenters will then struggle to design any experiments with sufficient power.

The analysis platform currently uses a one-way, user-based bootstrap to estimate the SE of e-commerce metrics. The approach mimics the randomisation process, and there is extensive evidence (both in [19] and within the organisation) that it produces CIs with the right coverage in A/A tests. The team responsible constantly evaluates whether a multi-way bootstrap will become necessary, especially for experiments with product-treatment interactions.³²

Further consideration - error vs runtime/cost tradeoff Bootstrap methods introduce their error (more precisely, *variance component*) to SE estimates during the resampling process. This is in addition to the (inflated) sampling uncertainty and generally scales along $O(1/\sqrt{B})$, where B is the number of bootstrap resamples [86]. Given the runtime of bootstrap methods scales along $O(B)$, at some point, it is no longer economically practical for a production system to draw further resamples to reduce the error to SE estimates.

Deciding where the tipping point is more an art than science and depends more on individual needs. Unlike many technology giants which perform analysis in a streaming fashion,

³² Section 4 of [19] shows that a one-way, user-based bootstrap may underestimate the SE of item-based metrics when there are product-treatment interactions. Note that “products” in this case study are referred to as “items” in [19] due to different business contexts.

our analysis platform features batch-processing pipelines due to temporal constraints from upstream data sources. As a result, our priority is to keep the processing jobs' runtime (and hence compute cost) low. The team finds $B \in [500, 1000]$ works well as it leads to a coefficient of variation (error of SE estimate divided by mean SE estimate) of $< 5\%$. That said, we believe the tradeoff also applies to streaming systems as the compute latency increases with a larger B , which impacts system availability.

5.6 Advances in Quasi-experiments and Natural Experiments

We move on to works related to quasi-experiments and natural experiments. While they differ regarding whether an experimenter can impose scientific control, they involve the same set of causal inference techniques to estimate the treatment effect. We thus combine the discussion on related methods and examples for the two experiment classes.

To date, few works are dedicated to reviewing the use of causal inference techniques in digital experimentation. Chapter 11 of [155] briefly outlined commonly used causal inference methods with examples, and [62] is a work-in-progress discussing its application in the online industry. That said, causal inference is an established discipline, and many excellent introductions and reviews exist:

- [122] and [131] provide an accessible introduction to the area and review some models and techniques. Both books stem from courses the respective authors taught in the past decades and are geared towards researchers and practitioners in the medical and social sciences, the latter sharing some scientific goals with digital experimentation.
- [216] focuses on the theoretical constructs (or representations) in causal inference – structural equation models, the potential outcomes framework (see Section 5.2 for the basics), and graphical models. The paper also proposed the structural causal model, which is seen as either a unifying or an alternative representation to those above.
- [290] extensively surveys the methods, datasets, software packages, and applications in (general) causal inference. Sections of note include the applications in advertising and recommendations (Sections 6.1–6.2). The paper is tailored to the computer science and

data mining communities, which substantially overlap with the digital experimentation community.

- [140] explores works in the intersection of causal inference and machine learning (CausalML), which has flourished in the past decade. At the same time, [267] outlines many of its applications in economics and marketing. Many experiments described in the latter are run offline, though most techniques can be directly translated to an online setting.

We keep our review close to digital experimentation and cover the following in the remainder of the section. We first revisit the conditional average treatment effect (CATE) and conditional exchangeability concepts in Section 5.6.1. Many causal inference methods seek to estimate the former and assume the latter. We then list several classes of popular methods in causal inference, followed by examples applying the said methods in digital experiments in Section 5.6.2. We finally provide a mini-spotlight in geo-experiments, a form of quasi-experiments that assign treatments based on the geographical region a user falls within, in Section 5.6.3.

5.6.1 CATE and Conditional Exchangeability Revisited

Many causal inference techniques seek to estimate the conditional average treatment effect (CATE), the average treatment effect among those sharing certain covariates $\mathcal{C}$ (those on which the treatment assignment $\mathcal{A}$ does not affect):³³

$$\mathbb{E}(Y(1) - Y(0) | \mathcal{C}) = \mathbb{E}(Y(1) | \mathcal{C}) - \mathbb{E}(Y(0) | \mathcal{C}). \quad (5.15)$$

The availability of covariates means one can produce a more expressive model $Y = g(\mathcal{A}, \mathcal{C})$ to capture the complex dynamics between the treatments and observed responses, achieved by drawing on the wealth of potentially complex statistical and machine learning models [137]. Of course, such a model may still only capture correlation/association instead of causation, as its dependent/target variables are not the potential responses.

³³ Here, we separate the covariates $\mathcal{C}$ and the treatment assignment indicator $\mathcal{A}$. Some in the causal inference community lump these two objects together when discussing CATE. Others may have a strict separation between these objects, with additional terminologies such as average treatment effect on the treated (ATT, i.e., CATE on $\mathcal{A} = 1$) and average treatment effect on the control/untreated (ATC/ATU, i.e. CATE on $\mathcal{A} = 0$) [110].

If, in addition, the covariates identify and account for all confounders, we noted that the potential responses are conditionally exchangeable, i.e., $Y(0), Y(1) \perp\!\!\!\perp \mathcal{A} \mid \mathcal{C}$.³⁴ Similar to its unconditional counterpart in Equation (5.6), if the potential responses are conditionally exchangeable, we have

$$\mathbb{E}(Y(a) \mid \mathcal{C}) = \mathbb{E}(Y(a) \mid \mathcal{A} = 1, \mathcal{C}) = \mathbb{E}(Y(a) \mid \mathcal{A} = 0, \mathcal{C}), \quad a \in \{0, 1\}. \quad (5.16)$$

This, combined with the definition of potential outcomes in Equation (5.3), enables us to estimate the potential responses using the observed responses Y conditioned on the covariates $\mathcal{C}$ and, thus, develop a model that captures the complex dynamics between the treatments and potential responses, i.e., causation.

At first glance, conditional exchangeability appears to be a huge assumption. This is mostly true but not insurmountable in practice. We live in a complex and dynamic environment; thus, it is difficult to identify and account for all confounders, which biases the treatment effect estimate as discussed in Section 5.2. On the other hand, many environmental variables have a tenuous link to our treatments and responses, and the bias they introduce may be practically insignificant. Identifying confounders is one of the main tasks in causal discovery and remains an active research topic, together with the study of the impact of hidden confounders [108, 217].

It is also worth emphasising that the decision to estimate CATE instead of ATE and the ability to perform random assignments are two orthogonal concepts. One can estimate CATE in randomised controlled trials (as in the case of heterogeneous treatment effect detection and uplift modelling discussed in Sections 5.4.1 and 4.3). They can also recover the population-level ATE from various CATE estimates by applying the law of total expectation:

$$\mathbb{E}(Y(1) - Y(0)) = \mathbb{E}_{\mathcal{C}}(\mathbb{E}(Y(1) - Y(0) \mid \mathcal{C})). \quad (5.17)$$

That said, it might be difficult in practice as the distribution of $\mathcal{C}$ is often unknown. While one can use the empirical distribution of $\mathcal{C}$, it may incur bias.

³⁴ Also known as *conditional unconfoundedness* or *strong ignorability* [62, 235, 240].

5.6.2 Methods and Examples

Given the underlying goal, we are in a position to list and briefly describe some popular causal inference techniques. [6] called the following approaches the “furious five” in econometrics, most of which are translatable to digital experiments:

- **Randomisation:** This has been extensively discussed in the previous sections.
- **(Outcome) regression:** A model that maps the relationship between the treatment assignment indicator, other covariates, and the responses. By manipulating the treatment assignment indicator within the fitted model, one can obtain the CATE for a particular covariate combination, provided the covariates present are sufficient for conditional exchangeability. See Chapter 9 of [100] for further discussions.
- **Instrumental variables:** Variables that only affect the responses via their effect on the treatment assignment indicator or other covariates (but not directly). They must also be independent of the error term. These variables are often employed when the treatment assignment indicator or other covariates correlate with the unobserved error terms in a regression model.
- **Regression discontinuity design:** Specifies a cut-off to a continuous eligibility index and assigns treatment based on which side of the cut-off an experimental unit falls on. By exploiting the assumption that there is no meaningful difference in experimental units’ characteristics or activity around the cut-off point before the treatment, one can estimate the local treatment effect as all others are roughly equal.
- **Difference-in-differences:** Calculates, per treatment, the difference in average response pre- and post-treatment and, then, another difference in the differences arising from different treatments. Under the parallel trend assumptions, i.e., the non-treatment variables affect the average response in both groups similarly across time, one is constructing a counterfactual when taking the per-treatment difference and estimating the treatment effect by comparing the counterfactual to the observed when taking the final difference.

Propensity scores To address hidden confounders between treatment and outcomes, some experimenters match the experimental units or reweigh the responses. This is usually done

via the *propensity score*, the probability that an experimental unit is assigned a certain treatment given the covariates associated with the unit [235]. The score is most straightforwardly estimated via a logistic regression model. Once calculated, it is used to match a unit receiving a treatment with other units receiving another treatment (known as *propensity score matching*) [235]. The estimated treatment effect is the difference in responses between the unit and its matched counterparts, averaged over all matches. Alternatively, one can reweigh the responses using the score and take the difference between the reweighted average response to obtain the treatment effect (known as *inverse probability/propensity weighting*) [121].

Another increasingly popular technique related to propensity scores is *doubly robust estimation* [98]. It combines outcome regression and propensity score models and yields a valuable property: the treatment effect estimates are consistent when at least one of the two models is correctly specified. This approach is taken by [287] to estimate the effect of changes to LinkedIn's mobile app, where the ability to perform random assignments is surrendered to the app store and users (as discussed in Section 5.3.2).

Mediation analyses Digital experimenters also turn to mediation analyses. They feature the use of mediators, variables directly influenced by the treatment assignment, and, in turn, directly influence the responses.³⁵ Using such intermediary variables enables one to attribute it to different causal pathways between the treatment assignment and change in responses. It also enables one to produce more accurate treatment effect estimates if the variables are unaffected by the same confounders that affect the treatment assignment indicator and responses. This is the case in [123], who, seeking to measure the effect of a display advert targeting algorithm on certain user behaviour, uses ad viewability³⁶ as a mediating event. Moreover, [274], in the quest to measure the impact of algorithmic changes to, e.g., predictive models or recommender systems, uses evaluation metrics from online A/B tests as media-

³⁵ Not to be confused with instruments. Explaining *why* and *how* a mediator and an instrument are different involves an introduction to graphical models in causal inference, which we have omitted for brevity. See, e.g., Chapters 2 & 3 of [218] for more details.

³⁶ Defined in [123] as whether at least half of the advert has come to a user's view for longer than a second. Treatment in display advertising is usually defined as serving an advert, i.e., the advert is planned/scheduled to appear on a user's screen as they scroll through the page. While such a definition leads to much easier tracking from an experimenter's perspective, the ad may not come into users' view for many reasons (e.g., the user stops scrolling or has the wrong browser dimensions). [123] observed that 45% of adverts have not come to view and hence cannot influence any change in user behaviour.

tors. They bridge the algorithmic changes' effect on evaluation metrics in offline experiments to that on business KPIs.

5.6.3 A mini spotlight on geo-experiments

A class of quasi-experiments gaining traction in the past decade is geo-experiments. They are heavily used in display advertising [20, 268] – the author also led a team at ASOS.com to build and maintain such a geo-experimentation platform to help its marketing function manage tens of millions of spend [17, 188].

A geo-experiment assigns treatments at a geographical region level. For example, it may assign all users based in Oxfordshire to control and all users based in Cambridgeshire to treatment. Such treatment assignment is necessitated by the inability to target and track individual users (due to privacy and data protection concerns or third-party platform restrictions) and, sometimes, by interference between individual users (see Section 5.4.1). The treatment assignment may or may not be random – while a randomised assignment will simplify the experiment to a straightforward cluster randomised OCE (see Section 5.4.1), experimenters may decline to do so due to various operational constraints (e.g., cross-region user activities, minimum advertising budget/coverage requirements).

We estimate the treatment effect by comparing the responses from regions receiving the treatment to a counterfactual, constructed using the responses from control regions, on how the same regions would have performed had there been no treatment. More precisely and in the language of potential outcomes, consider a simplified case where we have two regions, one with response X and the other with response Y . We call the former the control region and the latter the treatment region. We also assume we can access some covariates $\mathcal{C}$ shared between the two regions that address all confounding, i.e., conditional exchangeability $Y(0), Y(1) \perp\!\!\!\perp \mathcal{A} \mid \mathcal{C}$ holds.

Before the experiment-proper, we assign both regions to $\mathcal{A} = 0$ (i.e., control) and observe

$$X \mid \mathcal{A} = 0, \mathcal{C} \quad \text{and} \quad Y \mid \mathcal{A} = 0, \mathcal{C}. \quad (5.18)$$

We know from the definition of potential outcomes (see Section 5.2 / Equation (5.3)) that

Expression (5.18) is equivalent to

$$X(0) | \mathcal{A} = 0, \mathcal{C} \quad \text{and} \quad Y(0) | \mathcal{A} = 0, \mathcal{C}. \quad (5.19)$$

Suppose we have a mapping $g(\cdot)$ between the two observations under the same treatment assignment using the covariates, i.e., $Y = g(X, \mathcal{A} = 0, \mathcal{C})$. Given that the observations in Expressions (5.18) and (5.19) are equivalent, the following must also be true:

$$Y(0) = g(X, \mathcal{A} = 0, \mathcal{C}). \quad (5.20)$$

During the experiment, we assign the control region to $\mathcal{A} = 0$ and the treatment region to $\mathcal{A} = 1$ and observe³⁷

$$X | \mathcal{A} = 0, \mathcal{C} \quad \text{and} \quad Y | \mathcal{A} = 1, \mathcal{C}. \quad (5.21)$$

To estimate $Y(0)$, we require constructing a counterfactual using the observation X and the mapping $g(\cdot)$, i.e. $g(X, \mathcal{A} = 0, \mathcal{C}) | \mathcal{A} = 0, \mathcal{C}$. This enables us to take the expected difference between the observed Y and the counterfactual to obtain the CATE:

$$\begin{aligned} & \mathbb{E}(Y | \mathcal{A} = 1, \mathcal{C}) - \mathbb{E}(g(X, \mathcal{A} = 0, \mathcal{C}) | \mathcal{A} = 0, \mathcal{C}) \\ &= \mathbb{E}(Y | \mathcal{A} = 1, \mathcal{C}) - \mathbb{E}(Y(0) | \mathcal{A} = 0, \mathcal{C}) && \text{(RHS: By Eq. (5.20))} \\ &= \mathbb{E}(Y(1) | \mathcal{A} = 1, \mathcal{C}) - \mathbb{E}(Y(0) | \mathcal{A} = 0, \mathcal{C}) && \text{(LHS: By definition – Eq. (5.3))} \\ &= \mathbb{E}(Y(1) | \mathcal{C}) - \mathbb{E}(Y(0) | \mathcal{C}) && \text{(Conditional exchangeability – Eq. (5.16))} \\ &= \mathbb{E}(Y(1) - Y(0) | \mathcal{C}) && \text{(Linearity of conditional expectation).} \\ & && (5.22) \end{aligned}$$

The example above is grossly simplified. In reality, there are many copies of X and Y , both along the geographical (responses from multiple regions) and time dimensions (responses

³⁷ Readers may notice that it seems possible to observe responses under *both* treatments ($Y | \mathcal{A} = 0, \mathcal{C}$ in Expression (5.18) and $Y | \mathcal{A} = 1, \mathcal{C}$ in Expression (5.21)), which contradicts the premise of counterfactual reasoning. There is merely an artefact of oversimplifying the notations for illustration purposes – the two observed responses happened at different points in time and, thus, ought to be represented by different r.v.'s. Here, we pretend $Y | \mathcal{A} = 0, \mathcal{C}$ was never observed when we constructed the counterfactual.

at multiple time points). Given the many variables involved, experimenters often prefer using the more expressive structural equation model (instead of the potential outcomes framework) to express causal assumptions [20, 268].

Many models have been proposed for the mapping $g(\cdot)$, with the two most common being geo-based regression (GBR) and time-based regression (TBR) models. GBR uses the aggregated metric value of the regions before and during the experiment, plus the advertising spend as covariates [268, 269]. Thus, it allows experimenters to easily extract the return on ad spend (ROAS) metric from the fitted model. However, the statistical power of GBR is derived from the number of regions involved, which is problematic when the number of regions available for an experiment is small. The challenge is mitigated by TBR, which uses the value of the covariates at each time interval. TBR also employs Bayesian regression techniques, more precisely, Bayesian structural time series. This allows access to the full posterior distribution of the counterfactual and, thus, the calculation of a credible interval [39, 145].

That said, similar to many other Bayesian models, one needs to take care when specifying the prior distributions. The team at ASOS.com found during development that non-informative priors generate credible intervals (associated with the posterior distribution of the counterfactual response or, more generally, the posterior predictive) that are too narrow based on A/A tests.³¹ The team thus switched to informative priors that enabled them to calibrate the variance of the posterior predictive based on simulated experiments.

There are many other considerations in designing a geo-experiment. They include (1) the creation of regions and (2) the selection of regions into control and treatment groups. Often, regions are handcrafted based on expert input, though [232] showed it can be learnt from clustering past user queries using clickstream data. The optimal allocation of the regions into control and treatment groups is studied in [20] and [9], with the former focusing on finding the best region pair (i.e. one region each for control and treatment) and the latter finding the best disjoint subsets of regions available³⁸ using a hill-climbing algorithm.

Finally, an experimenter also needs to consider the duration for different stages of a geo-experiment [17, 188]. Unlike most digital OCEs, geo-experiments feature multiple stages:

³⁸ Consider the case where the experimenter can assign treatments over regions $\{1, 2, 3, \dots, 7\}$. The optimal allocation may assign $\{2, 4, 7\}$ to control and $\{3, 5\}$ to treatment, leaving $\{1, 6\}$ unused.

two – the period before the experiment where the data arising from two identical treatments is used to train the model and the period where the treatments differ – or more if an experimenter imposes one or more “cooldown” periods to avoid conflating effect from multiple treatments across time. The test power of geo-experiments is also less dependent on the sample size but more on the model fit, which reflects the variance of the counterfactual responses. This complicates power, sample size (or duration), and minimum detectable effect calculations. Often, one performs a grid search across all possible combinations of stage durations to determine the optimal experiment runtime.

5.7 A Brief Recap

We presented a brief review of recent advances in digital experimentation. These include advanced statistical methods that address the many sources of bias (e.g., post-selection inference, heterogeneous treatment effects, SUTVA violations) to experiments, improve experiment sensitivity, and produce treatment effect estimates using general causal inference techniques when one cannot perform a randomised controlled trial. To organise the research reported in the chapter, we have also introduced the potential outcomes framework and classified experiments into online randomised controlled trials, quasi-experiments, and natural experiments depending on the level of control an experimenter has on performing random assignments and imposing scientific control.

Interleaved between the reviews is an original case study on dependent responses in e-commerce OCEs, which arises as experimenters randomise by users to ensure a consistent user experience yet analyse by a more granular transaction- or item-based decision metric as is common in the industry. Using three real online retail/e-commerce transaction datasets, we provided evidence of the extent of the dependent responses in practice. We then highlighted its impact on test power and confidence interval coverage. Finally, we shared the lessons learned when incorporating relevant methods in the review above into the internal experiment analysis platform of ASOS.com, a global online fashion retail company. We also provided a spotlight on geo-experiments, an increasingly popular experimental design to measure the effect of display advertising.

Chapter 6

An Evaluation Framework for Personalisation Strategy Experiment Designs

This chapter is adapted from the research paper “*An Evaluation Framework for Personalization Strategy Experiment Designs*”, presented at and awarded Best Student Paper of *AdKDD 2020 Workshop (in conjunction with SIGKDD '20)* [176].

6.1 Motivation

Without mentioning experimental design, no discussion on statistical challenges in digital experimentation and measurement capabilities is complete. As alluded to in Section 5.3.4, in this chapter, we focus on the design of experiments that compare *personalisation strategies* – complex sets of targeted user interactions executed by e-commerce and digital marketing organisations that aim to create an individualised experience for every user to their websites. In addition to that mentioned in Section 5.3.4, examples of personalisation strategies include the scheduling, budgeting, and ordering of marketing activities directed at a user based on their browsing and purchase history.

Compared to one’s usual online controlled experiment (see Section 5.3.1 for examples), experiments for personalisation strategies face two unique challenges. Firstly, strategies are

often only applicable to a small fraction of the user base, leading to many simple experiment designs suffering from either a lack of test power due to low sample size or a diluted treatment effect due to the inclusion of irrelevant samples [63].

Secondly, users are not randomly assigned *a priori* but must qualify to be treated with a strategy via their actions or attributes. This leads to a one-sided non-compliance problem where experimenters do not have complete control over applying scientific control [135]. At best, they can randomise a user's eligibility to be treated and treat those who are both eligible and qualified [135, 270]. Under such a situation, groups of users subjected to different strategies cannot be assumed statistically equivalent and, hence, are not directly comparable.

While several variance reduction techniques, including stratification and control variates [69, 222] (see Section 5.4), partially address the challenges, the strata and control variates can vary dramatically from one personalisation strategy experiment to another, requiring many *ad hoc* adjustments. As a result, such techniques may not scale well when organisations design and run hundreds or thousands of experiments at any given time.

We argue that personalisation strategy experiments should focus on the assignment of users from the strategies they qualified for to the treatment/analysis groups. We call this mapping process an *experiment setup*. Identifying the best experiment setup increases the chance of detecting any treatment effect. An experimentation framework can also quickly reuse and switch between different setups with little custom input, ensuring the operation can scale. More importantly, the process does not hinder the subsequent application of variance reduction techniques, meaning we can still apply the techniques *post hoc* if required.

To date, many experiment setups exist to compare personalisation strategies. An increasingly popular approach is to compare the strategies using multiple control groups – Quantcast calls it a dual control [24], and Facebook calls it a multi-cell lift study [172, 270]. In the two-strategy case, this involves running two experiments on two random partitions of the user base in parallel, with each experiment further splitting the respective partition into treatment/control and measuring the *incrementality* (the change in a business metric as compared to the case where we do nothing) of each strategy. We then compare the incrementalities of the two strategies against each other.

Despite the setup above gaining traction in display advertising, there is a lack of literature

on whether it is a better setup – one with a higher sensitivity, presents a higher effect size, or both. While [172] noted that multi-cell lift studies require many users, they did not discuss how the number compares to other setups.¹ Identifying and adopting a better experiment setup can reduce the required sample size and enable more cost-effective experimentation.

We address the gap in the literature by introducing an evaluation framework that compares experiment setups given two personalisation strategies. The framework is designed to be flexible in dealing with a wide range of baselines and changes in user responses presented by any pair of strategies (*situations* hereafter). However, we also recognise the need to quickly compare typical setups and provide some simple rules of thumb for situations where one setup will be better. In particular, we outline the desirable conditions for employing treatment effect dilution and a multi-cell setup.

To summarise, our contributions are:

1. (Section 6.2) We develop a flexible evaluation framework for personalisation strategy experiments, where one can compare two experiment setups given the situation presented by two competing strategies;
2. (Section 6.3) We provide simple rules of thumb to enable experimenters who do not require the full flexibility of the framework to compare typical setups quickly; and
3. (Section 6.4) We make our results useful to practitioners by making the code that performs empirical verification publicly available.²

6.2 Evaluation Framework

We first present our evaluation framework for personalisation strategy experiments. The experiments compare two personalisation strategies, which we refer to as Strategy 1 and Strategy 2. Often, one is the existing strategy, and the other is a new strategy we intend to test and learn from. In this section, we introduce

¹ A single-cell lift study is often used to measure the incrementality of a single personalisation strategy and, hence, is not a representative comparison.

² The code is available on GitHub: https://github.com/liuchbryan/experiment_design_evaluation.




Figure 6.1: Venn diagram of the user groups in our evaluation framework. The outer, left inner (red), and right inner (blue) boxes represent the entire user base, those who qualify for Strategy 1, and those who qualify for Strategy 2, respectively.

1. How users qualifying themselves into strategies creates non-statistically equivalent groups,
2. How experimenters usually assign the users, and
3. When we consider an assignment to be better.

6.2.1 User grouping

As users qualify themselves into the two strategies, four disjoint groups emerge: those who qualify for neither strategy, those who qualify only for Strategy 1, those who qualify only for Strategy 2, and those who qualify for both strategies. We denote these groups (user) Groups 0, 1, 2, and 3, respectively (see Figure 6.1). It is perhaps obvious that we cannot assume those in different user groups are statistically equivalent and compare them directly.

We assume Groups 0, 1, 2, and 3 have n_0 , n_1 , n_2 , and n_3 users, respectively. We also assume responses from users (which we aggregate, often by taking the mean, to obtain our business metric) are distributed differently between groups, and within the same group, between the scenario where the group is subjected to the treatment associated with the corresponding strategy and where nothing is done (baseline). We list all group-scenario combinations in Table 6.1 and denote the mean and variance of the responses (μ_G, σ_G^2) for a combination G .³

³ For example, responses for Group 1 without any interventions have mean and variance $(\mu_{C1}, \sigma_{C1}^2)$, and that for Group 2 with the treatment prescribed under Strategy 2 have mean and variance $(\mu_{I2}, \sigma_{I2}^2)$.

Table 6.1: All group-scenario combinations in our evaluation framework for personalisation strategy experiments. The columns represent the groups described in Figure 6.1. The baseline represents the scenario where we do nothing. We assume those who qualify for both strategies (Group 3) can only receive treatment(s) associated with either strategy.

	Group 0	Group 1	Group 2	Group 3
Baseline (Control)	$C0$	$C1$	$C2$	$C3$
Under treatment (Intervention)	/	$I1$	$I2$	Under Strategy 1: $I\phi$ Under Strategy 2: $I\psi$




Figure 6.2: Experiment setups overlaid on the user grouping Venn diagram in Figure 6.1. The hatched boxes cover users included in the analysis, and the downward triangles and dots cover users subjected to the treatment(s) prescribed under Strategies 1 and 2, respectively. See Section 6.2.2 for a detailed description.

6.2.2 Experiment setups

Many experiment setups exist and are in use in different organisations. Here, we introduce four typical setups of various sophistication, which we also illustrate in Figure 6.2.

Setup 1 (Users in the intersection only) The setup considers only users who qualify for both strategies. The said users are randomly split (usually 50/50) into two (analysis) groups, A and B , and are prescribed the treatment specified by Strategies 1 and 2, respectively. The setup is easy to implement, though it is difficult to translate any learnings from the experiment to other user groups (e.g., those who qualify for one strategy only) [74].

Setup 2 (All samples) The setup is a simple A/B test that considers all users, regardless of whether they qualify for any strategy. The users are randomly split into two analysis groups, A and B , and are prescribed the treatment specified by Strategy 1(2) if (i) they qualify under the strategy and (ii) they are in Group $A(B)$. This setup is easiest to implement but usually suffers severely from dilution in treatment effect [63].

Setup 3 (Qualified users only) The setup is similar to Setup 2, except only those who qualified for at least one strategy (“triggered” users in some literature [63]) are included in the analysis groups. The setup sits between Setup 1 and Setup 2 in terms of user coverage. It has the advantage of capturing the greatest number of relevant samples yet having the least treatment effect dilution. However, the setup also prevents one from telling the incrementality of a strategy itself, but only the difference in incrementalities between the two strategies.

Setup 4 (Dual control / multi-cell lift test) As described in Section 6.1, the setup first split the users randomly into two randomisation groups. For the first randomisation group, we consider those who qualify for Strategy 1 and split them into Analysis Groups $A1$ and $A2$. Group $A2$ receives the treatment prescribed under Strategy 1, and Group $A1$ acts as a control. The incrementality for Strategy 1 is then the difference in the business metric between Groups $A2$ and $A1$. Next, We apply the same process to the second randomisation group, with Strategy 2 and Analysis Groups $B1$ and $B2$ in place. Finally, we compare the incrementality for Strategies 1 and 2. The setup allows one to obtain the incrementality of each strategy

and minimises treatment effect dilution. However, it also leaves many samples unused and creates additional analysis groups, and hence generally suffers from a low test power [172].

6.2.3 Evaluation criteria

There are several considerations when one evaluates competing experiment setups. These include technical considerations, such as the complexity of setting up the setups on an OCE framework, and business considerations, such as whether the incrementality of individual strategies is required.

Here, we focus on the statistical aspect and propose two evaluation metrics:

1. The actual average treatment effect size as presented by the two analysis groups in an experiment setup, and
2. The sensitivity of the experiment setup, represented by the minimum detectable effect (MDE) under a pre-specified test power.

Both evaluation metrics are necessary: the former indicates whether a setup suffers from treatment effect dilution, whereas the latter indicates whether the setup suffers from a lack of power or sample size.⁴ An ideal setup should yield a high actual effect size and a high sensitivity (i.e., a low MDE),⁵ though as we observe in the next section, it is usually a trade-off.

We formally define the two evaluation metrics. Let A and B be the two analysis groups in an experiment setup, with user responses randomly distributed with independent mean and variance (μ_A, σ_A^2) and (μ_B, σ_B^2) , respectively. We also assume the sample size, namely n_A and n_B for Groups A and B , respectively, is sufficient such that the mean response of the two groups approximately follows the normal distribution by the central limit theorem. This enables us to run a two-sample approximate z -test.

⁴ The evaluation metric on experiment sensitivity pertains to the *efficiency* of a statistical test [254]. However, existing literature on this topic often compares the minimum sample size required by competing statistical tests instead. As shown earlier in Section 3.6, under a given significance level α and power threshold $\pi_{\min}$, fixing the effect size will determine the sample size required and vice versa. Therefore, we opt to compare the MDE of competing personalisation strategy experiment designs because (1) it is generally easier to estimate the sample size and (2) it is on the same scale as the other evaluation metric.

⁵ We will use the terms “high(er) sensitivity” and “low(er) MDE” interchangeably.

The actual effect size is simply the difference in the mean of the two analysis groups:

$$\Delta \triangleq \mu_B - \mu_A. \quad (6.1)$$

It is worth noting that, unlike the means and variances of the user groups, the response means of the analysis groups, thus Δ , vary depending on the experiment setup. We will demonstrate this property in Section 6.3.1.

As covered in Section 3.6.6, the MDE of a two-sample, two-sided z -test is

$$\theta^* \triangleq \left(z_{1-\frac{\alpha}{2}} - z_{1-\pi_{\min}} \right) \sqrt{\frac{\sigma_A^2}{n_A} + \frac{\sigma_B^2}{n_B}}. \quad (6.2)$$

We finally define what it means to be better under these evaluation metrics. WLOG, we assume the actual effect sizes of the two competing experiment setups are positive⁶ and say a setup S is superior to another setup R if, all else being equal,

1. S produces a higher actual effect size ($\Delta_S > \Delta_R$) and a lower minimum detectable effect size ($\theta_S^* < \theta_R^*$), or
2. The gain in actual effect is greater than the loss in sensitivity:

$$\Delta_S - \Delta_R > \theta_S^* - \theta_R^*, \quad (6.3)$$

which means an actual effect still stands a higher chance to be observed under S .

6.3 Comparing Experiment Setups

Having described the evaluation framework above, in this section, we use the framework to compare the typical experiment setups described in Section 6.2.2. We will first derive the actual effect size and MDE for each setup in Section 6.3.1 and use the result to create rules of thumb on (1) whether diluting the treatment effect by including users who qualify

⁶ We swap the analysis groups if both actual effect sizes are negative. It is likely an error if the actual effect sizes are of opposite signs.

for neither strategy is beneficial (Section 6.3.2) and (2) if dual control is a better setup for personalisation strategy experiments (Section 6.3.3), two questions that are often discussed among e-commerce and marketing-focused experimenters.

6.3.1 Actual and minimum detectable effect sizes

We first present the four experiment setups' actual effect size and MDE. For each setup, we first compute each analysis group's sample size, mean response, and response variance, which arises as a mixture of user groups described in Section 6.2.1. We then substitute the quantities computed into the definitions of Δ (see Equation (6.1)) and θ^* (see Equation (6.2)) to obtain the setup-specific actual effect size and MDE. We assume all random splits are done 50/50 in these setups to maximise the test power.

Setup 1 (Users in the intersection only) We recall the setup, which considers only users who qualify for both personalisation strategies (i.e. the intersection), randomly splits User Group 3 into two analysis groups, A and B , each with the following number of samples:

$$n_A = n_B = \frac{n_3}{2}.$$

Users in Analysis Group A are provided treatment under Strategy 1, and users in Analysis Group B are provided treatment under Strategy 2. This leads to the groups' responses having the following mean and variance:

$$\mu_A = \mu_{I\phi}, \quad \mu_B = \mu_{I\psi}, \quad \sigma_A^2 = \sigma_{I\phi}^2, \quad \sigma_B^2 = \sigma_{I\psi}^2.$$

The actual effect size and MDE for Setup 1 are, hence,

$$\Delta_{S1} = \mu_{I\psi} - \mu_{I\phi}, \tag{6.4}$$

$$\theta_{S1}^* = (z_{1-\frac{\alpha}{2}} - z_{1-\pi_{\min}}) \sqrt{\frac{\sigma_{I\phi}^2}{\frac{n_3}{2}} + \frac{\sigma_{I\psi}^2}{\frac{n_3}{2}}}. \tag{6.5}$$

Setup 2 (All samples) This setup also contains two analysis groups, A and B , each taking half of all available users (regardless of whether they qualify for any strategy):

$$n_A = n_B = \frac{n_0 + n_1 + n_2 + n_3}{2}.$$

The mean response and response variance for Groups A and B are the weighted mean response and response variance of the constituent user groups, respectively, weighted by the constituent groups' size. As we only provide treatment to those who qualify for Strategy 1 in Group A and those who qualify for Strategy 2 in Group B , this leads to different responses in different constituent user groups:

$$\begin{aligned}\mu_A &= \frac{n_0\mu_{C0} + n_1\mu_{I1} + n_2\mu_{C2} + n_3\mu_{I\phi}}{n_0 + n_1 + n_2 + n_3}, \quad \mu_B = \frac{n_0\mu_{C0} + n_1\mu_{C1} + n_2\mu_{I2} + n_3\mu_{I\psi}}{n_0 + n_1 + n_2 + n_3}; \\ \sigma_A^2 &= \frac{n_0\sigma_{C0}^2 + n_1\sigma_{I1}^2 + n_2\sigma_{C2}^2 + n_3\sigma_{I\phi}^2}{n_0 + n_1 + n_2 + n_3}, \quad \sigma_B^2 = \frac{n_0\sigma_{C0}^2 + n_1\sigma_{C1}^2 + n_2\sigma_{I2}^2 + n_3\sigma_{I\psi}^2}{n_0 + n_1 + n_2 + n_3}.\end{aligned}$$

We then substitute the above into the actual effect size and MDE definitions and simplify the resultant expressions to obtain

$$\Delta_{S2} = \frac{n_1(\mu_{C1} - \mu_{I1}) + n_2(\mu_{I2} - \mu_{C2}) + n_3(\mu_{I\psi} - \mu_{I\phi})}{n_0 + n_1 + n_2 + n_3}, \quad (6.6)$$

$$\theta_{S2}^* = (z_{1-\frac{\alpha}{2}} - z_{1-\pi_{\min}}) \sqrt{\frac{2(n_0(2\sigma_{C0}^2) + n_1(\sigma_{I1}^2 + \sigma_{C1}^2) + n_2(\sigma_{C2}^2 + \sigma_{I2}^2) + n_3(\sigma_{I\phi}^2 + \sigma_{I\psi}^2))}{(n_0 + n_1 + n_2 + n_3)^2}}. \quad (6.7)$$

Setup 3 (Qualified users only) The setup is very similar to Setup 2, with members from User Group 0 excluded:

$$n_A = n_B = \frac{n_1 + n_2 + n_3}{2}.$$

The absence of members from User Group 0 means they are not featured in the weighted mean response and response variance of the two analysis groups:

$$\begin{aligned}\mu_A &= \frac{n_1\mu_{I1} + n_2\mu_{C2} + n_3\mu_{I\phi}}{n_1 + n_2 + n_3}, \quad \mu_B = \frac{n_1\mu_{C1} + n_2\mu_{I2} + n_3\mu_{I\psi}}{n_1 + n_2 + n_3}; \\ \sigma_A^2 &= \frac{n_1\sigma_{I1}^2 + n_2\sigma_{C2}^2 + n_3\sigma_{I\phi}^2}{n_1 + n_2 + n_3}, \quad \sigma_B^2 = \frac{n_1\sigma_{C1}^2 + n_2\sigma_{I2}^2 + n_3\sigma_{I\psi}^2}{n_1 + n_2 + n_3}.\end{aligned}$$

This leads to the following actual effect size and MDE for Setup 3:

$$\Delta_{S3} = \frac{n_1(\mu_{C1} - \mu_{I1}) + n_2(\mu_{I2} - \mu_{C2}) + n_3(\mu_{I\psi} - \mu_{I\phi})}{n_1 + n_2 + n_3}, \quad (6.8)$$

$$\theta_{S3}^* = (z_{1-\frac{\alpha}{2}} - z_{1-\pi_{\min}}) \sqrt{\frac{2(n_1(\sigma_{I1}^2 + \sigma_{C1}^2) + n_2(\sigma_{C2}^2 + \sigma_{I2}^2) + n_3(\sigma_{I\phi}^2 + \sigma_{I\psi}^2))}{(n_1 + n_2 + n_3)^2}}. \quad (6.9)$$

Setup 4 (Dual control) Setup 4 is unique amongst the experiment setups introduced as it has four analysis groups. Two analysis groups ($A1$ and $A2$) are drawn from those who qualified for Strategy 1 and are allocated into the first randomisation group. The other two ($B1$ and $B2$) are drawn from those who qualified for Strategy 2 and are allocated into the second randomisation group:

$$n_{A1} = n_{A2} = \frac{n_1 + n_3}{4}, \quad n_{B1} = n_{B2} = \frac{n_2 + n_3}{4}.$$

Each analysis group's mean response and response variance are the weighted mean response and response variance of the user groups involved, respectively:

$$\begin{aligned} \mu_{A1} &= \frac{n_1\mu_{C1} + n_3\mu_{C3}}{n_1 + n_3}, \quad \mu_{A2} = \frac{n_1\mu_{I1} + n_3\mu_{I\phi}}{n_1 + n_3}, \quad \mu_{B1} = \frac{n_2\mu_{C2} + n_3\mu_{C3}}{n_2 + n_3}, \quad \mu_{B2} = \frac{n_2\mu_{I2} + n_3\mu_{I\psi}}{n_2 + n_3}; \\ \sigma_{A1}^2 &= \frac{n_1\sigma_{C1}^2 + n_3\sigma_{C3}^2}{n_1 + n_3}, \quad \sigma_{A2}^2 = \frac{n_1\sigma_{I1}^2 + n_3\sigma_{I\phi}^2}{n_1 + n_3}, \quad \sigma_{B1}^2 = \frac{n_2\sigma_{C2}^2 + n_3\sigma_{C3}^2}{n_2 + n_3}, \quad \sigma_{B2}^2 = \frac{n_2\sigma_{I2}^2 + n_3\sigma_{I\psi}^2}{n_2 + n_3}. \end{aligned}$$

As the setup takes the difference of differences in the mean responses (i.e., the difference between the mean response for Groups $B2$ and $B1$ and the difference between the mean response for Groups $A2$ and $A1$),⁷ the actual effect size is specified, post-simplification, as follows:

$$\begin{aligned} \Delta_{S4} &= (\mu_{B2} - \mu_{B1}) - (\mu_{A2} - \mu_{A1}) \\ &= \frac{n_2(\mu_{I2} - \mu_{C2}) + n_3(\mu_{I\psi} - \mu_{C3})}{n_2 + n_3} - \frac{n_1(\mu_{I1} - \mu_{C1}) + n_3(\mu_{I\phi} - \mu_{C3})}{n_1 + n_3}. \end{aligned} \quad (6.10)$$

The MDE for Setup 4 is similar to that specified in the RHS of Equation (6.2), albeit with

⁷ Not to be confused with the difference-in-differences method, which captures the mean responses for the control and treatment groups at both the experiment's beginning (pre-intervention) and end (post-intervention). Here, we only capture the mean responses once at the end of the experiment.

more analysis groups featured in the square root term:

$$\begin{aligned}\theta_{S4}^* &= (z_{1-\frac{\alpha}{2}} - z_{1-\pi_{\min}}) \sqrt{\frac{\sigma_{A1}^2}{n_{A1}} + \frac{\sigma_{A2}^2}{n_{A2}} + \frac{\sigma_{B1}^2}{n_{B1}} + \frac{\sigma_{B2}^2}{n_{B2}}} \\ &= 2 \cdot (z_{1-\frac{\alpha}{2}} - z_{1-\pi_{\min}}) \sqrt{\frac{n_1(\sigma_{C1}^2 + \sigma_{I1}^2) + n_3(\sigma_{C3}^2 + \sigma_{I\phi}^2)}{(n_1 + n_3)^2} + \frac{n_2(\sigma_{C2}^2 + \sigma_{I2}^2) + n_3(\sigma_{C3}^2 + \sigma_{I\psi}^2)}{(n_2 + n_3)^2}}.\end{aligned}\quad (6.11)$$

6.3.2 Is dilution always bad?

The use of responses from users who do not qualify for any of the strategies we are comparing, an act known as treatment effect dilution, has stirred countless debates in experimentation teams. On the one hand, responses from these users make any treatment effect less pronounced by contributing exactly zero; on the other hand, it might be necessary as one does not know who actually qualifies for a strategy [172], or it might be desirable as they can be leveraged to reduce the variance of the treatment effect estimator [63]. Here, we are interested in whether we should engage in the act of dilution given the assumed user responses prior to an experiment. This can be clarified by understanding the conditions where Setup 3 would emerge superior (as defined in Section 6.2.3) to Setup 2. We relegate most of the intermediate algebraic work to Appendix A.2 for brevity.

By inspecting Equations (6.6) and (6.8), it is clear that $\Delta_{S3} > \Delta_{S2}$ if $n_0 > 0$. Thus, Setup 3 is superior to Setup 2 under the first criterion if $\theta_{S3}^* < \theta_{S2}^*$, which is the case if σ_{C0}^2 , the response variance of users who qualify for neither strategy, is large. This can be shown by substituting Equations (6.7) and (6.9) into $\theta_{S3}^* < \theta_{S2}^*$ and rearranging the terms to obtain

$$\frac{(n_1(\sigma_{I1}^2 + \sigma_{C1}^2) + n_2(\sigma_{C2}^2 + \sigma_{I2}^2) + n_3(\sigma_{I\phi}^2 + \sigma_{I\psi}^2)) \cdot (n_0 + 2n_1 + 2n_2 + 2n_3)}{2(n_1 + n_2 + n_3)^2} < \sigma_{C0}^2. \quad (6.12)$$

If we assume the response variances are similar across groups with users who qualified for at least one strategy, i.e. $\sigma_{I1}^2 \approx \sigma_{C1}^2 \approx \dots \approx \sigma_{I\psi}^2 \approx \sigma_G^2$, Inequality (6.12) can then be simplified as

$$\sigma_G^2 \left(\frac{n_0}{n_1 + n_2 + n_3} + 2 \right) < \sigma_{C0}^2, \quad (6.13)$$

which can be used quickly to determine if one should consider dilution.

If Inequality (6.12) does not hold (i.e., $\theta_{S3}^* \geq \theta_{S2}^*$), we should then consider when the second criterion (i.e., $\Delta_{S3} - \Delta_{S2} > \theta_{S3}^* - \theta_{S2}^*$) holds. Writing

$$\begin{aligned}\eta &= n_1(\mu_{C1} - \mu_{I1}) + n_2(\mu_{I2} - \mu_{C2}) + n_3(\mu_{I\psi} - \mu_{I\phi}), \\ \xi &= n_1(\sigma_{C1}^2 + \sigma_{I1}^2) + n_2(\sigma_{I2}^2 + \sigma_{C2}^2) + n_3(\sigma_{I\psi}^2 + \sigma_{I\phi}^2), \text{ and} \\ z &= z_{1-\frac{\alpha}{2}} - z_{1-\pi_{\min}},\end{aligned}\tag{6.14}$$

we can substitute Equations (6.6), (6.7), (6.8), and (6.9) into $\Delta_{S3} - \Delta_{S2} > \theta_{S3}^* - \theta_{S2}^*$ and rearrange the terms to obtain

$$\frac{n_1 + n_2 + n_3}{n_0} \sqrt{2n_0\sigma_{C0}^2 + \xi} > \frac{n_0 + n_1 + n_2 + n_3}{n_0} \sqrt{\xi} - \frac{\eta}{\sqrt{2}z}.\tag{6.15}$$

As the LHS of Inequality (6.15) is always positive, Setup 3 is superior if the RHS ≤ 0 . Noting

$$\Delta_{S3} = \frac{\eta}{n_1 + n_2 + n_3} \text{ and } \theta_{S3}^* = \frac{\sqrt{2} \cdot z \cdot \sqrt{\xi}}{n_1 + n_2 + n_3},$$

the trivial case is satisfied if

$$\frac{n_0 + n_1 + n_2 + n_3}{n_0} \theta_{S3}^* \leq \Delta_{S3}.\tag{6.16}$$

If the RHS of Inequality (6.15) is positive, we can safely square both sides and use the identities for Δ_{S3} and θ_{S3}^* to get

$$\frac{2\sigma_{C0}^2}{n_0} > \frac{\left(\theta_{S3}^* - \Delta_{S3} + \frac{n_1+n_2+n_3}{n_0}\theta_{S3}^*\right)^2 - \left(\frac{n_1+n_2+n_3}{n_0}\theta_{S3}^*\right)^2}{2z^2}.\tag{6.17}$$

As the LHS is always positive, the second criterion is met if

$$\theta_{S3}^* \leq \Delta_{S3}.\tag{6.18}$$

Note that this is a weaker, and thus more easily satisfiable condition than that introduced in Inequality (6.16). This suggests that an experiment setup is always superior to a diluted

alternative if the experiment is already adequately powered – introducing any dilution will make things worse.

Failing the condition in Inequality (6.18), we can always fall back to Inequality (6.17). While the inequality operates in squared space, it is essentially comparing the standard error of User Group 0 (LHS) – those who qualify for neither strategy – to the gap between the minimum detectable and actual effects ($\theta_{S3}^* - \Delta_{S3}$). The gap can be interpreted as the current noise level. Thus, a higher standard error means mixing in Group 0 users will introduce extra noise, and one is better off without them. Conversely, a smaller standard error means Group 0 users can lower the noise level, i.e. stabilise the treatment effect fluctuation, and one should take advantage of them.

To summarise, treatment effect dilution in a personalisation strategy experiment setup is *not* helpful if

1. Users who do not qualify for any strategies have a large response variance (see Inequality (6.12)), or
2. The experiment is already adequately powered (see Inequality (6.18)).

It could help if the experiment has yet to gain sufficient power and users who do not qualify for any strategy provide low-variance responses (i.e., when the complement of Inequality (6.17) holds), such that they exhibit stabilising effects when included in the analysis.

6.3.3 When is dual control more effective?

Often, when advertisers compare two personalisation strategies, the question of whether to use a dual control/multi-cell design comes up. Proponents of such an approach celebrate its ability to tell a story by making the incrementality of an individual strategy available, while opponents voice concerns about the complexity of setting up the design. Here, we are interested in whether Setup 4 (dual control) is superior to Setup 3 (a simple A/B test) under the prescribed evaluation criteria, and if so, under what circumstances.

We first observe $\theta_{S4}^* > \theta_{S3}^*$ always holds; hence, a dual control setup will never be superior to the more straightforward under the first criterion. This can be verified by substituting in

Equations (6.11) and (6.9) and rearranging the terms to show $\theta_{S4}^* > \theta_{S3}^*$ is equivalent to

$$\begin{aligned}
& 2 \left(\frac{n_1}{(n_1 + n_3)^2} (\sigma_{C1}^2 + \sigma_{I1}^2) + \frac{n_2}{(n_2 + n_3)^2} (\sigma_{C2}^2 + \sigma_{I2}^2) + \right. \\
& \quad \left. \frac{n_3}{(n_1 + n_3)^2} \sigma_{I\phi}^2 + \frac{n_3}{(n_2 + n_3)^2} \sigma_{I\psi}^2 + \left(\frac{n_3}{(n_1 + n_3)^2} + \frac{n_3}{(n_2 + n_3)^2} \right) \sigma_{C3}^2 \right) \\
& > \frac{n_1}{(n_1 + n_2 + n_3)^2} (\sigma_{C1}^2 + \sigma_{I1}^2) + \frac{n_2}{(n_1 + n_2 + n_3)^2} (\sigma_{C2}^2 + \sigma_{I2}^2) + \\
& \quad \frac{n_3}{(n_1 + n_2 + n_3)^2} \sigma_{I\phi}^2 + \frac{n_3}{(n_1 + n_2 + n_3)^2} \sigma_{I\psi}^2, \tag{6.19}
\end{aligned}$$

which, given all n -terms are non-negative and σ^2 -terms are positive, always holds : not only the coefficients of the σ^2 -terms are larger on the LHS than their RHS counterparts, the LHS also carries an additional σ_{C3}^2 term with non-negative coefficient and a factor of two.

Moving on to the second evaluation criterion, we recall that Setup 4 is superior if $\Delta_{S4} - \Delta_{S3} > \theta_{S4}^* - \theta_{S3}^*$. Otherwise, Setup 3 is superior under the same criterion. We can see the full flexibility of the model by substituting Equations (6.8), (6.9), (6.10), and (6.11) into $\Delta_{S4} - \Delta_{S3} > \theta_{S4}^* - \theta_{S3}^*$ and rearranging the terms to obtain

$$\begin{aligned}
& \frac{n_1 \frac{n_2(\mu_{I2} - \mu_{C2}) + n_3(\mu_{I\psi} - \mu_{C3})}{n_2 + n_3} - n_2 \frac{n_1(\mu_{I1} - \mu_{C1}) + n_3(\mu_{I\phi} - \mu_{C3})}{n_1 + n_3}}{\sqrt{n_1(\sigma_{C1}^2 + \sigma_{I1}^2) + n_2(\sigma_{C2}^2 + \sigma_{I2}^2) + n_3(\sigma_{I\phi}^2 + \sigma_{I\psi}^2)}} \\
& > \sqrt{2}z \left[\sqrt{2 \cdot \frac{\left(1 + \frac{n_2}{n_1 + n_3}\right)^2 [n_1(\sigma_{C1}^2 + \sigma_{I1}^2) + n_3(\sigma_{C3}^2 + \sigma_{I\phi}^2)] + \left(1 + \frac{n_1}{n_2 + n_3}\right)^2 [n_2(\sigma_{C2}^2 + \sigma_{I2}^2) + n_3(\sigma_{C3}^2 + \sigma_{I\psi}^2)]}{n_1(\sigma_{C1}^2 + \sigma_{I1}^2) + n_2(\sigma_{C2}^2 + \sigma_{I2}^2) + n_3(\sigma_{I\phi}^2 + \sigma_{I\psi}^2)}} - 1 \right], \tag{6.20}
\end{aligned}$$

where $z = z_{1-\frac{\alpha}{2}} - z_{1-\pi_{\min}}$.

A key observation from inspecting Inequality (6.20) is that the LHS of the inequality scales along $O(\sqrt{n})$, where n is the number of users, while the RHS remains constant. This leads to the insight that Setup 4 is more likely to be superior if the n -terms are large. Here, we assume the ratio $n_1 : n_2 : n_3$ remains unchanged when we scale the number of samples, an assumption that generally holds when an organisation increases their reach while maintaining its user mix. It is worth pointing out that our claim is stronger than that in previous work – we have shown that having a large user base not only fulfils the requirement of running a dual control experiment as described in [172], but it also makes a dual control experiment a

better setup than its simpler counterparts in terms of actual and minimum detectable effect sizes.

We can see the scaling relationship more clearly by simplifying the σ^2 - and n -terms. If we assume the response variances are similar across user groups (i.e., $\sigma_{C1}^2 \approx \sigma_{I1}^2 \approx \dots \approx \sigma_{I\psi}^2 \approx \sigma_G^2$), the RHS of Inequality (6.20) becomes

$$\sqrt{2}z \left[\sqrt{\frac{n_1 + n_2 + n_3}{n_1 + n_3} + \frac{n_1 + n_2 + n_3}{n_2 + n_3}} - 1 \right], \quad (6.21)$$

which remains a constant if the ratio $n_1 : n_2 : n_3$ remains unchanged. Separately, if we assume the number of users in Groups 1, 2, and 3 is similar (i.e., $n_1 \approx n_2 \approx n_3 \approx n$), the LHS of Inequality (6.20) becomes

$$\frac{\sqrt{n}((\mu_{I2} - \mu_{C2}) - (\mu_{I1} - \mu_{C1}) + \mu_{I\psi} - \mu_{I\phi})}{2\sqrt{\sigma_{C1}^2 + \sigma_{I1}^2 + \sigma_{C2}^2 + \sigma_{I2}^2 + \sigma_{I\phi}^2 + \sigma_{I\psi}^2}}, \quad (6.22)$$

which clearly scales along $O(\sqrt{n})$.

We conclude the section by indicating what a large n may look like. Suppose the response variances and the number of users are similar across user groups. In that case, we can rearrange Inequality (6.20) to make n the subject:

$$n > \left(2\sqrt{12} \left(\sqrt{6} - 1 \right) z \right)^2 \frac{\sigma_G^2}{\Delta^2}, \quad (6.23)$$

where $\Delta = (\mu_{I2} - \mu_{C2}) - (\mu_{I1} - \mu_{C1}) + \mu_{I\psi} - \mu_{I\phi}$ is the difference in actual effect sizes between Setups 4 and 3. Under a 5% significance level and 80% power, the first coefficient amounts to around 791, roughly 50 times the coefficient one would use to determine the sample size of a simple A/B test [194]. This suggests that a dual control setup is perhaps a luxury accessible only to the largest advertising platforms and their top advertisers. For example, consider an experiment to optimise the conversion rate where the baselines attain 20% (hence having a variance of $0.2(1 - 0.2) = 0.16$). If there is a 2.5% relative (i.e. 0.5% absolute) effect between the competing strategies, the dual control setup will only be superior if there are > 5 M users in each user group.

6.4 Empirical Verification

Having performed theoretical calculations for the actual and detectable effects and conditions where an experiment setup is superior to another, we verify those calculations using simulation results. We focus on the results presented in Section 6.3.1, as the rest of the results presented followed those calculations.

In each experiment setup evaluation, we randomly select the value of the μ - and σ -parameters from a uniform distribution,⁸ with the bounds chosen that are reflective of what advertisers usually see with decision metrics based on binary and count responses.⁹ We also select the n -parameters uniformly on a log scale, with the bounds chosen to represent usual user numbers in personalisation strategy experiments and to ensure the empirical verification runs within a practical timeframe.¹⁰ We favour sampling user numbers uniformly on a log scale (instead of a linear scale) due to the diminishing real-life experiment count and diminishing impact of having a single additional user as we increase the magnitude of user numbers. Each parameter is drawn independently from other parameters within the same evaluation and across different evaluations. We fix $\alpha = 5\%$ and $\pi_{\min} = 80\%$ across all evaluations as per the norm in digital experiments.

We then take 1,000 actual effect samples, each by

1. Sampling normally distributed responses from the user groups under the specified parameters,⁹
2. Computing the mean for the analysis groups, and
3. Taking the difference of the means.

We also take 100 MDE samples in separate evaluations, each by

1. Sampling a critical value under the null hypothesis,

⁸ Denoting $U(a, b)$ as a uniform distribution with bounds a and b , we draw $\mu_G \sim U(-10, 10)$ and $\sigma_G^2 \sim U(1, 10)$ for $G \in \{C0, C1, C2, C3, I1, I2, I\phi, I\psi\}$ (see Table 6.1 for how the groups are defined).

⁹ This is done to rule out any CLT convergence issues due to heavy-tailed distributed responses or low sample size. We acknowledge the limitations of such an approach in Section 7.2.

¹⁰ We draw $n_0, n_1, n_2, n_3 \sim 5 \times 10^{U(1, 3.5)}$.

Table 6.2: Number of evaluations where the theoretical value of the quantities (columns) falls between the 95% bootstrap confidence interval for each experiment setup (rows). See Section 6.4 for a detailed description of the empirical verification.

	Actual effect size	Minimum detectable effect
Setup 1	1049/1099 (95.45%)	66/81 (81.48%)
Setup 2	853/999 (85.38%)	87/106 (82.08%)
Setup 3	922/1099 (83.89%)	93/116 (80.18%)
Setup 4	240/333 (72.07%)	149/185 (80.54%)

2. Computing the test power under a large number of possible effect sizes, each using the critical value and sampled response means under the alternate hypothesis, and
3. Searching the effect size space for the value that gives the predefined power.

As the power vs effect size curve is noisy, given the use of simulated power samples, we use the bisection algorithm provided by the `noisyopt` package to perform the search. The algorithm dynamically adjusts the number of samples taken from the same point on the curve to ensure the noise does not send us down the wrong search space.

We expect the mean of the sampled actual effect and MDE to match the theoretical value. To verify this, we perform 1,000 bootstrap resamplings on the above samples to obtain an empirical bootstrap distribution of the sample mean in each evaluation. The 95% bootstrap resampling confidence interval (BRCI) should then contain the theoretical mean 95% of the time. Furthermore, the histogram of the theoretical quantity's percentile rank, in relation to the bootstrap samples, across multiple evaluations should also follow a uniform distribution [259].

The result is shown in Table 6.2. One can observe that more evaluations have their theoretical quantity lying outside the BRCI than expected. Upon further investigation, we observed a characteristic U-shape from the histograms of the percentile ranks for the actual effects. This suggests that the bootstrap samples may be under-dispersed but otherwise centred on the theoretical quantities.

We also observed the histograms for MDEs curving upward to the right. This suggests that

the theoretical value is a slight overestimate (of $< 1\%$ to the bootstrap mean in all cases). We believe this is likely due to a small bias in the bisection algorithm. The algorithm tests if the mean of the power samples is less than the target power to decide which half of the search space to continue along. Given that we can bisect up to 10 times in each evaluation, we will likely see a false positive even when we set the significance level for individual comparisons to 1%. This leads to the algorithm favouring a smaller MDE sample. Nevertheless, we are satisfied with the theoretical quantities for experiment design purposes since we have tested for a wide range of parameters and the overall bias is small.

6.5 A Brief Recap

We have addressed the problem of comparing experiment designs for personalisation strategies by presenting an evaluation framework that allows experimenters to evaluate which experiment setup they should adopt. The flexible framework can be easily extended to compare setups that compare more than two strategies by adding more user groups (i.e. new sets to the Venn diagram in Figure 6.1). A new setup can also be quickly incorporated as it is essentially a different weighting of user group-scenario combinations shown in Table 6.1. The framework also allows the development of simple rules of thumb, such as:

1. Treatment effect dilution should never be employed if the experiment already has sufficient power, though it can be helpful if the experiment is under-powered and the non-qualifying users provide a “stabilising effect”; and
2. A dual control setup is superior to simpler setups only if one has access to the user base of the largest organisations.

We have validated the theoretical results via simulations and made the code available² so that practitioners can benefit from the results immediately when designing their upcoming experiments.

Chapter 7

Conclusion

7.1 Summary of Thesis Achievements

We presented and addressed several statistical and data challenges experimenters face when building digital experimentation and measurement (DEM) capabilities from the ground up. They include

1. (Chapter 2) Valuing DEM capabilities to justify related investments via a novel ranking under lower uncertainty model that quantifies the value gained and risk when one lowers the measurement uncertainty of the items they are prioritising;
2. (Chapter 3) Understanding statistical tests and their alternatives to estimate the potential treatment impact via an introduction to statistical testing that aims to balance the theoretical foundations and the practical applications;
3. (Chapter 4) Recognising the availability of datasets to validate existing and develop new methods via the creation of the first ever taxonomy for digital experiment datasets, survey for publicly available online controlled experiment (OCE) datasets, and OCE dataset that can support experiments with adaptive stopping;
4. (Chapter 4) Identifying data collection requirements based on the statistical tests chosen (and vice versa) via a mapping between statistical tests and datasets, two concepts that are often studied in isolation;

5. (Chapter 5) Acquiring knowledge of advanced methods related to the design, running, analysis, and interpretation of digital experiments (while avoiding pitfalls) via a review of recent advances in methods and examples of OCEs, quasi-experiments, and natural experiments;
6. (Chapter 5) Implementing selected methods in the literature via a case study that critiques the effectiveness and practicalness of popular approaches that address dependent responses in OCEs (and geo-experiments to some extent); and
7. (Chapter 6) Designing experiments to compare competing personalisation strategies, known to suffer from low test power and lack of complete control in random assignments, via developing an evaluation framework that is both flexible and capable of generating simple rules of thumb.

Digital experimentation is a multi-disciplinary field with researchers and practitioners from many backgrounds. This is reflected in the diversity of statistical disciplines involved in this thesis. Addressing the challenges above requires the study of order statistics, statistical testing, data collection and processing, experimental design, and causal inference. In addition, the thesis also draws upon techniques in computational statistics and distributed computing to perform the large-scale simulations and empirical verification featured in the thesis. The breadth of topics makes it likely that a researcher or practitioner in digital experimentation can take away something from the text.

The thesis also reached considerable depth within each research topic. Most of the novel research contributions in the thesis have already been presented in and included in proceedings of highly-regarded conferences in applied statistics, machine learning, and data mining. These include the IEEE International Conference on Data Mining (ICDM), AdKDD Workshop (in conjunction with ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD), Conference on Neural Information Processing Systems (NeurIPS), and ACM Web Conference (TheWebConf, formerly WWW). The research has been further recognised via the inclusion of the extended ICDM paper into a journal special issue for “Highly-rated Short Papers for ICDM 2019” and the Best Student Paper award for AdKDD 2020.

Impact beyond academic research The research in this thesis is motivated by the need for and has directly contributed to the development of DEM capabilities at ASOS.com (a global

online fashion retail company) from its infancy:

- Results from the ranking under lower uncertainty model presented in Chapter 2, using parameters reflecting the company's measurement capability at the time, have formed one of the main arguments for the company to invest in initial DEM capabilities.
- Since then, the author has spent years engaging with and educating company colleagues from both technical and non-technical backgrounds on statistical testing, datasets, and experimental design in an online setting. Chapters 3, 4, and 6 reflect the distilled knowledge from that process. Striking a balance between theoretical rigour and applicability in practice also makes the text suitable for other pedagogical purposes.
- The author led a team that incorporates advanced statistical techniques in ASOS.com's internal experiment analysis platform. These include the mSPRT, a sequential test, to address peeking (see Section 3.8) and one-way/block bootstrap to address dependent responses as described in the case study in Section 5.5.4. The author also led another team that has built and is maintaining a geo-experimentation framework in-house that utilises techniques mentioned in Section 5.6.3.

7.2 Future Work

The research has also opened up many opportunities for further work in different statistical and data topics. Below, we outline them in the order they emerge in the thesis.

Ranking under lower uncertainty We believe the general ranking under lower uncertainty problem, introduced in Chapter 2, will interest many in the statistics and operations research community. Many exciting questions remain, not only for valuing DEM capabilities (e.g., understanding how the expected gain and risk change when we assume different item value distributions, see Section 2.8.1), but also for general prioritisation processes (e.g., the value gained when not all items in the prioritisation process can have their estimation noise reduced, see Section 2.8.2). It also yields an opportunity for us to gain a further understanding of order statistics via questions like the probability that two order statistics in the ranking under lower uncertainty model are generated from the same r.v. (see Appendix A.1).

Datasets for digital experiments The work in Chapter 4 has highlighted the need for more publicly available digital experiment datasets, both in quantity and variety, a call echoed by [159]. This is necessary to foster academic-industry collaboration and safeguard future methodological development. As more datasets become available, the survey, taxonomy, and mapping introduced in the chapter will require updating and extending.

Dependent responses in e-commerce experiments In Section 5.5.4, we observed that the two-way bootstrap, which addresses responses dependent on multiple types of units, produces impractically large standard error estimates and unnecessarily conservative confidence intervals under online retail/e-commerce transactions datasets. Further work is required to understand how and why this is the case. This enables experimenters to adjust the models and methods to estimate the degree of dependence between responses accurately.

Evaluation framework for personalisation strategy experiment designs The evaluation framework introduced in Chapter 6 assumes the responses from each user group-scenario combination are randomly distributed with the mean independent of the variance, with the evaluation criteria calculated assuming the decision metric is approximately normally distributed under the central limit theorem. While they are reasonable assumptions in practice, it begs the question of whether the evaluation framework and its results are robust to deviations from these assumptions, e.g., with binary responses (where the mean and variance correlate) and heavy-tailed distributed responses (where the sample mean converges to a normal distribution slowly).

In addition to the extensions outlined above, we propose the following related, high-potential research direction involving sub-disciplines not covered in the thesis. These sub-disciplines include (1) automatic experimentation and optimisation capabilities via bandit and Bayesian optimisation algorithms [117, 210] and (2) combining experimental and observational data to provide more accurate treatment effect estimates [141].

Automatic experiment prioritisation with decision makers in-the-loop It is increasingly common to deploy automatic experimentation and optimisation capabilities within individ-

ual business functions or digital products. However, the problem of finding the optimal sequencing of experimental interventions at a cross-functional level or across multiple digital products is often dealt with manually. We argue that a Bayesian optimisation framework, which balances exploitation (obtaining the highest valued intervention) and exploration (eliminating uncertainty in value estimates), is well suited to recommending the next interventions to be evaluated. The framework should be central to an iterative decision framework where we

1. Measure interventions' value via experiments,
2. Estimate other candidate interventions' value using causal inference techniques, and
3. Identify the next best interventions to be evaluated.

This can be roughly linked to the points on the objective function, the surrogate model, and the acquisition function. With the objective function being largely unknown, we propose researching into constructing novel surrogate models and acquisition functions that can address other challenges presented by prioritising experiments in real life. These include quantifying the uncertainty of interventions in a mixed experimental-observational setting with potential unmeasured confounders, having concept drift erode the confidence in past effect estimates over time, and incorporating preference from decision makers in addition to modelled effect estimates.

To sum up, the need to understand the cause and effect of decisions and test out different alternatives in the current digital era has led to a rich body of work in digital experimentation and measurement. The field has attracted substantial interdisciplinary interest in the past two decades and will remain fast-growing, as demonstrated by the open challenges above. This thesis has contributed to, and seeks to continue inspiring such growth.

References

- [1] Christopher P. Adams. “Empirical Bayesian Estimation of Treatment Effects”. In: *SSRN Electronic Journal* (2018). DOI: 10.2139/ssrn.3157819 (cited on p. 126).
- [2] Fares Alahdab. “Potential impact of large language models on academic writing”. In: *BMJ Evidence-Based Medicine* (2023). ISSN: 2515-446X. DOI: 10.1136/bmjebm-2023-112429 (cited on p. 1).
- [3] John Aldrich. “Correlations Genuine and Spurious in Pearson and Yule”. In: *Statistical Science* 10.4 (1995), pp. 364–376. DOI: 10.1214/ss/1177009870 (cited on p. 69).
- [4] Sadiq Alreemi. *Analyze AB Test Results* [Code repository]. 2019. URL: <https://github.com/SadiqAlreemi/Analyze-AB-Test-Results> (accessed 9/9/2021) (cited on p. 122).
- [5] Andrew D. Althouse. “Post Hoc Power: Not Empowering, Just Misleading”. In: *Journal of Surgical Research* 259 (2021), A3–A6. ISSN: 0022-4804. DOI: 10.1016/j.jss.2019.10.049. URL: [https://www.journalofsurgicalresearch.com/article/S0022-4804\(20\)30502-3/fulltext](https://www.journalofsurgicalresearch.com/article/S0022-4804(20)30502-3/fulltext) (cited on p. 96).
- [6] Joshua D. Angrist and Jörn-Steffen Pischke. *Mastering ‘Metrics: The Path from Cause to Effect*. Princeton University Press, 2014. ISBN: 9780691152844. URL: <https://www.masteringmetrics.com/> (cited on p. 167).
- [7] AQA. *A-Level Mathematics (7357): Specification*. Version 1.3. 2018. URL: <https://filestore.aqa.org.uk/resources/mathematics/specifications/AQA-7357-SP-2017.PDF> (cited on p. 68).
- [8] Anthony C. Atkinson. “Optimum Experimental Design”. In: *International Encyclopedia of Statistical Science*. Ed. by Miodrag Lovric. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 1037–1039. ISBN: 978-3-642-04898-2. DOI: 10.1007/978-3-642-04898-2_434 (cited on p. 21).
- [9] Timothy C. Au. “A Time-Based Regression Matched Markets Approach for Designing Geo Experiments”. 2018. URL: <https://research.google/pubs/pub48983.pdf> (cited on p. 171).
- [10] Florian Auer and Michael Felderer. “Current State of Research on Continuous Experimentation: A Systematic Mapping Study”. In: *2018 44th Euromicro Conference on Software Engineering and Advanced Applications (SEAA)*. Prague, Czech Republic: IEEE, 2018, pp. 335–344. ISBN: 978-1-5386-7383-6. DOI: 10.1109/SEAA.2018.00062 (cited on p. 118).
- [11] Florian Auer and Michael Felderer. *Characteristics of an Online Controlled Experiment: Preliminary Results of a Literature Review*. 2019. arXiv: 1912.01383 [cs.SE] (cited on p. 152).
- [12] Florian Auer, Rasmus Ros, Lukas Kaltenbrunner, Per Runeson, and Michael Felderer. “Controlled experimentation in continuous experimentation: Knowledge and challenges”. In: *Information and Software Technology* 134 (2021), p. 106551. ISSN: 0950-5849. DOI: 10.1016/j.infsof.2021.106551 (cited on pp. 118, 145).

- [13] Yamaç Eren Ay. *Synthetical A/B-Tests, Version 1 [Dataset]*. 2020. URL: <https://www.kaggle.com/yamaerenay/synthetical-abtests/version/1> (accessed 13/10/2021) (cited on pp. 121, 122).
- [14] Eduardo M. Azevedo, Alex Deng, José L. Montiel Olea, and E. Glen Weyl. “Empirical Bayes Estimation of Treatment Effects with Many A/B Tests: An Overview”. In: *AEA Papers and Proceedings* 109 (2019), pp. 43–47. doi: 10.1257/pandp.20191003 (cited on pp. 126, 147).
- [15] Rasmus Bååth and Bret Romero. *Mobile Games A/B Testing with Cookie Cats [MOOC]*. URL: <https://www.datacamp.com/projects/184> (accessed 25/8/2021) (cited on pp. 121, 122).
- [16] R. H. Baayen, D. J. Davidson, and D. M. Bates. “Mixed-effects modeling with crossed random effects for subjects and items”. In: *Journal of Memory and Language* 59.4 (2008). Special Issue: Emerging Data Analysis, pp. 390–412. ISSN: 0749-596X. doi: <https://doi.org/10.1016/j.jml.2007.12.005> (cited on p. 155).
- [17] Lilia Ben Baccar and Clara Dionet. *Geo Experimentation at ASOS: Measuring Advertising Effectiveness for Data Driven Marketing*. Presentation in Data Science Festival Mayday 2023, London, UK. 2023. URL: <https://www.youtube.com/watch?v=ZoRVrurPnnw> (cited on pp. 169, 171).
- [18] Lars Backstrom and Jon Kleinberg. “Network Bucket Testing”. In: *Proceedings of the 20th International Conference on World Wide Web. WWW ’11*. Hyderabad, India: ACM, 2011, pp. 615–624. ISBN: 9781450306324. doi: 10.1145/1963405.1963492 (cited on pp. 37, 150).
- [19] Eytan Bakshy and Dean Eckles. “Uncertainty in Online Experiments with Dependent Data: An Evaluation of Bootstrap Methods”. In: *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. KDD ’13*. Chicago, Illinois, USA: ACM, 2013, pp. 1303–1311. ISBN: 9781450321747. doi: 10.1145/2487575.2488218 (cited on pp. 37, 124, 149, 150, 155, 157, 158, 161–163).
- [20] Joel Barajas, Tom Zidar, and Mert Bay. *Advertising Incrementality Measurement using Controlled Geo-Experiments: The Universal App Campaign Case Study*. Presented in AdKDD 2020 Workshop (in conjunction with KDD ’20). 2020. URL: <http://papers.adkdd.org/2020/papers/adkdd20-barajas-advertising.pdf> (cited on pp. 34, 169, 171).
- [21] M. J. Bayarri and James O. Berger. “Quantifying Surprise in the Data and Model Verification”. In: *Bayesian Statistics 6: Proceedings of the Sixth Valencia International Meeting*. Ed. by José M. Bernardo, James O. Berger, A. P. Dawid, and Adrian F. M. Smith. Oxford University Press, 1998, pp. 53–82 (cited on p. 90).
- [22] David E. Bell. “Regret in Decision Making Under Uncertainty”. In: *Operations Research* 30.5 (1982), pp. 961–981. ISSN: 0030-364X. doi: 10.1287/opre.30.5.961 (cited on p. 37).
- [23] Gerald van Belle. *Statistical Rules of Thumb*. Wiley Series in Probability and Statistics. John Wiley & Sons, 2008. ISBN: 9780470377963. doi: 10.1002/9780470377963 (cited on p. 99).
- [24] Brooke Bengier and Amanda Knupp. *Selling More Stuff: The What, Why & How of Incrementality Testing*. Blog post. 2017. URL: <https://www.quantcast.com/blog/selling-more-stuff-the-what-why-how-of-incrementality-testing/> (cited on p. 174).
- [25] Yoav Benjamini and Yosef Hochberg. “Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing”. In: *Journal of the Royal Statistical Society: Series B (Methodological)* 57.1 (1995), pp. 289–300. doi: 10.1111/j.2517-6161.1995.tb02031.x (cited on pp. 119, 146).
- [26] Joseph Berkson. “Some Difficulties of Interpretation Encountered in the Application of the Chi-Square Test”. In: *Journal of the American Statistical Association* 33.203 (1938), pp. 526–536. ISSN: 0162-1459. doi: 10.1080/01621459.1938.10502329 (cited on p. 68).

- [27] Gunnar Blom. “Statistical estimates and transformed beta-variables”. PhD thesis. Stockholm College, 1958 (cited on pp. 42, 43).
- [28] Robert C. Bolles. “The Difference between Statistical Hypotheses and Scientific Hypotheses”. In: *Psychological Reports* 11.3 (1962), pp. 639–645. doi: 10.2466/pr0.1962.11.3.639 (cited on p. 70).
- [29] Brooks Bell. *The State of Experimentation Report: 2019 - 2020 Edition*. White paper. 2020. URL: https://www.brooksbell.com/wp-content/uploads/2020/07/State-of-Experimentation-Report_Final.pdf (cited on p. 152).
- [30] Will Browne and Mike Swarbrick Jones. *What Works in e-commerce - a Meta-analysis of 6700 Online Experiments*. White paper. 2017. URL: <https://web.archive.org/web/20210331032057/http://www.qubit.com/wp-content/uploads/2017/12/qubit-research-meta-analysis.pdf> (accessed 15/5/2019) (cited on pp. 34, 36, 40, 45, 57, 59, 60, 152).
- [31] Andreas Buja, Dianne Cook, Heike Hofmann, Michael Lawrence, Eun-Kyung Lee, Deborah F. Swayne, and Hadley Wickham. “Statistical inference for exploratory data analysis and model diagnostics”. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 367.1906 (2009), pp. 4361–4383. ISSN: 1364-503X. doi: 10.1098/rsta.2009.0120 (cited on p. 73).
- [32] A. Colin Cameron, Jonah B. Gelbach, and Douglas L. Miller. “Robust Inference With Multiway Clustering”. In: *Journal of Business & Economic Statistics* 29.2 (2011), pp. 238–249. doi: 10.1198/jbes.2010.07136 (cited on p. 155).
- [33] David Campos, Chester Ismay, and Shon Inouye. *A/B Testing in R — DataCamp [MOOC]*. URL: <https://www.datacamp.com/courses/ab-testing-in-r> (accessed 25/8/2021) (cited on pp. 121, 122).
- [34] David Campos, Chester Ismay, and Shon Inouye. *Data Visualization Website - April 2018 [Dataset]*. URL: https://assets.datacamp.com/production/repositories/2292/datasets/b502094e5de478105cccea959d4f91data_viz_website_2018_04.csv (accessed 25/8/2021) (cited on p. 122).
- [35] David Campos, Chester Ismay, and Shon Inouye. *Experiment dataset [Dataset]*. URL: https://assets.datacamp.com/production/repositories/2292/datasets/52b52cb1ca28ce10f9a09689325c4d94d889a6da/experiment_data.csv (accessed 25/8/2021) (cited on p. 122).
- [36] George Casella and Roger L. Berger. *Statistical Inference*. 2nd. Duxbury advanced series in statistics and decision sciences. Duxbury, 2002. ISBN: 9780534243128 (cited on pp. 69, 78, 90).
- [37] Charlie Casper, Thomas Cook, and Oscar A. Perez. *ldbounds: Lan-DeMets Method for Group Sequential Boundaries*. R package version 2.0.1. Based on FORTRAN program ld98. 2023. URL: <https://CRAN.R-project.org/package=ldbounds> (cited on p. 109).
- [38] Horng-Jinh Chang, Kuo-Chung Huang, and Chao-Hsien Wu. “Determination of Sample Size in Using Central Limit Theorem for Weibull Distribution”. In: *International Journal of Information and Management Sciences* 17.3 (2006), pp. 31–46. URL: <https://tpl.ncl.edu.tw/NclService/JournalContentDetail?SysId=A06083289> (cited on p. 85).
- [39] Aiyu Chen and Timothy C. Au. *Robust Causal Inference for Incremental Return on Ad Spend with Randomized Paired Geo Experiments*. 2021. arXiv: 1908.02922 [stat.ME] (cited on p. 171).
- [40] Daqing Chen, Sai Laing Sain, and Kun Guo. “Data mining for the online retail industry: A case study of RFM model-based customer segmentation using data mining”. In: *Journal of Database Marketing & Customer Strategy Management* 19.3 (2012), pp. 197–208. ISSN: 1741-2447. doi: 10.1057/dbm.2012.17 (cited on p. 156).

- [41] Nanyu Chen, Min Liu, and Ya Xu. "How A/B Tests Could Go Wrong: Automatic Diagnosis of Invalid Online Experiments". In: *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*. WSDM '19. Melbourne VIC, Australia: ACM, 2019, pp. 501–509. ISBN: 9781450359405. DOI: 10.1145/3289600.3291000 (cited on pp. 130, 152, 153).
- [42] Sisi Chen. *A/B tests project — Udacity DAND P2*. Blog post. 2020. URL: <https://rachelchen0104.medium.com/a-b-tests-project-udacity-dand-p2-55fadd4d5e0> (accessed 9/9/2021) (cited on p. 122).
- [43] Yuhui Chen and Timothy E. Hanson. "Bayesian nonparametric k-sample tests for censored and uncensored data". In: *Computational Statistics & Data Analysis* 71 (2014), pp. 335–346. ISSN: 0167-9473. DOI: <https://doi.org/10.1016/j.csda.2012.11.003> (cited on p. 127).
- [44] Yuan Shih Chow and Henry Teicher. "Binomial Random Variables". In: *Probability Theory: Independence, Interchangeability, Martingales*. New York, NY: Springer New York, 1997, pp. 30–53. ISBN: 978-1-4612-1950-7. DOI: 10.1007/978-1-4612-1950-7_2 (cited on p. 141).
- [45] Dominic Coey and Tom Cunningham. "Improving Treatment Effect Estimators Through Experiment Splitting". In: *The World Wide Web Conference*. WWW '19. San Francisco, CA, USA: ACM, 2019, pp. 285–295. ISBN: 9781450366748. DOI: 10.1145/3308558.3313452 (cited on p. 147).
- [46] Jacob Cohen. *Statistical Power Analysis for the Behavioral Sciences*. Taylor & Francis, 1988. ISBN: 9781134742707. DOI: 10.4324/9780203771587 (cited on p. 74).
- [47] Jacob Cohen. "The earth is round ($p < .05$)". In: *American Psychologist* 49.12 (1994), pp. 997–1003. ISSN: 1935-990X. DOI: 10.1037/0003-066X.49.12.997 (cited on pp. 68, 70, 71, 109).
- [48] College Board. *AP Statistics: Course and Exam Description, Effective Fall 2020*. 2020. URL: <https://apcentral.collegeboard.org/media/pdf/ap-statistics-course-and-exam-description.pdf> (cited on p. 68).
- [49] Committee for Proprietary Medicinal Products. "Points to consider on switching between superiority and non-inferiority". In: *British Journal of Clinical Pharmacology* 52.3 (2001), pp. 223–228. ISSN: 0306-5251. DOI: 10.1046/j.1365-2125.2001.01397-3.x (cited on p. 70).
- [50] Guido Consonni, Dimitris Fouskakis, Brunero Liseo, and Ioannis Ntzoufras. "Prior Distributions for Objective Bayesian Analysis". In: *Bayesian Analysis* 13.2 (2018), pp. 627–679. DOI: 10.1214/18-BA1103 (cited on p. 112).
- [51] David R. Cox. *Planning of Experiments*. Wiley Series in Probability and Statistics - Applied Probability and Statistics Section. Wiley, 1958. ISBN: 9780471181835 (cited on p. 149).
- [52] David R. Cox and David V. Hinkley. *Theoretical Statistics*. London: Chapman and Hall, 1974. ISBN: 978-0412124204 (cited on p. 75).
- [53] Thomas Crook, Brian Frasca, Ron Kohavi, and Roger Longbotham. "Seven Pitfalls to Avoid when Running Controlled Experiments on the Web". In: *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. KDD '09. Paris, France: ACM, 2009, pp. 1105–1114. ISBN: 978-1-60558-495-9. DOI: 10.1145/1557019.1557139 (cited on pp. 152, 153).
- [54] F. N. David and N. L. Johnson. "Statistical Treatment of Censored Data: Part I. Fundamental Formulae". In: *Biometrika* 41.1-2 (1954), pp. 228–240. ISSN: 0006-3444. DOI: 10.1093/biomet/41.1-2.228 (cited on pp. 48, 56).
- [55] H. A. David and H. N. Nagaraja. "Concomitants of order statistics". In: *Order Statistics: Theory & Methods*. Ed. by N. Balakrishnan and C. R. Rao. Vol. 16. Handbook of Statistics. Elsevier, 1998, pp. 487–513. ISBN: 978-0-444-82091-4. DOI: 10.1016/S0169-7161(98)16020-0 (cited on pp. 48, 50).

- [56] H. A. David and H. N. Nagaraja. *Order Statistics*. Wiley Series in Probability and Statistics. Hoboken, NJ, USA: John Wiley & Sons, Inc., 2003. ISBN: 9780471722168. DOI: 10.1002/0471722162 (cited on pp. 42, 47, 49, 51).
- [57] Ahmed Dawoud. *E-commerce A/B testing, Version 1 [Dataset]*. 2021. URL: <https://www.kaggle.com/ahmedmohameddawoud/ecommerce-ab-testing/version/1> (accessed 9/9/2021) (cited on p. 122).
- [58] Anirban Deb, Suman Bhattacharya, Jeremy Gu, Tianxia Zhou, Eva Feng, and Mandie Liu. *Under the Hood of Uber's Experimentation Platform*. 2018. URL: <https://www.uber.com/en-GB/blog/xp/> (cited on p. 73).
- [59] Rajeev H. Dehejia and Sadek Wahba. "Causal Effects in Nonexperimental Studies: Reevaluating the Evaluation of Training Programs". In: *Journal of the American Statistical Association* 94.448 (1999), pp. 1053–1062. ISSN: 0162-1459. DOI: 10.1080/01621459.1999.10473858 (cited on p. 121).
- [60] David L. DeMets and K. K. Gordon Lan. "Interim analysis: The alpha spending function approach". In: *Statistics in Medicine* 13.13-14 (1994), pp. 1341–1352. DOI: 10.1002/sim.4780131308 (cited on p. 109).
- [61] Alex Deng. "Objective Bayesian Two Sample Hypothesis Testing for Online Controlled Experiments". In: *Proceedings of the 24th International Conference on World Wide Web. WWW '15 Companion*. Florence, Italy: ACM, 2015, pp. 923–928. ISBN: 9781450334730. DOI: 10.1145/2740908.2742563 (cited on pp. 110, 112).
- [62] Alex Deng. *Causal Inference and Its Applications in Online Industry*. Work in progress. 2021. URL: <https://alex deng.github.io/causal/> (cited on pp. 137, 164, 166).
- [63] Alex Deng and Victor Hu. "Diluted Treatment Effect Estimation for Trigger Analysis in Online Controlled Experiments". In: *Proceedings of the Eighth ACM International Conference on Web Search and Data Mining. WSDM '15*. Shanghai, China: ACM, 2015, pp. 349–358. ISBN: 9781450333177. DOI: 10.1145/2684822.2685307 (cited on pp. 151, 174, 178, 184).
- [64] Alex Deng, Ulf Knoblich, and Jiannan Lu. "Applying the Delta Method in Metric Analytics: A Practical Guide with Novel Ideas". In: *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. KDD '18*. London, United Kingdom: ACM, 2018, pp. 233–242. ISBN: 9781450355520. DOI: 10.1145/3219819.3219919 (cited on pp. 124, 149, 155, 161, 162).
- [65] Alex Deng, Yicheng Li, Jiannan Lu, and Vivek Ramamurthy. "On Post-selection Inference in A/B Testing". In: *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining. KDD '21*. Virtual Event, Singapore: ACM, 2021, pp. 2743–2752. ISBN: 9781450383325. DOI: 10.1145/3447548.3467129 (cited on p. 147).
- [66] Alex Deng, Jiannan Lu, and Shouyuan Chen. "Continuous Monitoring of A/B Tests without Pain: Optional Stopping in Bayesian Testing". In: *2016 IEEE International Conference on Data Science and Advanced Analytics (DSAA)*. 2016, pp. 243–252. DOI: 10.1109/DSAA.2016.33 (cited on pp. 110, 112, 124, 129).
- [67] Alex Deng, Jiannan Lu, and Jonthan Litz. "Trustworthy Analysis of Online A/B Tests: Pitfalls, challenges and solutions". In: *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining. WSDM '17*. Cambridge, United Kingdom: ACM, 2017, pp. 641–649. ISBN: 978-1-4503-4675-7. DOI: 10.1145/3018661.3018677 (cited on pp. 148, 149, 156, 161).
- [68] Alex Deng and Xiaolin Shi. "Data-Driven Metric Development for Online Controlled Experiments: Seven Lessons Learned". In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. KDD '16*. San Francisco, California, USA: ACM, 2016, pp. 77–86. ISBN: 978-1-4503-4232-2. DOI: 10.1145/2939672.2939700 (cited on p. 152).

- [69] Alex Deng, Ya Xu, Ron Kohavi, and Toby Walker. “Improving the Sensitivity of Online Controlled Experiments by Utilizing Pre-Experiment Data”. In: *Proceedings of the Sixth ACM International Conference on Web Search and Data Mining*. WSDM ’13. Rome, Italy: ACM, 2013, pp. 123–132. ISBN: 9781450318693. doi: 10.1145/2433396.2433413 (cited on pp. 151, 174).
- [70] Shaojie Deng, Roger Longbotham, Toby Walker, and Ya Xu. “Choice of the Randomization Unit in Online Controlled Experiment”. In: *JSM Proceedings*. Miami Beach, Florida: American Statistical Association, 2011. ISBN: 9780983937500. URL: <https://alex deng.github.io/public/files/jsm2011-deng.pdf> (cited on p. 148).
- [71] James K. Dias. *Selecting Statistical Procedures*. Handout from the “Biostatistics for Librarians” presentation in the 2010 Annual Meeting of the Georgia Health Sciences Library Association – Continuing Education Programs – Augusta, GA. 2010. URL: <https://ghsla.org/wp-content/uploads/2013/12/Dias2010b.pdf> (cited on p. 73).
- [72] Eustache Diemert, Artem Betlei, Christophe Renaudin, and Amini Massih-Reza. *A Large Scale Benchmark for Uplift Modeling*. Presented in 2018 AdKDD & TargetAd Workshop (in conjunction with KDD ’18). 2018. URL: <http://ama.imag.fr/~amini/Publis/large-scale-benchmark.pdf> (cited on p. 121).
- [73] Directorate for Financial and Enterprise Affairs Competition Committee, Organisation for Economic Co-operation and Development. *Personalised Pricing in the Digital Era – Note by the United Kingdom*. 2018. URL: [https://one.oecd.org/document/DAF/COMP/WD\(2018\)127/en/pdf](https://one.oecd.org/document/DAF/COMP/WD(2018)127/en/pdf) (cited on p. 143).
- [74] Pavel Dmitriev, Somit Gupta, Dong Woo Kim, and Garnet Vaz. “A Dirty Dozen: Twelve Common Metric Interpretation Pitfalls in Online Controlled Experiments”. In: *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. KDD ’17. Halifax, NS, Canada: ACM, 2017, pp. 1427–1436. ISBN: 978-1-4503-4887-4. doi: 10.1145/3097983.3098024 (cited on pp. 37, 104, 152, 153, 178).
- [75] Yadolah Dodge. “Kolmogorov–Smirnov Test”. In: *The Concise Encyclopedia of Statistics*. New York, NY: Springer New York, 2008, pp. 283–287. ISBN: 978-0-387-32833-1. doi: 10.1007/978-0-387-32833-1_214 (cited on p. 101).
- [76] Nianbo Dong and Rebecca Maynard. “PowerUp!: A Tool for Calculating Minimum Detectable Effect Sizes and Minimum Required Sample Sizes for Experimental and Quasi-Experimental Design Studies”. In: *Journal of Research on Educational Effectiveness* 6.1 (2013), pp. 24–67. ISSN: 1934-5747. doi: 10.1080/19345747.2012.673143 (cited on p. 69).
- [77] Joseph Doyle, Sarah Abraham, Laura Feeney, Sarah Reimer, and Amy Finkelstein. *Clinical decision support for high-cost imaging: a randomized clinical trial [Dataset]*. Version V3. 2018. doi: 10.7910/DVN/BRKDVQ (cited on p. 121).
- [78] Dheeru Dua and Casey Graff. *UCI Machine Learning Repository*. 2019. URL: <http://archive.ics.uci.edu/ml> (cited on pp. 116, 118, 156).
- [79] Wouter Duivesteijn, Tara Farzami, Thijs Putman, Evertjan Peer, Hilde J. P. Weerts, Jasper N. Adegeest, Gerson Foks, and Mykola Pechenizkiy. “Have It Both Ways—From A/B Testing to A&B Testing with Exceptional Model Mining”. In: *Machine Learning and Knowledge Discovery in Databases*. Ed. by Y. Altun, K. Das, T. Mielikäinen, D. Malerba, J. Stefanowski, J. Read, M. Žitnik, M. Ceci, and S. Džeroski. Vol. 10536. Lecture Notes in Computer Science. Cham: Springer International Publishing, 2017, pp. 114–126. ISBN: 978-3-319-71272-7. doi: 10.1007/978-3-319-71273-4_10 (cited on p. 147).

- [80] Olive Jean Dunn. "Multiple Comparisons among Means". In: *Journal of the American Statistical Association* 56.293 (1961), pp. 52–64. ISSN: 0162-1459. DOI: 10.1080/01621459.1961.10482090 (cited on p. 146).
- [81] Pascaline Dupas, Vivian Hoffman, Michael Kremer, and Alix Peterson Zwane. *Targeting health subsidies through a non-price mechanism: A randomized controlled trial in Kenya [Dataset]*. Version V1. 2016. DOI: 10.7910/DVN/PBLJXJ (cited on p. 121).
- [82] Arpit Dwivedi. *Cookie Cats, Version 1 [Dataset]*. 2020. URL: <https://www.kaggle.com/arpitdw/cookie-cats/version/1> (accessed 25/8/2021) (cited on p. 122).
- [83] Dean Eckles, Brian Karrer, and Johan Ugander. "Design and Analysis of Experiments in Networks: Reducing Bias from Interference". In: *Journal of Causal Inference* 5.1 (2017), p. 20150021. DOI: 10.1515/jci-2015-0021 (cited on p. 155).
- [84] Ward Edwards, Harold Lindman, and Leonard J. Savage. "Bayesian statistical inference for psychological research". In: *Psychological Review* 70.3 (1963), pp. 193–242. DOI: 10.1037/h0044139 (cited on p. 73).
- [85] Bradley Efron. "Tweedie's Formula and Selection Bias". In: *Journal of the American Statistical Association* 106.496 (2011), pp. 1602–1614. ISSN: 0162-1459. DOI: 10.1198/jasa.2011.tm11181 (cited on p. 147).
- [86] Bradley Efron and Robert J. Tibshirani. *An Introduction to the Bootstrap*. Monographs on Statistics and Applied Probability 57. Boca Raton, Florida, USA: Chapman & Hall/CRC, 1994. DOI: 10.1201/9780429246593 (cited on pp. 84, 163).
- [87] Paul D. Ellis. *The Essential Guide to Effect Sizes: Statistical Power, Meta-Analysis, and the Interpretation of Research Results*. Cambridge: Cambridge University Press, 2010. ISBN: 9780511761676. DOI: 10.1017/CB09780511761676 (cited on pp. 69, 95).
- [88] Osuolale Emmanuel. *Ad A/B Testing, Version 1 [Dataset]*. 2020. URL: <https://www.kaggle.com/osuolaleemmanuel/ad-ab-testing/version/1> (accessed 25/8/2021) (cited on pp. 121, 122).
- [89] Maria Esteller-Cucala, Vicenc Fernandez, and Diego Villuendas. "Experimentation Pitfalls to Avoid in A/B Testing for Online Personalization". In: *Adjunct Publication of the 27th Conference on User Modeling, Adaptation and Personalization*. UMAP'19 Adjunct. Larnaca, Cyprus: ACM, 2019, pp. 153–159. ISBN: 9781450367110. DOI: 10.1145/3314183.3323853 (cited on pp. 152, 153).
- [90] Aleksander Fabijan, Pavel Dmitriev, Helena Holmstrom Olsson, and Jan Bosch. "Online Controlled Experimentation at Scale: An Empirical Survey on the Current State of A/B Testing". In: *2018 44th Euromicro Conference on Software Engineering and Advanced Applications (SEAA)*. Prague, Czech Republic: IEEE, 2018, pp. 68–72. ISBN: 978-1-5386-7383-6. DOI: 10.1109/SEAA.2018.00021 (cited on pp. 20, 118).
- [91] Aleksander Fabijan, Pavel Dmitriev, Helena Holmstrom Olsson, and Jan Bosch. "The Online Controlled Experiment Lifecycle". In: *IEEE Software* 37.2 (2020), pp. 60–67. ISSN: 0740-7459. DOI: 10.1109/MS.2018.2875842 (cited on p. 152).
- [92] Aleksander Fabijan, Pavel Dmitriev, Helena Holmstrom Olsson, and Jan Bosch. "The Evolution of Continuous Experimentation in Software Product Development: From Data to a Data-Driven Organization at Scale". In: *2017 IEEE/ACM 39th International Conference on Software Engineering (ICSE)*. Buenos Aires, Argentina: IEEE, May 2017, pp. 770–780. ISBN: 978-1-5386-3868-2. DOI: 10.1109/ICSE.2017.76 (cited on p. 152).

- [93] Aleksander Fabijan, Jayant Gupchup, Somit Gupta, Jeff Omhover, Wen Qin, Lukas Vermeer, and Pavel Dmitriev. “Diagnosing Sample Ratio Mismatch in Online Controlled Experiments: A Taxonomy and Rules of Thumb for Practitioners”. In: *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. KDD ’19. Anchorage, AK, USA: ACM, 2019, pp. 2156–2164. ISBN: 9781450362016. DOI: 10.1145/3292500.3330722 (cited on pp. 105, 152, 153).
- [94] Franz Faul, Edgar Erdfelder, Albert-Georg Lang, and Axel Buchner. “G*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences”. In: *Behavior Research Methods* 39.2 (2007), pp. 175–191. ISSN: 1554-3528. DOI: 10.3758/BF03193146 (cited on p. 96).
- [95] Ronald A. Fisher. “On the Mathematical Foundations of Theoretical Statistics”. In: *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character* 222.594–604 (1922), pp. 309–368. DOI: 10.1098/rsta.1922.0009 (cited on p. 123).
- [96] Joseph L. Fleiss, Bruce Levin, and Myunghee Cho Paik. “Determining Sample Sizes Needed to Detect a Difference between Two Proportions”. In: *Statistical Methods for Rates and Proportions*. Hoboken, NJ, USA: John Wiley & Sons, Inc., 2003. Chap. 4, pp. 64–85. ISBN: 9780471445425. DOI: 10.1002/0471445428.ch4 (cited on p. 69).
- [97] Rina Foygel Barber and Emmanuel J. Candès. “Controlling the false discovery rate via knockoffs”. In: *The Annals of Statistics* 43.5 (2015), pp. 2055–2085. DOI: 10.1214/15-AOS1337 (cited on p. 146).
- [98] Michele Jonsson Funk, Daniel Westreich, Chris Wiesen, Til Stürmer, M. Alan Brookhart, and Marie Davidian. “Doubly Robust Estimation of Causal Effects”. In: *American Journal of Epidemiology* 173.7 (2011), pp. 761–767. ISSN: 0002-9262. DOI: 10.1093/aje/kwq439 (cited on p. 168).
- [99] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, and Kate Crawford. *Datasheets for Datasets*. 2020. arXiv: 1803.09010 [cs.DB] (cited on p. 128).
- [100] Andrew Gelman and Jennifer Hill. *Data Analysis Using Regression and Multilevel/Hierarchical Models*. Analytical Methods for Social Research. Cambridge University Press, 2006. DOI: 10.1017/CB09780511790942 (cited on p. 167).
- [101] Georgi Z. Georgiev. *Statistical Methods in Online A/B Testing*. Independently Published, 2019. ISBN: 978-1694079725. URL: <https://www.abtestingstats.com/> (cited on p. 69).
- [102] Zoubin Ghahramani. “A Brief Overview of Nonparametric Bayesian Models”. NIPS 2009 Workshop on Nonparametric Bayesian Methods. 2009. URL: <https://mlg.eng.cam.ac.uk/zoubin/talks/nips09npb.pdf> (cited on p. 72).
- [103] Sourobh Ghosh, Stefan H Thomke, and Hazjier Pourkhalkhali. “The Effects of Hierarchy on Learning and Performance in Business Experimentation”. Harvard Business School Working Paper 20-081. 2020. URL: https://www.hbs.edu/faculty/Publication%20Files/20-081_9d2608f9-5a49-4bc3-a902-307d3427477b.pdf (cited on p. 152).
- [104] Phillip I. Good and James W. Hardin. *Common Errors in Statistics (And How to Avoid Them)*. John Wiley & Sons, Ltd, 2012. ISBN: 9781118360125. DOI: 10.1002/9781118360125 (cited on p. 70).
- [105] Brett R. Gordon, Florian Zettelmeyer, Neha Bhargava, and Dan Chapsky. “A Comparison of Approaches to Advertising Measurement: Evidence from Big Field Experiments at Facebook”. In: *Marketing Science* 38.2 (2019), pp. 193–225. DOI: 10.1287/mksc.2018.1135 (cited on pp. 34, 60).
- [106] William Sealy (Student) Gosset. “The Probable Error of a Mean”. In: *Biometrika* 6.1 (1908), pp. 1–25. ISSN: 0006-3444. DOI: 10.1093/biomet/6.1.1 (cited on p. 82).

- [107] Judith Greene and Manuela D'Oliveira. *Learning To Use Statistical Tests In Psychology*. 3rd ed. Open guides to psychology. Open University Press, 2006. ISBN: 9780335216802 (cited on p. 73).
- [108] Sander Greenland, Judea Pearl, and James M. Robins. "Confounding and Collapsibility in Causal Inference". In: *Statistical Science* 14.1 (1999), pp. 29–46. DOI: 10.1214/ss/1009211805 (cited on p. 166).
- [109] Sander Greenland, Stephen J. Senn, Kenneth J. Rothman, John B. Carlin, Charles Poole, Steven N. Goodman, and Douglas G. Altman. "Statistical tests, P values, confidence intervals, and power: a guide to misinterpretations". In: *European Journal of Epidemiology* 31 (2016), pp. 337–350. DOI: 10.1007/s10654-016-0149-3 (cited on pp. 76, 117).
- [110] Noah Greifer and Elizabeth A. Stuart. *Choosing the Causal Estimand for Propensity Score Analysis of Observational Studies*. 2023. arXiv: 2106.10577 [stat.ME] (cited on p. 165).
- [111] Carrie Grimes, Caroline Buckey, and Diane Tang. *A/B Testing — Udacity Free Courses [MOOC]*. URL: <https://www.udacity.com/course/ab-testing--ud257> (accessed 25/8/2021) (cited on pp. 121, 122).
- [112] Ryan Grossman. *AB Testing Results [Dataset]*. URL: https://assets.datacamp.com/production/repositories/1646/datasets/2751adce60684a03d8b4132adeadab8a0b95ee56/AB_testing_exercise.csv (accessed 13/10/2021) (cited on p. 122).
- [113] Ryan Grossman. *Customer Analytics and A/B Testing in Python — DataCamp [MOOC]*. URL: <https://www.datacamp.com/courses/customer-analytics-and-ab-testing-in-python> (accessed 13/10/2021) (cited on pp. 121, 122).
- [114] Huan Gui, Ya Xu, Anmol Bhasin, and Jiawei Han. "Network A/B Testing: From Sampling to Estimation". In: *Proceedings of the 24th International Conference on World Wide Web. WWW '15*. Florence, Italy: International World Wide Web Conferences Steering Committee, 2015, pp. 399–409. ISBN: 9781450334693. DOI: 10.1145/2736277.2741081 (cited on p. 155).
- [115] F. Richard Guo, James McQueen, and Thomas S. Richardson. *Empirical Bayes for Large-scale Randomized Experiments: a Spectral Approach*. 2020. arXiv: 2002.02564 [stat.ME] (cited on p. 126).
- [116] Somit Gupta, Ronny Kohavi, Diane Tang, Ya Xu, Reid Andersen, Eytan Bakshy, Niall Cardin, Sumita Chandran, Nanyu Chen, Dominic Coey, Mike Curtis, Alex Deng, Weitao Duan, Peter Forbes, Brian Frasca, Tommy Guy, Guido W. Imbens, Guillaume Saint Jacques, Pranav Kantawala, Ilya Katsev, Moshe Katzwer, Mikael Konutgan, Elena Kunakova, Minyong Lee, MJ Lee, Joseph Liu, James McQueen, Amir Najmi, Brent Smith, Vivek Trehana, Lukas Vermeer, Toby Walker, Jeffrey Wong, and Igor Yashkov. "Top Challenges from the first Practical Online Controlled Experiments Summit". In: *ACM SIGKDD Explorations Newsletter* 21.1 (2019), pp. 20–35. ISSN: 1931-0145. DOI: 10.1145/3331651.3331655 (cited on p. 145).
- [117] Benjamin Han and Jared Gabor. *Contextual Bandits for Advertising Budget Allocation*. Presented in AdKDD 2020 Workshop (in conjunction with KDD '20). 2020. URL: <http://papers.adkdd.org/2020/papers/adkdd20-han-exploration.pdf> (cited on p. 196).
- [118] Anna Hart. "Mann-Whitney test is not just a test of medians: differences in spread can be important". In: *BMJ* 323.7309 (2001), pp. 391–393. ISSN: 0959-8138. DOI: 10.1136/bmj.323.7309.391 (cited on p. 103).
- [119] H. Leon Harter. "Expected values of normal order statistics". In: *Biometrika* 48.1-2 (1961), pp. 151–165. ISSN: 0006-3444. DOI: 10.1093/biomet/48.1-2.151 (cited on p. 42).

- [120] Larry V. Hedges and Ingram Olkin. “Chapter 5 - Estimation of a Single Effect Size: Parametric and Nonparametric Methods”. In: *Statistical Methods for Meta-Analysis*. San Diego: Academic Press, 1985, pp. 75–106. ISBN: 978-0-08-057065-5. DOI: 10.1016/B978-0-08-057065-5.50010-5 (cited on p. 74).
- [121] Miguel A. Hernán and James M. Robins. “Estimating causal effects from epidemiological data”. In: *Journal of Epidemiology & Community Health* 60.7 (2006), pp. 578–586. ISSN: 0143-005X. DOI: 10.1136/jech.2004.029496 (cited on p. 168).
- [122] Miguel A. Hernán and James M. Robins. *Causal Inference: What If*. Boca Raton, FL: Chapman & Hall/CRC, 2020. URL: https://www.hsph.harvard.edu/miguel-hernan/wp-content/uploads/sites/1268/2023/07/hernanrobins_WhatIf_19jul23.pdf (cited on pp. 135, 137, 138, 164).
- [123] Daniel N. Hill, Robert Moakler, Alan E. Hubbard, Vadim Tsemekhman, Foster Provost, and Kiril Tsemekhman. “Measuring Causal Impact of Online Actions via Natural Experiments: Application to Display Advertising”. In: *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. KDD ’15. Sydney, NSW, Australia: ACM, 2015, pp. 1839–1847. ISBN: 9781450336642. DOI: 10.1145/2783258.2788622 (cited on pp. 34, 168).
- [124] Kevin Hillstrom. *The MineThatData E-Mail Analytics And Data Mining Challenge*. Blog post. 2008. URL: <https://blog.minethatdata.com/2008/03/minethatdata-e-mail-analytics-and-data.html> (cited on p. 121).
- [125] John M. Hoenig and Dennis M. Heisey. “The Abuse of Power”. In: *The American Statistician* 55.1 (2001), pp. 19–24. DOI: 10.1198/000313001300339897 (cited on p. 96).
- [126] Henning Hohnhold, Deirdre O’Brien, and Diane Tang. “Focusing on the Long-Term: It’s Good for Users and Business”. In: *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. KDD ’15. Sydney, NSW, Australia: ACM, 2015, pp. 1849–1858. ISBN: 9781450336642. DOI: 10.1145/2783258.2788583 (cited on pp. 37, 152, 153).
- [127] Chris C. Holmes, François Caron, Jim E. Griffin, and David A. Stephens. “Two-sample Bayesian Non-parametric Hypothesis Testing”. In: *Bayesian Analysis* 10.2 (2015), pp. 297–320. DOI: 10.1214/14-BA914 (cited on p. 127).
- [128] Sally Hopewell, Susan Dutton, Ly-Mee Yu, An-Wen Chan, and Douglas G Altman. “The quality of reports of randomised trials in 2000 and 2006: comparative study of articles indexed in PubMed”. In: *BMJ* 340 (2010). ISSN: 0959-8138. DOI: 10.1136/bmj.c723 (cited on p. 118).
- [129] Henry Hsu and Peter A. Lachenbruch. “Paired t Test”. In: *Encyclopedia of Biostatistics*. Ed. by Peter Armitage and Theodore Colton. John Wiley & Sons, Ltd, 2005. ISBN: 9780470011812. DOI: 10.1002/0470011815.b2a15112 (cited on p. 82).
- [130] William A. Huber. *Justification of one-tailed hypothesis testing*. Cross Validated (Version: 2011-03-03). User handle: whuber. 2011. URL: <https://stats.stackexchange.com/q/7856> (accessed 17/7/2023) (cited on p. 90).
- [131] Guido W. Imbens and Donald B. Rubin. *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*. Cambridge University Press, 2015. DOI: 10.1017/CB09781139025751 (cited on pp. 137, 149, 164).
- [132] International Baccalaureate Organization. *Diploma Programme — Mathematics: applications and interpretation guide*. International Baccalaureate Organization (UK) Ltd, 2019 (cited on p. 68).

- [133] Shan Jiang, John Martin, and Christo Wilson. "Who's the Guinea Pig? Investigating Online A/B/n Tests in-the-Wild". In: *Proceedings of the Conference on Fairness, Accountability, and Transparency*. FAT* '19. Atlanta, GA, USA: ACM, 2019, pp. 201–210. ISBN: 9781450361255. DOI: 10.1145/3287560.3287565 (cited on pp. 120, 143).
- [134] Ramesh Johari, Pete Koomen, Leonid Pekelis, and David Walsh. "Peeking at A/B Tests: Why It Matters, and What to Do about It". In: *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. KDD '17. Halifax, NS, Canada: ACM, 2017, pp. 1517–1525. ISBN: 9781450348874. DOI: 10.1145/3097983.3097992 (cited on pp. 34, 105, 107, 124, 126, 129).
- [135] Garrett A. Johnson. "Inferno: A guide to field experiments in online display advertising". In: *Journal of Economics & Management Strategy* 32.3 (2023), pp. 469–490. DOI: 10.1111/jems.12513 (cited on p. 174).
- [136] Garrett A. Johnson, Randall A. Lewis, and Elmar I. Nubbemeyer. "The Online Display Ad Effectiveness Funnel & Carry-Over: A Meta-Study of Ghost Ad Experiments". In: *SSRN Electronic Journal* (2017). ISSN: 1556-5068. DOI: 10.2139/ssrn.2701578 (cited on pp. 36, 40, 45, 57, 58, 60, 61, 152).
- [137] Richard A. Johnson and Dean W. Wichern. *Applied Multivariate Statistical Analysis*. Pearson Prentice Hall, 2007. ISBN: 978-0-13-187715-3 (cited on p. 165).
- [138] Valen E. Johnson and David Rossell. "On the use of Non-Local Prior Densities in Bayesian Hypothesis Tests". In: *Journal of the Royal Statistical Society Series B: Statistical Methodology* 72.2 (2010), pp. 143–170. ISSN: 1369-7412. DOI: 10.1111/j.1467-9868.2009.00730.x (cited on p. 112).
- [139] David S. Jones and Scott H. Podolsky. "The history and fate of the gold standard". In: *The Lancet* 385.9977 (2015), pp. 1502–1503. DOI: 10.1016/S0140-6736(15)60742-5 (cited on p. 142).
- [140] Jean Kaddour, Aengus Lynch, Qi Liu, Matt J. Kusner, and Ricardo Silva. *Causal Machine Learning: A Survey and Open Problems*. 2022. arXiv: 2206.15475 [cs.LG] (cited on p. 165).
- [141] Nathan Kallus, Aahlad Manas Puli, and Uri Shalit. "Removing Hidden Confounding by Experimental Grounding". In: *Advances in Neural Information Processing Systems*. Ed. by S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett. Vol. 31. Curran Associates, Inc., 2018. URL: https://proceedings.neurips.cc/paper_files/paper/2018/file/566f0ea4f6c2e947f36795c8f58ba901-Paper.pdf (cited on p. 196).
- [142] Brian Karrer, Liang Shi, Monica Bhole, Matt Goldman, Tyrone Palmer, Charlie Gelman, Mikael Konutgan, and Feng Sun. "Network Experimentation at Scale". In: *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. KDD '21. Virtual Event, Singapore: ACM, 2021, pp. 3106–3116. ISBN: 9781450383325. DOI: 10.1145/3447548.3467091 (cited on p. 150).
- [143] Robert E. Kass and Adrian E. Raftery. "Bayes Factors". In: *Journal of the American Statistical Association* 90.430 (1995), pp. 773–795. ISSN: 0162-1459. DOI: 10.1080/01621459.1995.10476572 (cited on pp. 110, 129).
- [144] Robert E. Kass and Larry Wasserman. "The Selection of Prior Distributions by Formal Rules". In: *Journal of the American Statistical Association* 91.435 (1996), pp. 1343–1370. ISSN: 0162-1459. DOI: 10.1080/01621459.1996.10477003 (cited on p. 112).
- [145] Jouni Kerman, Peng Wang, and Jon Vaver. "Estimating Ad Effectiveness using Geo Experiments in a Time-Based Regression Framework". 2017. URL: <https://research.google/pubs/pub45950.pdf> (cited on p. 171).
- [146] Norbert L. Kerr. "HARKing: Hypothesizing After the Results are Known". In: *Personality and Social Psychology Review* 2.3 (1998), pp. 196–217. DOI: 10.1207/s15327957pspr0203_4 (cited on pp. 116, 123).

- [147] Tetiana Klimonova. *Grocery website data for AB test [Dataset]*. 2020. URL: <https://www.kaggle.com/tklimonova/grocery-website-data-for-ab-test/version/1> (accessed 25/8/2021) (cited on pp. 121, 122).
- [148] Ron Kohavi, Alex Deng, Brian Frasca, Roger Longbotham, Toby Walker, and Ya Xu. "Trustworthy Online Controlled Experiments: Five Puzzling Outcomes Explained". In: *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. KDD '12. Beijing, China: ACM, 2012, pp. 786–794. ISBN: 9781450314626. DOI: 10.1145/2339530.2339653 (cited on pp. 152, 153).
- [149] Ron Kohavi, Alex Deng, Brian Frasca, Toby Walker, Ya Xu, and Nils Pohlmann. "Online Controlled Experiments at Large Scale". In: *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. KDD '13. Chicago, Illinois, USA: ACM, 2013, pp. 1168–1176. ISBN: 9781450321747. DOI: 10.1145/2487575.2488217 (cited on pp. 18, 34, 152).
- [150] Ron Kohavi, Alex Deng, Roger Longbotham, and Ya Xu. "Seven Rules of Thumb for Web Site Experimenters". In: *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. KDD '14. New York, New York, USA: ACM, 2014, pp. 1857–1866. ISBN: 9781450329569. DOI: 10.1145/2623330.2623341 (cited on pp. 86, 100, 151).
- [151] Ron Kohavi, Alex Deng, and Lukas Vermeer. "A/B Testing Intuition Busters: Common Misunderstandings in Online Controlled Experiments". In: *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. KDD '22. Washington DC, USA: ACM, 2022, pp. 3168–3177. ISBN: 9781450393850. DOI: 10.1145/3534678.3539160 (cited on p. 96).
- [152] Ron Kohavi and Roger Longbotham. "Unexpected Results in Online Controlled Experiments". In: *ACM SIGKDD Explorations Newsletter* 12.2 (2011), pp. 31–35. ISSN: 1931-0145. DOI: 10.1145/1964897.1964905 (cited on pp. 152, 153).
- [153] Ron Kohavi and Roger Longbotham. "Online Controlled Experiments and A/B Testing". In: *Encyclopedia of Machine Learning and Data Mining*. Ed. by Claude Sammut and Geoffrey I. Webb. Boston, MA: Springer US, 2017, pp. 922–929. ISBN: 978-1-4899-7687-1. DOI: 10.1007/978-1-4899-7687-1_891 (cited on p. 145).
- [154] Ron Kohavi, Roger Longbotham, Dan Sommerfield, and Randal M. Henne. "Controlled experiments on the web: survey and practical guide". In: *Data Mining and Knowledge Discovery* 18 (2009), pp. 140–181. ISSN: 1573-756X. DOI: 10.1007/s10618-008-0114-1 (cited on p. 152).
- [155] Ron Kohavi, Diane Tang, and Ya Xu. *Trustworthy Online Controlled Experiments: A Practical Guide to A/B Testing*. 1st Edition. Cambridge University Press, 2020. ISBN: 9781108724265. URL: <https://experimentguide.com/> (cited on pp. 17, 69, 81, 99, 104, 119, 132, 139, 145, 152, 154, 164).
- [156] Ron Kohavi, Diane Tang, Ya Xu, Lars G. Hemkens, and John P. A. Ioannidis. "Online randomized controlled experiments at scale: lessons and extensions to medicine". In: *Trials* 21 (2020), p. 150. ISSN: 1745-6215. DOI: 10.1186/s13063-020-4084-y (cited on p. 18).
- [157] William H. Kruskal and W. Allen Wallis. "Use of Ranks in One-Criterion Variance Analysis". In: *Journal of the American Statistical Association* 47.260 (1952), pp. 583–621. ISSN: 0162-1459. DOI: 10.1080/01621459.1952.10483441 (cited on p. 101).
- [158] K. K. Gordon Lan and David L. DeMets. "Discrete sequential boundaries for clinical trials". In: *Biometrika* 70.3 (Dec. 1983), pp. 659–663. ISSN: 0006-3444. DOI: 10.1093/biomet/70.3.659 (cited on p. 108).

- [159] Nicholas Larsen, Jonathan Stallrich, Srijan Sengupta, Alex Deng, Ron Kohavi, and Nathaniel T. Stevens. "Statistical Challenges in Online Controlled Experiments: A Review of A/B Testing Methodology". In: *The American Statistician* (2023). issn: 0003-1305. doi: 10.1080/00031305.2023.2257237 (cited on pp. 145, 148, 150, 196).
- [160] Insa Lawler and Georg Zimmermann. "Misalignment Between Research Hypotheses and Statistical Hypotheses: A Threat to Evidence-Based Medicine?" In: *Topoi* 40 (2021), pp. 307–318. issn: 1572-8749. doi: 10.1007/s11245-019-09667-0 (cited on p. 70).
- [161] Minyong R. Lee and Milan Shen. "Winner's Curse: Bias Estimation for Total Effects of Features in Online Controlled Experiments". In: *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. KDD '18. London, United Kingdom: ACM, 2018, pp. 491–499. isbn: 9781450355520. doi: 10.1145/3219819.3219905 (cited on pp. 34, 147).
- [162] Erich L. Lehmann. "The Fisher, Neyman-Pearson Theories of Testing Hypotheses: One Theory or Two?" In: *Journal of the American Statistical Association* 88.424 (1993), pp. 1242–1249. issn: 0162-1459. doi: 10.1080/01621459.1993.10476404 (cited on pp. 71, 94).
- [163] Erich L. Lehmann and Howard J. M. D'Abrera. *Nonparametrics: statistical methods based on ranks*. Holden-Day series in probability and statistics. San Francisco: Holden-Day, 1975. isbn: 9780070370739 (cited on p. 102).
- [164] Erich L. Lehmann and Joseph P. Romano. *Testing Statistical Hypotheses*. Springer Texts in Statistics. New York, NY: Springer, 2006. isbn: 9780387276052. doi: 10.1007/0-387-27605-X (cited on p. 78).
- [165] Johannes Lenhard. "Models and Statistical Inference: The Controversy between Fisher and Neyman-Pearson". In: *The British Journal for the Philosophy of Science* 57.1 (2006), pp. 69–91. issn: 0007-0882. doi: 10.1093/bjps/axi152 (cited on p. 71).
- [166] Michael J. Lew. "A Reckless Guide to P-values". In: *Good Research Practice in Non-Clinical Pharmacology and Biomedicine*. Ed. by Anton Bespalov, Martin C. Michel, and Thomas Steckler. Cham: Springer International Publishing, 2020, pp. 223–256. isbn: 978-3-030-33656-1. doi: 10.1007/164_2019_286 (cited on pp. 68–70, 72, 106).
- [167] Hannah Li, Geng Zhao, Ramesh Johari, and Gabriel Y. Weintraub. "Interference, Bias, and Variance in Two-Sided Marketplace Experimentation: Guidance for Platforms". In: *Proceedings of the ACM Web Conference 2022*. WWW '22. Virtual Event, Lyon, France: ACM, 2022, pp. 182–192. isbn: 9781450390965. doi: 10.1145/3485447.3512063 (cited on p. 150).
- [168] Kevin Liou and Sean J. Taylor. "Variance-Weighted Estimators to Improve Sensitivity in Online Experiments". In: *Proceedings of the 21st ACM Conference on Economics and Computation*. EC '20. Virtual Event, Hungary: ACM, 2020, pp. 837–850. isbn: 9781450379755. doi: 10.1145/3391403.3399542 (cited on p. 151).
- [169] C. H. Bryan Liu. "On overlapping community-based networks: generation, detection, and their applications". MEng thesis. London, UK: Imperial College London, 2016. url: https://www.doc.ic.ac.uk/~mpd37/theses/2016_meng_liu.pdf (cited on p. 150).
- [170] C. H. Bryan Liu. *Probability that the same r.v. generates the rth order statistic in one noise-added set, and the sth order statistic in another noise-added set*. Cross Validated (Version: 2023-04-03). User handle: B.Liu. 2020. url: <https://stats.stackexchange.com/q/448802> (accessed 3/4/2023) (cited on p. 223).
- [171] C. H. Bryan Liu. *Variance of Normal Order Statistics*. Cross Validated (Version: 2022-09-25). User handle: B.Liu. 2022. url: <https://stats.stackexchange.com/q/394960> (cited on p. 48).

- [172] C. H. Bryan Liu, Elaine M. Bettaney, and Benjamin Paul Chamberlain. *Designing Experiments to Measure Incrementality on Facebook*. Presented in 2018 AdKDD & TargetAd Workshop (in conjunction with KDD '18). 2018. arXiv: 1806.02588 [stat.ME] (cited on pp. 60, 174, 175, 179, 184, 187).
- [173] C. H. Bryan Liu, Ângelo Cardoso, Paul Couturier, and Emma J. McCoy. "Datasets for Online Controlled Experiments". In: *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*. Ed. by J. Vanschoren and S. Yeung. Vol. 1. 2021. URL: <https://datasets-benchmarks-proceedings.neurips.cc/paper/2021/file/274ad4786c3abca69fa097b85867d9a4-Paper-round2.pdf> (cited on pp. 67, 115).
- [174] C. H. Bryan Liu and Benjamin Paul Chamberlain. "What is the Value of Experimentation & Measurement?" In: *2019 IEEE International Conference on Data Mining (ICDM)*. IEEE, 2019, pp. 1222–1227. ISBN: 978-1-7281-4604-1. DOI: 10.1109/ICDM.2019.00151 (cited on p. 33).
- [175] C. H. Bryan Liu, Benjamin Paul Chamberlain, and Emma J. McCoy. "What is the Value of Experimentation and Measurement?: Quantifying the Value and Risk of Reducing Uncertainty to Make Better Decisions". In: *Data Science and Engineering* 5 (2020), pp. 152–167. ISSN: 2364-1541. DOI: 10.1007/s41019-020-00121-5 (cited on pp. 33, 223).
- [176] C. H. Bryan Liu and Emma J. McCoy. *An Evaluation Framework for Personalization Strategy Experiment Designs*. Presented in AdKDD 2020 Workshop (In conjunction with KDD'20). 2020. arXiv: 2007.11638 [stat.ME] (cited on pp. 67, 118, 132, 173).
- [177] C. H. Bryan Liu and Emma J. McCoy. "Measuring E-Commerce Metric Changes in Online Experiments". In: *Companion Proceedings of the ACM Web Conference 2023*. WWW '23 Companion. Austin, TX, USA: ACM, 2023, pp. 495–499. ISBN: 9781450394192. DOI: 10.1145/3543873.3584654 (cited on pp. 67, 135).
- [178] Edward H. Livingston and Laura Cassidy. "Statistical Power and Estimation of the Number of Required Subjects for a Study Based on the t-Test: A Surgeon's Primer". In: *Journal of Surgical Research* 126.2 (2005), pp. 149–159. ISSN: 0022-4804. DOI: 10.1016/j.jss.2004.12.013 (cited on p. 69).
- [179] Luo Lu and Chuang Liu. "Separation strategies for three pitfalls in A/B testing". Presented in The Second Workshop on User Engagement Optimization at KDD 2014, New York City, USA. 2014. URL: http://www.ueo-workshop.com/wp-content/uploads/2014/04/Separation-strategies-for-three-pitfalls-in-AB-testing_withacknowledgments.pdf (cited on pp. 152, 153).
- [180] Jana Malesova. *Analyze A/B Test Results*. Code notebook. 2020. URL: <https://janamalesova.github.io/Udacity-Data-Analyst-Nanodegree/P2/> (accessed 9/9/2021) (cited on p. 122).
- [181] H. B. Mann and D. R. Whitney. "On a Test of Whether one of Two Random Variables is Stochastically Larger than the Other". In: *The Annals of Mathematical Statistics* 18.1 (1947), pp. 50–60. DOI: 10.1214/aoms/1177730491 (cited on p. 102).
- [182] Jialiang Mao and Iavor Bojinov. *Quantifying the Value of Iterative Experimentation*. 2021. arXiv: 2111.02334 [stat.AP] (cited on pp. 36, 37).
- [183] Paulo C. Marques F. *Size of a test and level of significance*. Cross Validated (Version: 2020-12-03). User handle: Zen. 2020. URL: <https://stats.stackexchange.com/q/51252> (accessed 17/7/2023) (cited on p. 93).
- [184] Ellen Marshall and Elizabeth Boggis. *The Statistics Tutor's Quick Guide to Commonly Used Statistical Tests*. 2016. URL: <https://www.statstutor.ac.uk/resources/uploaded/tutorsquickguidetostatistics.pdf> (cited on p. 73).

- [185] J. Nathan Matias, Kevin Munger, Marianne Aubin Le Quere, and Charles Ebersole. "The Upworthy Research Archive, a time series of 32,487 experiments in U.S. media". In: *Scientific Data* 8 (2021), p. 195. ISSN: 2052-4463. DOI: 10.1038/s41597-021-00934-7 (cited on pp. 121, 122, 127).
- [186] George Mavrotas and Olena Pechak. "The trichotomic approach for dealing with uncertainty in project portfolio selection: combining MCDA, mathematical programming and Monte Carlo simulation". In: *International Journal of Multicriteria Decision Making* 3.1 (2013), pp. 79–96. URL: <https://ideas.repec.org/a/ids/ijmcdm/v3y2013i1p79-96.html> (cited on p. 38).
- [187] Andrew McAfee and Erik Brynjolfsson. "Big Data: The Management Revolution". In: *Harvard Business Review* 90.10 (2012), pp. 60–68. URL: <https://hbr.org/2012/10/big-data-the-management-revolution> (cited on p. 17).
- [188] Conor McCabe. *Accelerating Geo-testing at ASOS*. Blog post. 2021. URL: <https://medium.com/asos-techblog/accelerating-geo-testing-at-asos-edbdf123a98b> (cited on pp. 169, 171).
- [189] Guy R. McPherson. "Teaching & Learning the Scientific Method". In: *The American Biology Teacher* 63.4 (2001), pp. 242–245. ISSN: 0002-7685. DOI: 10.2307/4451093 (cited on p. 70).
- [190] William Mendenhall, Robert J. Beaver, and Barbara M. Beaver. *Introduction to Probability and Statistics*. 15th. Cengage Learning, 2018. ISBN: 9781337554428 (cited on p. 69).
- [191] Merriam-Webster. *Measure Definition & Meaning - Merriam-Webster*. URL: <https://www.merriam-webster.com/dictionary/measure> (accessed 17/7/2023) (cited on p. 74).
- [192] Marius Messerschmied. *Udacity A/B Testing—Final Course Project, Version 11*. Code notebook. 2019. URL: <https://www.kaggle.com/mariusmesserschmied/udacity-a-b-testing-final-course-project?scriptVersionId=21267929> (accessed 9/9/2021) (cited on p. 122).
- [193] Alex Miller and Kartik Hosanagar. "A Meta-Analytic Investigation of p-hacking in E-commerce Experimentation". Working paper. 2020. URL: https://mackinstitute.wharton.upenn.edu/wp-content/uploads/2020/11/FP0398b_WP_2020Mar.pdf (cited on p. 129).
- [194] Evan Miller. *How Not To Run an A/B Test*. Blog post. 2010. URL: <https://www.evanmiller.org/how-not-to-run-an-ab-test.html> (cited on pp. 69, 99, 124, 188).
- [195] Evan Miller. *Simple Sequential A/B Testing*. Blog post. 2015. URL: <https://www.evanmiller.org/sequential-ab-testing.html> (cited on p. 106).
- [196] Ministry for Education, Youth and Sport of the State of Brandenburg (Ministerium für Bildung, Jugend und Sport des Landes Brandenburg). *Rahmenlehrplan für die gymnasiale Oberstufe – Teil C – Mathematik [Curriculum Framework for the Upper Level of the Gymnasium – Part C – Mathematics]*. 2022. URL: https://bildungsserver.berlin-brandenburg.de/fileadmin/bbb/unterricht/rahmenlehrplaene/gymnasiale_oberstufe/curricula/2022/Teil_C_RLP_GOST_2022_Mathematik.pdf (cited on p. 68).
- [197] Ministry of Education of the People's Republic of China (中华人民共和国教育部). *普通高中数学课程标准 (2017 年版 2020 年修订) [Curriculum Standard for Ordinary High School Mathematics (2017 Edition, Revised 2020)]*. People's Education Press (人民教育出版社), 2018. ISBN: 978-7-107-31842-9. URL: <https://www.ncet.edu.cn/Uploads/File/2022/04/28/%E6%99%AE%E9%80%9A%E9%AB%98%E4%B8%AD%E6%95%B0%E5%AD%A6%E8%AF%BE%E7%A8%8B%E6%A0%87%E5%87%86%EF%BC%882017%E5%B9%B4%E7%89%882020%E5%B9%B4%E4%BF%AE%E8%AE%A2%EF%BC%89.20220428174241.pdf> (cited on p. 68).

- [198] Ministry of Education, Culture, Sports, Science and Technology of Japan (文部科学省). 【数学編 理数編】 高等学校学習指導要領 (平成 30 年告示) 解説 [Commentary to the Senior High School Curriculum Guidelines (Announced 2018) - Mathematics Edition and Sciences & Mathematics Edition]. GAKKOTOSHO Co., Ltd. (学校図書), 2018. ISBN: 978-4762505355. URL: https://www.mext.go.jp/content/20230217-mxt_kyoiku02-100002620_05.pdf (cited on p. 68).
- [199] Richard D. Morey, Rink Hoekstra, Jeffrey N. Rouder, Michael D. Lee, and Eric-Jan Wagenmakers. “The fallacy of placing confidence in confidence intervals”. In: *Psychonomic Bulletin & Review* 23 (2016), pp. 103–123. ISSN: 1531-5320. DOI: 10.3758/s13423-015-0947-8 (cited on p. 76).
- [200] Denton E. Morrison and Ramon E. Henkel, eds. *The Significance Test Controversy: A Reader*. Chicago, USA: Aldine Publishing Co., 1970. ISBN: 978-0202300689 (cited on p. 68).
- [201] Stephen Mumford and Rani Lill Anjum. *Correlation is not causation*. Blog post. 2013. URL: <https://blog.oup.com/2013/11/correlation-is-not-causation/> (cited on pp. 137, 139).
- [202] Randall Munroe. *xkcd: Significant*. 2011. URL: <https://xkcd.com/882/> (accessed 25/8/2021) (cited on pp. 117, 146).
- [203] Randall Munroe. *xkcd: P-Values*. 2015. URL: <https://xkcd.com/1478/> (accessed 25/8/2021) (cited on p. 117).
- [204] Randall Munroe. *xkcd: Statistics*. 2020. URL: <https://xkcd.com/2400/> (accessed 14/7/2023) (cited on p. 73).
- [205] Purva Nahar. *Udacity Free Trial Screener Model Analysis, Version 1 [Dataset]*. 2021. URL: <https://www.kaggle.com/purvanahar/udacity-free-trial-screener-model-analysis/version/1> (accessed 9/9/2021) (cited on p. 122).
- [206] National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research, United States. *The Belmont report: Ethical principles and guidelines for the protection of human subjects of research*. 1979. URL: https://www.hhs.gov/ohrp/sites/default/files/the-belmont-report-508c_FINAL.pdf (cited on p. 132).
- [207] Jerzy Neyman. “Outline of a Theory of Statistical Estimation Based on the Classical Theory of Probability”. In: *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences* 236.767 (1937), pp. 333–380. DOI: 10.1098/rsta.1937.0005 (cited on p. 77).
- [208] Office for Human Research Protections, U.S. Department of Health and Human Services et al. *Federal Policy for the Protection of Human Subjects ('Common Rule')*. 2018. URL: <https://www.hhs.gov/ohrp/regulations-and-policy/regulations/common-rule/> (cited on p. 132).
- [209] Olist and André Sionek. *Brazilian E-Commerce Public Dataset by Olist [Dataset]*. 2018. DOI: 10.34740/KAGGLE/DSV/195341 (cited on p. 156).
- [210] Optimizely. *Multi-Armed Bandit*. 2021. URL: <https://www.optimizely.com/optimization-glossary/multi-armed-bandit/> (accessed 30/9/2023) (cited on p. 196).
- [211] Optimizely. *Statistical significance*. URL: <https://www.optimizely.com/optimization-glossary/statistical-significance/> (accessed 19/7/2023) (cited on p. 94).
- [212] Donald B. Owen. “Tables for Computing Bivariate Normal Probabilities”. In: *The Annals of Mathematical Statistics* 27.4 (1956), pp. 1075–1090. DOI: 10.1214/aoms/1177728074 (cited on p. 232).
- [213] Donald B. Owen. “A table of normal integrals”. In: *Communications in Statistics - Simulation and Computation* 9.4 (1980), pp. 389–419. DOI: 10.1080/03610918008812164 (cited on p. 231).

- [214] Mike Patefield and David Tandy. “Fast and Accurate Calculation of Owen’s T Function”. In: *Journal of Statistical Software* 5.5 (2000), pp. 1–25. ISSN: 1548-7660. DOI: 10.18637/jss.v005.i05 (cited on p. 232).
- [215] Grant Patience. *Udacity Data Analysis Nanodegree — Analyze AB Test Results*. Code repository. 2019. URL: https://github.com/patiegm/Udacity_Data_Analysis_Nanodegree/tree/master/Analyze%20AB%20Test%20Results (accessed 9/9/2021) (cited on p. 122).
- [216] Judea Pearl. “Causal inference in statistics: An overview”. In: *Statistics Surveys* 3 (2009), pp. 96–146. ISSN: 1935-7516. DOI: 10.1214/09-SS057 (cited on p. 164).
- [217] Judea Pearl. “Simpson’s Paradox, Confounding, and Collapsibility”. In: *Causality*. 2nd ed. Cambridge University Press, 2009, pp. 173–200. DOI: 10.1017/CB09780511803161.008 (cited on p. 166).
- [218] Judea Pearl, Madelyn Glymour, and Nicholas P. Jewell. *Causal Inference in Statistics: A Primer*. John Wiley & Sons, Inc., 2016. ISBN: 978-1-119-18684-7. URL: <https://www.wiley.com/en-us/Causal+Inference+in+Statistics%3A+A+Primer-p-9781119186847> (cited on p. 168).
- [219] Pearson Education. *Pearson Edexcel Level 3 Advanced GCE in Statistics (9ST0): Specification*. Issue 1. 2017. URL: <https://qualifications.pearson.com/content/dam/pdf/A%20Level/statistics/2017/Specification%20and%20Sample%20assessment%20material/A-level-Statistics-Specification.pdf> (cited on p. 68).
- [220] Stuart J. Pocock. “Group sequential methods in the design and analysis of clinical trials”. In: *Biometrika* 64.2 (1977), pp. 191–199. ISSN: 0006-3444. DOI: 10.1093/biomet/64.2.191 (cited on p. 108).
- [221] Jean Pouget-Abadie, Kevin Aydin, Warren Schudy, Kay Brodersen, and Vahab Mirrokni. “Variance Reduction in Bipartite Experiments through Correlation Clustering”. In: *Advances in Neural Information Processing Systems*. Ed. by H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett. Vol. 32. Curran Associates, Inc., 2019. URL: https://proceedings.neurips.cc/paper_files/paper/2019/file/bc047286b224b7bfa73d4cb02de1238d-Paper.pdf (cited on pp. 130, 150).
- [222] Alexey Poyarkov, Alexey Drutsa, Andrey Khalyavin, Gleb Gusev, and Pavel Serdyukov. “Boosted Decision Tree Regression Adjustment for Variance Reduction in Online Controlled Experiments”. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. KDD ’16. San Francisco, California, USA: ACM, 2016, pp. 235–244. ISBN: 9781450342322. DOI: 10.1145/2939672.2939688 (cited on pp. 34, 151, 174).
- [223] probabilityislogic et al. *Approximate order statistics for normal random variables*. Cross Validated (Version: 2015-05-15). User name unavailable. 2015. URL: <https://stats.stackexchange.com/a/9010/239410> (accessed 10/2/2024) (cited on p. 42).
- [224] Simon A. Queenborough. *FES 720 Introduction to R — Choosing a Statistical Test*. 2016. URL: <http://www.simonqueenborough.info/R/statistics/which-test.html> (cited on p. 73).
- [225] Federico Quin, Danny Weyns, Matthias Galster, and Camila Costa Silva. *A/B Testing: A Systematic Literature Review*. 2023. arXiv: 2308.04929 [cs.SE] (cited on p. 145).
- [226] Nicholas J. Radcliffe. “Using control groups to target on predicted lift: Building and assessing uplift model”. In: *Direct Marketing Analytics Journal* (2007), pp. 14–21 (cited on p. 121).
- [227] Mohsin Raza. *Analyze A/B Test Results, Version 1*. Code notebook. 2020. URL: <https://www.kaggle.com/razamh/analyze-a-b-test-results?scriptVersionId=41223618> (accessed 9/9/2021) (cited on p. 122).

- [228] Thomas S. Richardson, Yu Liu, James McQueen, and Doug Hains. “A Bayesian Model for Online Activity Sample Sizes”. In: *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*. Ed. by Gustau Camps-Valls, Francisco J. R. Ruiz, and Isabel Valera. Vol. 151. Proceedings of Machine Learning Research. PMLR, 2022, pp. 1775–1785. URL: <https://proceedings.mlr.press/v151/richardson22a.html> (cited on pp. 99, 159).
- [229] Herbert Robbins. “Statistical Methods Related to the Law of the Iterated Logarithm”. In: *The Annals of Mathematical Statistics* 41.5 (1970), pp. 1397–1409. DOI: 10.1214/aoms/1177696786 (cited on p. 106).
- [230] Christian P. Robert. *The Bayesian Choice: From Decision-Theoretic Foundations to Computational Implementation*. 2nd ed. Springer Texts in Statistics. New York, NY: Springer, 2007. ISBN: 978-0-387-71599-5. DOI: 10.1007/0-387-71599-1 (cited on p. 60).
- [231] John Rogers and Andrea Révész. “Experimental and quasi-experimental designs”. In: *The Routledge Handbook of Research Methods in Applied Linguistics*. Ed. by Jim McKinley and Heath Rose. London, UK: Routledge, 2019, pp. 133–143. ISBN: 9780367824471. DOI: 10.4324/9780367824471-12 (cited on p. 143).
- [232] David Rolnick, Kevin Aydin, Jean Pouget-Abadie, Shahab Kamali, Vahab Mirrokni, and Amir Najmi. “Randomized Experimental Design via Geographic Clustering”. In: *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. KDD ’19. Anchorage, AK, USA: ACM, 2019, pp. 2745–2753. ISBN: 9781450362016. DOI: 10.1145/3292500.3330778 (cited on p. 171).
- [233] Rasmus Ros and Per Runeson. “Continuous Experimentation and A/B Testing: A Mapping Study”. In: *Proceedings of the 4th International Workshop on Rapid Continuous Software Engineering*. RCoSE ’18. Gothenburg, Sweden: ACM, 2018, pp. 35–41. ISBN: 9781450357456. DOI: 10.1145/3194760.3194766 (cited on p. 118).
- [234] Paul R. Rosenbaum. “Interference Between Units in Randomized Experiments”. In: *Journal of the American Statistical Association* 102.477 (2007), pp. 191–200. ISSN: 0162-1459. DOI: 10.1198/016214506000001112 (cited on p. 149).
- [235] Paul R. Rosenbaum and Donald B. Rubin. “The central role of the propensity score in observational studies for causal effects”. In: *Biometrika* 70.1 (Apr. 1983), pp. 41–55. ISSN: 0006-3444. DOI: 10.1093/biomet/70.1.41 (cited on pp. 166, 168).
- [236] Tammy Rotem. *AB Testing With Python - Walkthrough Udacity's Course Final Project, Version 29*. Code notebook. 2018. URL: <https://www.kaggle.com/tammyrotem/ab-tests-with-python?scriptVersionId=2109873> (accessed 9/9/2021) (cited on p. 122).
- [237] Cynthia Rudin and Joanna Radin. “Why Are We Using Black Box Models in AI When We Don’t Need To? A Lesson From an Explainable AI Competition”. In: *Harvard Data Science Review* 1.2 (2019). DOI: 10.1162/99608f92.5a8a3a3d (cited on p. 19).
- [238] F. E. Satterthwaite. “An Approximate Distribution of Estimates of Variance Components”. In: *Biometrics Bulletin* 2.6 (1946), pp. 110–114. ISSN: 0099-4987. DOI: 10.2307/3002019 (cited on p. 83).
- [239] Martin Saveski, Jean Pouget-Abadie, Guillaume Saint-Jacques, Weitao Duan, Souvik Ghosh, Ya Xu, and Edoardo M. Airolidi. “Detecting Network Effects: Randomizing Over Randomized Experiments”. In: *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. KDD ’17. Halifax, NS, Canada: ACM, 2017, pp. 1027–1035. ISBN: 9781450348874. DOI: 10.1145/3097983.3098192 (cited on p. 150).
- [240] Fredrik Sävje. *Randomization does not imply unconfoundedness*. Presented at The Neglected Assumptions in Causal Inference Workshop @ ICML 2021. 2021. arXiv: 2107.14197 [stat.ME] (cited on pp. 140, 166).

- [241] Martin Schader and Friedrich Schmid. “Two Rules of Thumb for the Approximation of the Binomial Distribution by the Normal Distribution”. In: *The American Statistician* 43.1 (1989), pp. 23–24. ISSN: 0003-1305. DOI: 10.1080/00031305.1989.10475601 (cited on p. 89).
- [242] Møarten Schultzberg and Sebastian Ankargren. *Choosing a Sequential Testing Framework – Comparisons and Discussions*. Blog post. 2023. URL: <https://engineering.atspotify.com/2023/03/choosing-sequential-testing-framework-comparisons-and-discussions/> (cited on pp. 105, 106).
- [243] Helmut Schütz. *Power Analysis*. 2022. URL: <https://bebac.at/articles/Power-Analysis.phtml> (accessed 15/1/2024) (cited on p. 96).
- [244] Douglas Selent, Thanaporn Patikorn, and Neil Heffernan. “ASSISTments Dataset from Multiple Randomized Controlled Experiments”. In: *Proceedings of the Third (2016) ACM Conference on Learning @ Scale*. L@S ’16. Edinburgh, Scotland, UK: ACM, 2016, pp. 181–184. ISBN: 9781450337267. DOI: 10.1145/2876034.2893409 (cited on pp. 121, 122, 127).
- [245] M. Shakhshi-Niaei, S. A. Torabi, and S. H. Iranmanesh. “A Comprehensive Framework for Project Selection Problem Under Uncertainty and Real-world Constraints”. In: *Computers & Industrial Engineering* 61.1 (2011), pp. 226–237. ISSN: 0360-8352. DOI: 10.1016/j.cie.2011.03.015 (cited on p. 38).
- [246] William F. Sharpe. “Mutual Fund Performance”. In: *The Journal of Business* 39.1 (1966), pp. 119–138. ISSN: 0021-9398. URL: <http://www.jstor.org/stable/2351741> (cited on pp. 34, 39).
- [247] Zacks Shen. *Udacity AB Testing by Google Datasets, Version 1 [Dataset]*. 2021. URL: <https://www.kaggle.com/zacksshshen/udacity-ab-testing-by-google-datasets/version/1> (accessed 9/9/2021) (cited on p. 122).
- [248] Haolun Shi and Guosheng Yin. “Reconnecting p-Value and Posterior Probability Under One- and Two-Sided Tests”. In: *The American Statistician* 75.3 (2021), pp. 265–275. ISSN: 0003-1305. DOI: 10.1080/00031305.2020.1717621 (cited on p. 91).
- [249] Zbyněk Šidák. “Rectangular Confidence Regions for the Means of Multivariate Normal Distributions”. In: *Journal of the American Statistical Association* 62.318 (1967), pp. 626–633. ISSN: 0162-1459. DOI: 10.1080/01621459.1967.10482935 (cited on p. 146).
- [250] Betsy Sinclair, Margaret McConnell, and Donald P. Green. “Detecting Spillover Effects: Design and Analysis of Multilevel Experiments”. In: *American Journal of Political Science* 56.4 (2012), pp. 1055–1069. DOI: 10.1111/j.1540-5907.2012.00592.x (cited on p. 149).
- [251] Ashudeep Singh, David Kempe, and Thorsten Joachims. “Fairness in Ranking under Uncertainty”. In: *Advances in Neural Information Processing Systems*. Ed. by M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan. Vol. 34. Curran Associates, Inc., 2021, pp. 11896–11908. URL: https://proceedings.neurips.cc/paper_files/paper/2021/file/63c3ddcc7b23daa1e42dc41f9a44a873-Paper.pdf (cited on p. 38).
- [252] George Davey Smith and Shah Ebrahim. “Data dredging, bias, or confounding”. In: *BMJ* 325.7378 (2002), pp. 1437–1438. ISSN: 0959-8138. DOI: 10.1136/bmj.325.7378.1437 (cited on p. 92).
- [253] Michael E. Sobel. “What Do Randomized Studies of Housing Mobility Demonstrate?” In: *Journal of the American Statistical Association* 101.476 (2006), pp. 1398–1407. ISSN: 0162-1459. DOI: 10.1198/016214506000000636 (cited on p. 149).
- [254] The European Mathematical Society. “Efficiency, asymptotic”. In: *Encyclopedia of Mathematics*. 2011. URL: https://encyclopediaofmath.org/wiki/Efficiency,_asymptotic (cited on pp. 21, 179).

- [255] Mohamed A. Soliman and Ihab F. Ilyas. “Ranking with Uncertain Scores”. In: *2009 IEEE 25th International Conference on Data Engineering*. 2009, pp. 317–328. doi: 10.1109/ICDE.2009.102 (cited on p. 37).
- [256] Tim Stobierski. *The Advantages of Data-driven Decision-making*. Blog post on Harvard Business School Online. 2019. URL: <https://online.hbs.edu/blog/post/data-driven-decision-making> (cited on p. 17).
- [257] Aurelia Sui. *Mobile Games: A/B Testing, Version 1 [Dataset]*. 2019. URL: <https://www.kaggle.com/yufengsui/mobile-games-ab-testing/version/1> (accessed 25/8/2021) (cited on p. 122).
- [258] Lisa Sullivan. *Hypothesis Testing for Means & Proportions*. Part of Introduction to Biostatistics (SPH BS 704) in Boston University School of Public Health. 2017. URL: https://sphweb.bumc.bu.edu/otlt/MPH-Modules/BS/BS704_HypothesisTest-Means-Proportions/ (cited on p. 85).
- [259] Sean Talts, Michael Betancourt, Daniel Simpson, Aki Vehtari, and Andrew Gelman. *Validating Bayesian Inference Algorithms with Simulation-Based Calibration*. 2018. arXiv: 1804.06788 [stat.ME] (cited on pp. 54, 56, 190).
- [260] Diane Tang, Ashish Agarwal, Deirdre O’Brien, and Mike Meyer. “Overlapping Experiment Infrastructure: More, Better, Faster Experimentation”. In: *Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. KDD ’10. Washington, DC, USA: ACM, 2010, pp. 17–26. ISBN: 9781450300551. doi: 10.1145/1835804.1835810 (cited on pp. 18, 34, 152).
- [261] Thyago Tenório, Ig Ibert Bittencourt, Seiji Isotani, Alan Pedro, and Patrícia Ospina. “A gamified peer assessment model for on-line learning environments in a competitive context”. In: *Computers in Human Behavior* 64 (2016), pp. 247–263. ISSN: 0747-5632. doi: 10.1016/j.chb.2016.06.049 (cited on p. 122).
- [262] Thyago Tenório, Ig Ibert Bittencourt, Seiji Isotani, Alan Pedro, Patrícia Ospina, and Daniel Tenório. “Dataset of two experiments of the application of gamified peer assessment model into online learning environment MeuTutor”. In: *Data in Brief* 12 (2017), pp. 433–437. ISSN: 2352-3409. doi: 10.1016/j.dib.2017.04.032 (cited on pp. 121, 122, 127).
- [263] Stefan Thomke. *Experimentation Works: The Surprising Power of Business Experiments*. Boston, MA: Harvard Business Review Press, 2020. ISBN: 9781633697102. URL: <https://www.hbs.edu/faculty/Pages/item.aspx?num=57045> (cited on p. 145).
- [264] Udacity. *Online Data Analyst Course [MOOC]*. URL: <https://www.udacity.com/course/data-analyst-nanodegree--nd002> (accessed 25/8/2021) (cited on pp. 121, 122).
- [265] Graham Upton and Ian Cook. *A Dictionary of Statistics*. 3rd ed. Oxford University Press, 2014. ISBN: 9780199679188. doi: 10.1093/acref/9780199679188.001.0001 (cited on p. 112).
- [266] Joaquin Vanschoren, Jan N. van Rijn, Bernd Bischl, and Luis Torgo. “OpenML: Networked Science in Machine Learning”. In: *ACM SIGKDD Explorations Newsletter* 15.2 (2014), pp. 49–60. ISSN: 1931-0145. doi: 10.1145/2641190.2641198 (cited on pp. 116, 118).
- [267] Hal R. Varian. “Causal inference in economics and marketing”. In: *Proceedings of the National Academy of Sciences* 113.27 (2016), pp. 7310–7315. ISSN: 0027-8424. doi: 10.1073/pnas.1510479113 (cited on p. 165).
- [268] Jon Vaver and Jim Koehler. “Measuring Ad Effectiveness Using Geo Experiments”. 2011. URL: <https://storage.googleapis.com/pub-tools-public-publication-data/pdf/38355.pdf> (cited on pp. 169, 171).
- [269] Jon Vaver and Jim Koehler. “Periodic Measurement of Advertising Effectiveness Using Multiple-Test-Period Geo Experiments”. 2012. URL: <https://static.googleusercontent.com/media/research.google.com/en//pubs/archive/38356.pdf> (cited on p. 171).

- [270] Caio Waisman and Brett R. Gordon. *A multi-cell experimental design to recover policy relevant treatment effects, with an application to online advertising*. 2023. arXiv: 2302.13857 [econ.EM] (cited on pp. 144, 174).
- [271] Abraham Wald. “Sequential Tests of Statistical Hypotheses”. In: *The Annals of Mathematical Statistics* 16.2 (1945), pp. 117–186. DOI: 10.1214/aoms/1177731118 (cited on pp. 106, 124).
- [272] Esteban Walker and Amy S. Nowacki. “Understanding Equivalence and Noninferiority Testing”. In: *Journal of General Internal Medicine* 26 (2011), pp. 192–196. ISSN: 1525-1497. DOI: 10.1007/s11606-010-1513-8 (cited on p. 70).
- [273] Ella Wan. *A/B Testing for an Pre-enrollment Intervention in Udacity [Code repository]*. 2019. URL: <https://github.com/moggirain/A-B-Testing-for-Udacity> (accessed 9/9/2021) (cited on p. 122).
- [274] Zenan Wang, Xuan Yin, Tianbo Li, and Liangjie Hong. “Causal Meta-Mediation Analysis: Inferring Dose-Response Function From Summary Statistics of Many Randomized Experiments”. In: *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. KDD ’20. Virtual Event, CA, USA: ACM, 2020, pp. 2625–2635. ISBN: 9781450379984. DOI: 10.1145/3394486.3403313 (cited on pp. 130, 168).
- [275] Larry Wasserman. *All of Statistics: A Concise Course in Statistical Inference*. Springer Texts in Statistics. New York, NY: Springer New York, 2004. ISBN: 978-0-387-21736-9. DOI: 10.1007/978-0-387-21736-9 (cited on pp. 69, 85).
- [276] Ronald L. Wasserstein and Nicole A. Lazar. “The ASA Statement on p-Values: Context, Process, and Purpose”. In: *The American Statistician* 70.2 (2016), pp. 129–133. ISSN: 0003-1305. DOI: 10.1080/00031305.2016.1154108 (cited on pp. 68, 90, 92).
- [277] Ronald L. Wasserstein, Allen L. Schirm, and Nicole A. Lazar. “Moving to a World Beyond “ $p < 0.05$ ””. In: *The American Statistician* 73.sup1 (2019), pp. 1–19. ISSN: 0003-1305. DOI: 10.1080/00031305.2019.1583913 (cited on p. 68).
- [278] Martin Weber. “Decision making with incomplete information”. In: *European Journal of Operational Research* 28.1 (1987), pp. 44–57. ISSN: 0377-2217. DOI: 10.1016/0377-2217(87)90168-8. URL: <https://www.sciencedirect.com/science/article/pii/0377221787901688> (cited on p. 37).
- [279] Bernard L. Welch. “The Generalization of ‘Student’s’ Problem When Several Different Population Variances are Involved”. In: *Biometrika* 34.1-2 (Jan. 1947), pp. 28–35. ISSN: 0006-3444. DOI: 10.1093/biomet/34.1-2.28 (cited on p. 82).
- [280] Wider Funnel. *The State of Experimentation Maturity 2018*. White paper. 2018. URL: <https://web.archive.org/web/20220503123305/https://www.widerfunnel.com/wp-content/uploads/2018/04/State-of-Experimentation-2018-Original-Research-Report.pdf> (cited on p. 152).
- [281] Frank Wilcoxon. “Individual Comparisons by Ranking Methods”. In: *Biometrics Bulletin* 1.6 (1945), pp. 80–83. ISSN: 0099-4987. DOI: 10.2307/3001968 (cited on p. 101).
- [282] Morgan Wild and Marini Thorne. *A price of one’s own: An investigation into personalised pricing in essential markets*. 2018. URL: <https://www.citizensadvice.org.uk/Global/CitizensAdvice/Consumer%20publications/A%20price%20of%20one’s%20own%20final.pdf> (cited on p. 143).
- [283] Jeffrey Wong, Jasmine Nettiksimmons, Jiannan Lu, and Katherine Livins. *Addressing Hidden Imperfections in Online Experimentation*. Presented in 2021 Conference on Digital Experimentation @ MIT (CODE@MIT). 2022. arXiv: 2209.00649 [cs.SE] (cited on p. 150).

- [284] Kevin Wright. *agridat: Agricultural Datasets*. R package version 1.17. 2020. URL: <https://CRAN.R-project.org/package=agridat> (cited on p. 121).
- [285] Karl L. Wuensch. *An Overview of Power Analysis*. Part of Psychological Statistics (PSYC 2101) in East Carolina University Department of Psychology. 2020. URL: https://core.ecu.edu/wuenschk/stathelp/poweranalysis_overview.pdf (accessed 15/1/2024) (cited on p. 96).
- [286] Huizhi Xie and Juliette Aurisset. “Improving the Sensitivity of Online Controlled Experiments: Case Studies at Netflix”. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. KDD ’16. San Francisco, California, USA: ACM, 2016, pp. 645–654. ISBN: 9781450342322. DOI: 10.1145/2939672.2939733 (cited on pp. 34, 100, 151).
- [287] Ya Xu and Nanyu Chen. “Evaluating Mobile Apps with A/B and Quasi A/B Tests”. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. KDD ’16. San Francisco, California, USA: ACM, 2016, pp. 313–322. ISBN: 9781450342322. DOI: 10.1145/2939672.2939703 (cited on pp. 142, 168).
- [288] Ya Xu, Nanyu Chen, Addrian Fernandez, Omar Sinno, and Anmol Bhasin. “From Infrastructure to Culture: A/B Testing Challenges in Large Scale Social Networks”. In: *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. KDD ’15. Sydney, NSW, Australia: ACM, 2015, pp. 2227–2236. ISBN: 9781450336642. DOI: 10.1145/2783258.2788602 (cited on pp. 18, 34, 152).
- [289] Yueping Xu, Yeou-Koung Tung, Jia Li, and Shaofeng Niu. “Alternative risk measure for decision-making under uncertainty in water management”. In: *Progress in Natural Science* 19.1 (2009), pp. 115–119. ISSN: 1002-0071. DOI: <https://doi.org/10.1016/j.pnsc.2008.07.005> (cited on p. 37).
- [290] Liuyi Yao, Zhixuan Chu, Sheng Li, Yaliang Li, Jing Gao, and Aidong Zhang. “A Survey on Causal Inference”. In: *ACM Transactions on Knowledge Discovery from Data* 15.5 (2021). ISSN: 1556-4681. DOI: 10.1145/3444944 (cited on p. 164).
- [291] Mürşide Yarkin. *Mobile Games A/B Testing - Cookie Cats, Version 1 [Dataset]*. 2021. URL: <https://www.kaggle.com/mursideyarkin/mobile-games-ab-testing-cookie-cats/version/1> (accessed 25/8/2021) (cited on p. 122).
- [292] Scott W. H. Young. “Improving Library User Experience with A/B Testing: Principles and Process”. In: *Weave: Journal of Library User Experience* 1.1 (2014). DOI: 10.3998/weave.12535642.0001.101 (cited on p. 122).
- [293] Scott W. H. Young. *Improving Library User Experience with A/B Testing: Principles and Process [Dataset]*. 2014. DOI: 10.15788/m2rp42 (cited on pp. 121, 122).
- [294] Sheena Yu. *A/B Testing for Udacity Free Trial Screener*. Code notebook. 2017. URL: https://rstudio-pubs-static.s3.amazonaws.com/294845_0216f40bdc6b49c6a7d3fa0e5e0c5470.html (accessed 9/9/2021) (cited on p. 122).
- [295] Karen K. Yuen. “The two-sample trimmed t for unequal population variances”. In: *Biometrika* 61.1 (1974), pp. 165–170. ISSN: 0006-3444. DOI: 10.1093/biomet/61.1.165 (cited on p. 82).
- [296] Jing Zhou, Jiannan Lu, and Anas Shallah. “All about Sample-Size Calculations for A/B Testing: Novel Extensions & Practical Guide”. In: *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*. CIKM ’23. Birmingham, United Kingdom: ACM, 2023, pp. 3574–3583. ISBN: 9798400701245. DOI: 10.1145/3583780.3614779 (cited on p. 99).

- [297] Stephen Ziliak and Deirdre N. McCloskey. *The Cult of Statistical Significance: How the Standard Error Costs Us Jobs, Justice, and Lives*. Economics, Cognition, And Society. University of Michigan Press, 2008. ISBN: 978-0-472-05007-9. DOI: 10.3998/mpub.186351 (cited on p. 68).
- [298] Or Zuk, Liat Ein-Dor, and Eytan Domany. "Ranking under Uncertainty". In: *Proceedings of the Twenty-Third Conference on Uncertainty in Artificial Intelligence*. UAI'07. Vancouver, BC, Canada: AUAI Press, 2007, pp. 466–474. ISBN: 0-9749039-3-0. URL: <https://arxiv.org/ftp/arxiv/papers/1206/1206.5280.pdf> (cited on pp. 37, 38).

Appendix A

Miscellaneous Mathematical Results

A.1 Ranking Under Lower Uncertainty: Probability in Generating Specific Order Statistic Pairs

We consider the problem posed in Section 2.5.4, namely the probability that two order statistics in the ranking under lower uncertainty problem, one from the high-noise set and the other from the low-noise set, are generated from the same random variable.

This appendix is adapted from the appendix of the research paper “What is the Value of Experimentation and Measurement?” published in *Data Science and Engineering* [175]. It also incorporates relevant discussions on Cross Validated Stack Exchange [170].

We restate the setup and problem formally, which is also re-illustrated in Figure A.1. Let

$$\mathcal{V}_i \stackrel{\text{i.i.d.}}{\sim} F_{\mathcal{V}}(\cdot), \text{ where } \mathbb{E}(\mathcal{V}_i) = \mu_{\mathcal{V}}, \text{Var}(\mathcal{V}_i) = \sigma_{\mathcal{V}}^2, \quad (*2.5)$$

$$\epsilon_{1i} \stackrel{\text{i.i.d.}}{\sim} F_{\epsilon_1}(\cdot), \text{ where } \mathbb{E}(\epsilon_{1i}) = \mu_{\epsilon}, \text{Var}(\epsilon_{1i}) = \sigma_1^2,$$

$$\epsilon_{2i} \stackrel{\text{i.i.d.}}{\sim} F_{\epsilon_2}(\cdot), \text{ where } \mathbb{E}(\epsilon_{2i}) = \mu_{\epsilon}, \text{Var}(\epsilon_{2i}) = \sigma_2^2, \quad (*2.18)$$

where $\mathcal{V}_i \perp \epsilon_{1j}$, $\mathcal{V}_i \perp \epsilon_{2j}$, and $\epsilon_{1i} \perp \epsilon_{2j} \forall i, j$. $\mathcal{V}_i$ represents the true value of the i^{th} item under consideration in the ranking under lower uncertainty problem, while ϵ_{1i} and ϵ_{2i} represents




Figure A.1: (L) Figure 2.3 – The generative model for the ranking under lower uncertainty problem in plate notation. (R) RHS of Figure 2.5 – Relationship between different quantities in a three-item generative model. $\mathcal{V}_i$, $\mathcal{H}_i/\mathcal{L}_i$, and $\mathcal{H}_{(r)}/\mathcal{L}_{(s)}$, $i, r, s \in \{1, 2, 3\}$, represent the true value, the unranked noisy estimates, and the ranked noisy estimates of the items, respectively. Here $\mathcal{H}_{(r)}$ and $\mathcal{L}_{(s)}$ may be generated by the same $\mathcal{V}_i$ for some r and s combinations.

noise under different noise levels. We further let

$$\mathcal{H}_i = \mathcal{V}_i + \epsilon_{1i}, \quad \mathcal{L}_i = \mathcal{V}_i + \epsilon_{2i} \quad \forall i. \quad (*2.19)$$

Here, $\mathcal{H}_i$ and $\mathcal{L}_i$ are the estimated values for $\mathcal{V}_i$ under high and low estimation noise, respectively. We rank $\mathcal{H}$ and $\mathcal{L}$ separately to obtain two sets of order statistics: $\mathcal{H}_{(r)}$ and $\mathcal{L}_{(s)}$. We then introduce two functions, $\mathcal{I} : \{1, \dots, N\} \rightarrow \{1, \dots, N\}$ and $\mathcal{J} : \{1, \dots, N\} \rightarrow \{1, \dots, N\}$, which we use to map the ranks of $\mathcal{H}$ and $\mathcal{L}$ back to their original indices (and thus the indices of $\mathcal{V}$), respectively. We are interested in $\mathbb{P}(\mathcal{I}(r) = \mathcal{J}(s))$, the probability that the r^{th} ranked $\mathcal{H}$ and the s^{th} ranked $\mathcal{L}$ are generated by the same $\mathcal{V}_i$.

Before deriving the probability, we obtain some intuition on what we are dealing with via some simulations. In each run within a simulation, we:

1. Given N , μ_γ , μ_ϵ , σ_γ^2 , σ_1^2 , and σ_2^2 , generate $\mathcal{V}_i$, $\mathcal{H}_i$, and $\mathcal{L}_i \forall i \in \{1, \dots, N\}$ as specified above.
2. Obtain the rank for each realised $\mathcal{H}$ while preserving the order of the data array (i.e., the indices). Do the same for $\mathcal{L}$.
3. For each index i , obtain the rank of $\mathcal{H}_i$ as r and the rank of $\mathcal{L}_i$ as s . The pair (r, s) is said to




Figure A.2: Heatmap showing the empirical distribution of $\mathbb{P}(\mathcal{I}(r) = \mathcal{J}(s))$ across different r and s , generated from 200 000 simulation runs with $N = 25$, $\mu_{\mathcal{V}} = \mu_{\epsilon} = 0$, $\sigma_{\mathcal{V}}^2 = 1$, $\sigma_1^2 = 0.5$, and $\sigma_2^2 = 0.4$. Here, all $\mathcal{V}_i$, $\mathcal{H}_i$, and $\mathcal{L}_i$ follow the normal distribution. Every row and column in the heatmap sums to one.

have gained a “hit”.

After many runs, we obtain an empirical distribution of $\mathbb{P}(\mathcal{I}(r) = \mathcal{J}(s))$ across all possible (r, s) pairs by dividing the number of “hits” for each rank pair by the total number of simulation runs.¹ We observe that each (r, s) pair gets N chances to gain a “hit” within each simulation run, with the restriction that each r (and each s) would ultimately gain one “hit”. The latter owes to the fact that each rank within $\mathcal{H}$ and $\mathcal{L}$ should appear once and only once.

Figures A.2 and A.3 show the empirical $\mathbb{P}(\mathcal{I}(r) = \mathcal{J}(s))$ for a particular ranking under lower uncertainty model. We observe from Figure A.2 that the probability is clearly not uniform

¹ The statement “the r^{th} ranked $\mathcal{H}$ and the s^{th} ranked $\mathcal{L}$ are generated by the same $\mathcal{V}_i$ ” is equivalent to “the same $\mathcal{V}_i$ generates the r^{th} ranked $\mathcal{H}$ and the s^{th} ranked $\mathcal{L}$ ”, and thus Step 3 above provides a more computationally efficient way to obtain the empirical distribution – we do not need to determine the index mapping functions $\mathcal{I}(\cdot)$ and $\mathcal{J}(\cdot)$.




Figure A.3: The marginal distribution of $\mathbb{P}(\mathcal{I}(r) = \mathcal{J}(s))$ across different s for a given r , generated from 200 000 simulation runs with $N = 25$, $\mu_{\mathcal{V}} = \mu_{\epsilon} = 0$, $\sigma_{\mathcal{V}}^2 = 1$, $\sigma_1^2 = 0.5$, and $\sigma_2^2 = 0.4$. Here, all $\mathcal{V}_i$, $\mathcal{H}_i$, and $\mathcal{L}_i$ follow the normal distribution. Each plot represents a column in the heatmap in Figure A.2 and has a different scale on the y-axis to emphasise the relative difference in the probability across different s for a given r .

and higher when r and s are close to each other (i.e., along the $r = s$ diagonal), reaching the maximum when $r = s = 1$ and $r = s = N$. This agrees with the intuition that if $\mathcal{H}_i$ is the largest (smallest) amongst $\mathcal{H}$, chances are that the associated $\mathcal{V}_i$ is quite large (small) as well, and thus the associated $\mathcal{L}_i$ stands a high chance of becoming the largest (smallest) amongst $\mathcal{L}$. On the other hand, the probability is lower when r and s are far apart, reaching the minimum when $(r, s) = (1, N)$ or $(N, 1)$. This also makes sense – for $\mathcal{V}_i$ to generate both $\mathcal{H}_{(N)}$ and $\mathcal{L}_{(1)}$, one requires a large realisation for ϵ_{1i} and a large but negative realisation for ϵ_{2i} . The two realisations rarely happen together.

We also note the similarity between the marginal distributions in Figure A.3 (i.e., the probability across s for a given r) and the family of beta-binomial distributions. This is not surprising: beta distributions are closely connected to order statistics. As shown below, the observation will shape our approach to obtaining the probability.

A.1.1 Representing the probability as a certain number of successes in multiple Bernoulli trials

To obtain $\mathbb{P}(\mathcal{I}(r) = \mathcal{J}(s))$, one can express it as follow:

$$\begin{aligned} \mathbb{P}(\mathcal{I}(r) = \mathcal{J}(s)) &= \sum_{k=1}^N \mathbb{P}(\mathcal{I}(r) = k \cap \mathcal{J}(s) = k) && \text{(law of total prob.)} \\ &= \sum_{k=1}^N \mathbb{P}(\mathcal{J}(s) = k | \mathcal{I}(r) = k) \mathbb{P}(\mathcal{I}(r) = k) && \text{(conditional prob.)} \end{aligned} \quad (\text{A.1.1})$$

We first focus on the event $\mathcal{J}(s) = k$. Given any k , the event is equivalent to precisely $(s - 1)$ $\{\mathcal{L}_j\}_{j \neq k}$ being less than $\mathcal{L}_k$, which makes $\mathcal{L}_k$ the s^{th} ranked $\mathcal{L}$. Thus, we can express Equation (A.1.1) as

$$\mathbb{P}(\mathcal{I}(r) = \mathcal{J}(s)) = \sum_{k=1}^N \mathbb{P} \left(\sum_{j=1, j \neq k}^N \mathbb{I}_{\{\mathcal{L}_j < \mathcal{L}_k\}} = s - 1 \mid \mathcal{I}(r) = k \right) \mathbb{P}(\mathcal{I}(r) = k). \quad (\text{A.1.2})$$

We then recognise:

1. $\mathbb{P}(\mathcal{I}(r) = k) = \frac{1}{N} \forall k$. This is because all $\mathcal{V}_i$ and $\mathcal{H}_i$ are i.i.d. As a result, each $\mathcal{H}_i$ is equally

likely to have generated the order statistic $\mathcal{H}_{(r)}$;

2. The conditional probability within the outer summation does not depend on the index k but the rank r ; and
3. Given $k = \mathcal{I}(r)$, this particular $\mathcal{L}_k$ is now different from all the other $\{\mathcal{L}_j\}_{j \neq k}$, as the former now carries information on its possible values from $\mathcal{H}_{(r)}$ via $\mathcal{V}_{\mathcal{I}(r)}$. To stress the difference, we perform the index substitution here and refer to the random variable as $\mathcal{L}_{\mathcal{I}(r)}$ from now on.

This enables us to further simplify Equation (A.1.2) as:

$$\begin{aligned} \mathbb{P}(\mathcal{I}(r) = \mathcal{J}(s)) &= \sum_{k=1}^N \left(\mathbb{P} \left(\sum_{j=1, j \neq \mathcal{I}(r)}^N \mathbb{I}_{\{\mathcal{L}_j < \mathcal{L}_{\mathcal{I}(r)}\}} = s - 1 \right) \cdot \frac{1}{N} \right) \\ &= \mathbb{P} \left(\underbrace{\sum_{j=1, j \neq \mathcal{I}(r)}^N \mathbb{I}_{\{\mathcal{L}_j < \mathcal{L}_{\mathcal{I}(r)}\}}}_{\triangleq C} = s - 1 \right), \end{aligned} \quad (\text{A.1.3})$$

where we also define C (the count) as the summation of the indicator variables.

From (A.1.3), we observe that C represents the number of successes in $(N - 1)$ independent Bernoulli trials (i.e., $\mathcal{L}_j < \mathcal{L}_{\mathcal{I}(r)}$). Moreover, the success probability of each Bernoulli trial is not fixed but depends on the value of $\mathcal{L}_{\mathcal{I}(r)}$, itself a random variable. We know that for any given $l \in \mathbb{R}$, the probability of $\mathcal{L}_j$ being less than l is, by definition, $F_{\mathcal{L}_j}(l)$. Thus, we can define another random variable $\mathcal{P} = F_{\mathcal{L}_j}(\mathcal{L}_{\mathcal{I}(r)})$ to show how the success probability in the Bernoulli trials is distributed, with the probability density function (PDF) being

$$f_{\mathcal{P}}(p | r, \dots) = \frac{f_{\mathcal{L}_{\mathcal{I}(r)}}(F_{\mathcal{L}_j}^{-1}(p) | r, \dots)}{f_{\mathcal{L}_j}(F_{\mathcal{L}_j}^{-1}(p) | \dots)}, \quad (\text{A.1.4})$$

where we omit other parameters in the ranking under lower uncertainty model such as $\mu_{\mathcal{V}}$, $\sigma_{\mathcal{V}}^2, \dots$, for clarity.

This PDF arises from two standard results:

1. (Transformation of random variables) If $Y = g(X)$ and $g^{-1}(\cdot)$ exists, then

$$f_Y(y) = f_X(g^{-1}(x)) \left| \frac{dg^{-1}(x)}{dx} \right|; \text{ and}$$

2. (Derivative of inverse functions) $\frac{dg^{-1}(x)}{dx} = \frac{1}{g'(g^{-1}(x))}$ for all functions $g(\cdot)$ if both the inverse $g^{-1}(\cdot)$ and the derivative $g'(\cdot)$ exist (which they do for $F_{\mathcal{L}_j}(\cdot)$).

C then follows the following compound probability distribution:

$$f_C(c | N - 1, r, \dots) = \int_0^1 \text{Bin}(c | N - 1, p) f_{\mathcal{P}}(p | r, \dots) dp, \quad (\text{A.1.5})$$

where $\text{Bin}(c | N - 1, p)$ represents the PDF of the binomial distribution with $N - 1$ trials and success probability p .

A.1.2 Modelling the success count distribution using beta-binomials

Deriving the exact distribution for $\mathcal{P}$ and C is beyond the scope of this work, and we believe the distributions are analytically intractable in many cases. To estimate the covariance in (2.37), we will model $\mathcal{P}$ as beta distributions, with the parameters $\alpha_{\mathcal{P}}$ and $\beta_{\mathcal{P}}$ obtained via the method of moments. We believe beta distributions are a natural choice for $\mathcal{P}$ as they are closely related to order statistics. Moreover, beta distributions are also conjugate priors to binomial distributions, which eases the computation of the probability masses for C .

To obtain the beta distribution parameters, we first require the mean and variance for $\mathcal{L}_{\mathcal{I}(r)}$ and $F_{\mathcal{L}_j}(\mathcal{L}_{\mathcal{I}(r)})$. We know that $\mathcal{L}_{\mathcal{I}(r)} = \mathcal{V}_{\mathcal{I}(r)} + \epsilon_{2(\mathcal{I}(r))}$, $\mathcal{V}_{\mathcal{I}(r)} \perp \epsilon_{2(\mathcal{I}(r))}$ from (2.19), and hence we have

$$\mathbb{E}(\mathcal{L}_{\mathcal{I}(r)}) = \mathbb{E}(\mathcal{V}_{\mathcal{I}(r)}) + \mathbb{E}(\epsilon_{2(\mathcal{I}(r))}) \approx \mu_{\mathcal{V}} + \frac{\sigma_{\mathcal{V}}^2}{\sigma_{\mathcal{V}}^2 + \sigma_1^2} \left(F_{\mathcal{H}}^{-1}\left(\frac{r - \alpha}{N - 2\alpha + 1}\right) - (\mu_{\mathcal{V}} + \mu_{\epsilon}) \right) + \mu_{\epsilon}, \quad (\text{A.1.6})$$

$$\begin{aligned} \text{Var}(\mathcal{L}_{\mathcal{I}(r)}) &= \text{Var}(\mathcal{V}_{\mathcal{I}(r)}) + \text{Var}(\epsilon_{2(\mathcal{I}(r))}) \\ &\approx \frac{\sigma_1^2 \sigma_{\mathcal{V}}^2}{\sigma_{\mathcal{V}}^2 + \sigma_1^2} + \frac{\sigma_{\mathcal{V}}^4}{(\sigma_{\mathcal{V}}^2 + \sigma_1^2)^2} \frac{r(N - r + 1)}{(N + 1)^2(N + 2)} \frac{1}{(f_{\mathcal{H}}(F_{\mathcal{H}}^{-1}(\frac{r}{N+1})))^2} + \sigma_2^2, \end{aligned} \quad (\text{A.1.7})$$

where $\mathbb{E}(\mathcal{V}_{\mathcal{I}(r)})$ and $\text{Var}(\mathcal{V}_{\mathcal{I}(r)})$ are obtained from (2.9) and (2.25), respectively.

We can then approximate the expected value and variance of $F_{\mathcal{L}_j}(\mathcal{L}_{\mathcal{I}(r)})$ using Taylor series

expansion:

$$\mathbb{E} (F_{\mathcal{L}_j}(\mathcal{L}_{\mathcal{I}(r)})) \approx F_{\mathcal{L}_j}(\mathbb{E}(\mathcal{L}_{\mathcal{I}(r)})) + \frac{1}{2} f'_{\mathcal{L}_j}(\mathbb{E}(\mathcal{L}_{\mathcal{I}(r)})) \cdot \text{Var}(\mathcal{L}_{\mathcal{I}(r)}) + \dots, \quad (\text{A.1.8})$$

$$\begin{aligned} \text{Var} (F_{\mathcal{L}_j}(\mathcal{L}_{\mathcal{I}(r)})) &\approx (f_{\mathcal{L}_j}(\mathbb{E}(\mathcal{L}_{\mathcal{I}(r)})))^2 \cdot \text{Var}(\mathcal{L}_{\mathcal{I}(r)}) + \\ &\quad \frac{1}{4} \left(f'_{\mathcal{L}_j}(\mathbb{E}(\mathcal{L}_{\mathcal{I}(r)})) \right)^2 \cdot \text{Var} \left((\mathcal{L}_{\mathcal{I}(r)} - \mathbb{E}(\mathcal{L}_{\mathcal{I}(r)}))^2 \right) + \dots, \end{aligned} \quad (\text{A.1.9})$$

where $f_{\mathcal{L}_j}(\cdot)$ and $f'_{\mathcal{L}_j}(\cdot)$ are the probability density function and its derivative for $\mathcal{L}_j$, respectively. We observe that the first-order approximation (a special case of the delta method) is insufficiently accurate when compared against simulation results (see Section A.1.4). This is likely due to $F_{\mathcal{L}_j}$ being non-linear. We thus recommend using a second or higher-order approximation.

Finally, we denote $\mu_{\mathcal{P}} \triangleq \mathbb{E} (F_{\mathcal{L}_j}(\mathcal{L}_{\mathcal{I}(r)}))$, $\sigma_{\mathcal{P}}^2 \triangleq \text{Var} (F_{\mathcal{L}_j}(\mathcal{L}_{\mathcal{I}(r)}))$ and obtain the beta distribution parameters $\alpha_{\mathcal{P}}$ and $\beta_{\mathcal{P}}$ via the method of moments:

$$\alpha_{\mathcal{P}} = \left(\frac{1 - \mu_{\mathcal{P}}}{\sigma_{\mathcal{P}}^2} - \frac{1}{\mu_{\mathcal{P}}} \right) \mu_{\mathcal{P}}^2, \quad \beta_{\mathcal{P}} = \alpha_{\mathcal{P}} \left(\frac{1}{\mu_{\mathcal{P}}} - 1 \right). \quad (\text{A.1.10})$$

A.1.3 Estimation under normal assumptions

To complement the main text, we also discuss how the quantities derived above behave under normal assumptions. Firstly, (A.1.6) and (A.1.7) are now

$$\mathbb{E}(\mathcal{L}_{\mathcal{I}(r)}) \approx \mu_{\mathcal{V}} + \mu_{\epsilon} + \frac{\sigma_{\mathcal{V}}^2}{\sqrt{\sigma_{\mathcal{V}}^2 + \sigma_1^2}} \Phi^{-1} \left(\frac{r - \alpha}{N - 2\alpha + 1} \right), \quad (\text{A.1.11})$$

$$\text{Var}(\mathcal{L}_{\mathcal{I}(r)}) \approx \frac{\sigma_1^2 \sigma_{\mathcal{V}}^2}{\sigma_{\mathcal{V}}^2 + \sigma_1^2} + \frac{\sigma_{\mathcal{V}}^4}{\sigma_{\mathcal{V}}^2 + \sigma_1^2} \frac{r(N - r + 1)}{(N + 1)^2 (N + 2)} \frac{1}{(\phi(\Phi^{-1}(\frac{r}{N+1})))^2} + \sigma_2^2. \quad (\text{A.1.12})$$

We then recall from (2.19) that $\mathcal{L}_j \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(\mu_{\mathcal{V}} + \mu_{\epsilon}, \sigma_{\mathcal{V}}^2 + \sigma_2^2)$, and hence

$$F_{\mathcal{L}_j}(l) = \Phi \left(\frac{l - (\mu_{\mathcal{V}} + \mu_{\epsilon})}{\sqrt{\sigma_{\mathcal{V}}^2 + \sigma_2^2}} \right). \quad (\text{A.1.13})$$

This opens up multiple pathways to estimate $\mu_{\mathcal{P}}$ and $\sigma_{\mathcal{P}}^2$. Firstly, we can proceed with the Taylor series expansion approach, using the quantities derived above and noting that the

first and second derivatives of (A.1.13) are

$$f_{\mathcal{L}_j}(l) = \frac{1}{\sqrt{\sigma_{\mathcal{L}}^2 + \sigma_2^2}} \phi \left(\frac{l - (\mu_{\mathcal{V}} + \mu_{\epsilon})}{\sqrt{\sigma_{\mathcal{V}}^2 + \sigma_2^2}} \right), \text{ and} \quad (\text{A.1.14})$$

$$f'_{\mathcal{L}_j}(l) = -\frac{l - (\mu_{\mathcal{V}} + \mu_{\epsilon})}{(\sigma_{\mathcal{V}}^2 + \sigma_2^2)^{\frac{3}{2}}} \phi \left(\frac{l - (\mu_{\mathcal{V}} + \mu_{\epsilon})}{\sqrt{\sigma_{\mathcal{V}}^2 + \sigma_2^2}} \right). \quad (\text{A.1.15})$$

However, while the estimate for $\mathbb{E}(F_{\mathcal{L}_j}(\mathcal{L}_{\mathcal{I}(r)}))$ under this approach is reasonably accurate, we find the estimates for $\text{Var}(F_{\mathcal{L}_j}(\mathcal{L}_{\mathcal{I}(r)}))$ unsatisfactory in some cases, even when we involve the *sixth*-order Taylor polynomial (see Section A.1.4).

Instead, we estimate the quantities using Owen's work on integrals of Gaussian functions [213].

We first define $\mathcal{L}^*$, being $\mathcal{L}_{\mathcal{I}(r)}$ normalised by the parameters of $\mathcal{L}_j$:

$$\mathcal{L}^* \triangleq \frac{\mathcal{L}_{\mathcal{I}(r)} - (\mu_{\mathcal{V}} + \mu_{\epsilon})}{\sqrt{\sigma_{\mathcal{V}}^2 + \sigma_2^2}}. \quad (\text{A.1.16})$$

Since $\mathcal{L}_{\mathcal{I}(r)}$ is approximately normally distributed, $\mathcal{L}^*$ is also approximately normally distributed with mean and variance

$$\mu_{\mathcal{L}^*} = \frac{\mathbb{E}(\mathcal{L}_{\mathcal{I}(r)}) - (\mu_{\mathcal{V}} + \mu_{\epsilon})}{\sqrt{\sigma_{\mathcal{V}}^2 + \sigma_2^2}}, \quad \sigma_{\mathcal{L}^*}^2 = \frac{1}{\sigma_{\mathcal{V}}^2 + \sigma_2^2} \text{Var}(\mathcal{L}_{\mathcal{I}(r)}), \quad (\text{A.1.17})$$

where $\mathbb{E}(\mathcal{L}_{\mathcal{I}(r)})$ and $\text{Var}(\mathcal{L}_{\mathcal{I}(r)})$ are approximated in (A.1.11) and (A.1.12), respectively. This allows us to represent $\mathcal{L}^*$ by scaling a standard normal r.v. Z :

$$\mathcal{L}^* \approx \mu_{\mathcal{L}^*} + \sigma_{\mathcal{L}^*} Z. \quad (\text{A.1.18})$$

We can then write $F_{\mathcal{L}_j}(\mathcal{L}_{\mathcal{I}(r)})$ as

$$F_{\mathcal{L}_j}(\mathcal{L}_{\mathcal{I}(r)}) = \Phi(\mathcal{L}^*) \approx \Phi(\mu_{\mathcal{L}^*} + \sigma_{\mathcal{L}^*} Z) \quad (\text{A.1.19})$$

by substituting, in turn, (A.1.13), (A.1.16) and (A.1.18) into the LHS of the equation.

We then make use of the following identities provided by Owen [213] (Equations 10010.8

and 20010.4):

$$\mathbb{E}(\Phi(\mu_{\mathcal{L}^*} + \sigma_{\mathcal{L}^*}Z)) = \Phi\left(\frac{\mu_{\mathcal{L}^*}}{\sqrt{1 + \sigma_{\mathcal{L}^*}^2}}\right), \quad (\text{A.1.20})$$

$$\mathbb{E}((\Phi(\mu_{\mathcal{L}^*} + \sigma_{\mathcal{L}^*}Z))^2) = \Phi\left(\frac{\mu_{\mathcal{L}^*}}{\sqrt{1 + \sigma_{\mathcal{L}^*}^2}}\right) - 2 \cdot T\left(\frac{\mu_{\mathcal{L}^*}}{\sqrt{1 + \sigma_{\mathcal{L}^*}^2}}, \frac{1}{\sqrt{1 + 2\sigma_{\mathcal{L}^*}^2}}\right), \quad (\text{A.1.21})$$

where $T(\cdot, \cdot)$ represents Owen's T function [212], of which a fast numerical algorithm is available from [214]. While the original work does not provide an identity for the variance of $\Phi(\mathcal{L}^*)$, we can obtain so from (A.1.20) and (A.1.21) easily:

$$\begin{aligned} & \text{Var}(\Phi(\mu_{\mathcal{L}^*} + \sigma_{\mathcal{L}^*}Z)) \\ &= \mathbb{E}((\Phi(\mu_{\mathcal{L}^*} + \sigma_{\mathcal{L}^*}Z))^2) - (\mathbb{E}(\Phi(\mu_{\mathcal{L}^*} + \sigma_{\mathcal{L}^*}Z)))^2 \\ &= \Phi\left(\frac{\mu_{\mathcal{L}^*}}{\sqrt{1 + \sigma_{\mathcal{L}^*}^2}}\right) \left(1 - \Phi\left(\frac{\mu_{\mathcal{L}^*}}{\sqrt{1 + \sigma_{\mathcal{L}^*}^2}}\right)\right) - 2 \cdot T\left(\frac{\mu_{\mathcal{L}^*}}{\sqrt{1 + \sigma_{\mathcal{L}^*}^2}}, \frac{1}{\sqrt{1 + 2\sigma_{\mathcal{L}^*}^2}}\right). \end{aligned} \quad (\text{A.1.22})$$

We finally substitute (A.1.20) and (A.1.22) into (A.1.10) to obtain the beta distribution parameters under normal assumptions.

A.1.4 Some basic validations

We close the investigation on the supposedly minor quantity in the ranking under lower uncertainty problem by validating the theoretical results derived above. We have accepted in Section A.1.2 that the probability $\mathbb{P}(\mathcal{I}(r) = \mathcal{J}(s))$ and its associated distributions are analytically intractable in many cases. As a result, a full validation across many parameter combinations (similar to what we have done for quantities such as $\mathbb{E}(\mathcal{V}_{\mathcal{I}(r)})$ and $\text{Var}(\mathcal{W}_{\mathcal{L}})$ in the main text) for what are best estimates is overkill. Instead, we will perform visual checks on whether the fitted distributions are reasonable approximation to the empirical distribution for a handful of parameter combinations under normal assumptions. This is not to say we have fully addressed the original question posed at the top of this chapter – there is still plenty to uncover for the said probability in terms of estimation method and interpretation – it is just not the primary quantity of interest for our ranking under lower uncertainty problem.

We also calculate the mean Kullback-Leibler (KL) divergence between the empirical distri-

bution and distributions produced from different modelling approaches to compare them quantitatively. Here, we define the mean KL divergence as

$$\begin{aligned} & \frac{1}{N} \sum_{r=1}^N D_{KL} \left(\hat{f}_{C|N-1,r,\dots} \parallel f_{\hat{C}|N-1,r,\dots} \right) \\ &= \frac{1}{N} \sum_{r=1}^N \sum_{s=1}^N \hat{f}_C(s-1 | N-1, r, \dots) \log \left(\frac{\hat{f}_C(s-1 | N-1, r, \dots)}{f_{\hat{C}}(s-1 | N-1, r, \dots)} \right), \end{aligned} \quad (\text{A.1.23})$$

where $\hat{f}_C(\cdot)$ denotes the probability mass function (PMF) of the empirical marginal distribution and $f_{\hat{C}}(\cdot)$ denotes the PMF of the fitted distribution. In other words, we calculate the KL divergence of each marginal distribution shown in Figure A.3 and average it over the N marginal distributions.²

We observe that the approach using Owen's integrals on Gaussian functions produces marginal distributions produces better fit on the empirical distribution than the approaches using Taylor series expansion up to the sixth-order Taylor polynomial, both visually (see Figure A.4) and in terms of the mean KL divergence (see Table A.1). We also observe the fit using Owen's integrals is less ideal (yet still reasonable) in cases where:

1. r or s is close to 1 or N ,
2. N is small, and
3. σ_X^2 and σ_1^2 are similar in magnitude.

These cases lead to a higher variance in order statistic and rank estimators and greater deviation from the normal assumptions, which make accurate estimation more difficult.

² We use the empirical marginal distribution as the target distribution for KL divergence as we believe the true distribution is analytically intractable in many cases. This means the average KL divergence may vary slightly from one simulation to another due to using fresh samples. We ensure such variance is negligible by performing hundreds of thousands of simulation runs when generating each empirical marginal distributions.




Figure A.4: The fitted marginal distributions – that using the Taylor series expansion up to the second-order (blue dotted line), fourth-order (olive dashed line), and sixth-order (purple dot-dashed line), as well as that using Owen's integrals of Gaussian functions (red solid line) – overlaid on Figure A.3, which shows the empirical marginal distributions of $\mathbb{P}(\mathcal{I}(r) = \mathcal{J}(s))$ (grey bars).




Figure A.5: The fitted marginal distributions – that using the Taylor series expansion up to the second-order (blue dotted line), fourth-order (olive dashed line), and sixth-order (purple dot-dashed line), as well as that using Owen's integrals of Gaussian functions (red solid line) – overlaid on the empirical marginal distributions of $\mathbb{P}(\mathcal{I}(r) = \mathcal{J}(s))$ (grey bars), with $N = 10$, $\mu_V = \mu_\epsilon = 0$, $\sigma_V^2 = 1$, $\sigma_1^2 = 0.5$, and $\sigma_2^2 = 0.4$.

Modelled distribution	Mean KL divergence	
	$N = 10$	$N = 25$
Uniform	0.24582 ± 0.00025	0.28358 ± 0.00018
Beta-binomial (2nd order Taylor SE)	0.05924 ± 0.00017	0.06190 ± 0.00007
Beta-binomial (4th order Taylor SE)	0.01070 ± 0.00004	0.01930 ± 0.00004
Beta-binomial (6th order Taylor SE)	0.03414 ± 0.00013	0.02029 ± 0.00004
Beta-binomial (Owen's integrals)	0.00784 ± 0.00006	0.00317 ± 0.00002

Table A.1: Mean Kullback-Leibler (KL) divergence between the empirical marginal distribution of $\mathbb{P}(\mathcal{I}(r) = \mathcal{J}(s))$ and distributions produced from different modelling approaches, with $\mu_V = \mu_\epsilon = 0$, $\sigma_V^2 = 1$, $\sigma_1^2 = 0.5$, and $\sigma_2^2 = 0.4$. Each entry records the mean and standard deviation of the mean KL divergence across multiple simulations.

A.2 Calculations Related to Dilution in Personalisation Strategy Experiment Designs

This appendix expands the exposition in Section 6.3.2 by providing the detailed algebraic manipulations of the inequalities involved. In the section, we discussed, under the context of designing experiments for personalisation strategies, the conditions where an experiment setup with treatment effect dilution (aka Setup 2) will emerge superior to one without treatment effect dilution (aka Setup 3) and vice versa.

We recap the mathematical arguments presented in the section. Under the evaluation framework introduced in Chapter 6, we can establish superiority of an experiment setup via two criteria. Setup 3 is superior to Setup 2 if at least one of the following inequalities hold:

First criterion: $\theta_{S3}^* < \theta_{S2}^*$ (as $\Delta_{S3} > \Delta_{S2}$ trivially holds),

Second criterion: $\Delta_{S3} - \Delta_{S2} > \theta_{S3}^* - \theta_{S2}^*$,

where Δ_{S2} , θ_{S2}^* , Δ_{S3} , and θ_{S3}^* are that stated in Equations (6.6), (6.7), (6.8), and (6.9) respectively – we will restate the equations below where appropriate. If neither inequality holds (and both sides are not equal), we consider Setup 2 as superior to Setup 3 under the second criterion as the following holds:

$$\Delta_{S3} - \Delta_{S2} < \theta_{S3}^* - \theta_{S2}^* \iff \Delta_{S2} - \Delta_{S3} > \theta_{S2}^* - \theta_{S3}^*.$$

This appendix only deals with the detailed algebraic manipulation. We refer readers back to Section 6.3.2 in the main text for the interpretations of and the insights from the resultant inequalities.

A.2.1 The first criterion

We first show $\theta_{S3}^* < \theta_{S2}^*$, where

$$\theta_{S2}^* = (z_{1-\frac{\alpha}{2}} - z_{1-\pi_{\min}}) \sqrt{\frac{2(n_0(2\sigma_{C0}^2) + n_1(\sigma_{I1}^2 + \sigma_{C1}^2) + n_2(\sigma_{C2}^2 + \sigma_{I2}^2) + n_3(\sigma_{I\phi}^2 + \sigma_{I\psi}^2))}{(n_0 + n_1 + n_2 + n_3)^2}}, \quad (*6.7)$$

$$\theta_{S3}^* = (z_{1-\frac{\alpha}{2}} - z_{1-\pi_{\min}}) \sqrt{\frac{2(n_1(\sigma_{I1}^2 + \sigma_{C1}^2) + n_2(\sigma_{C2}^2 + \sigma_{I2}^2) + n_3(\sigma_{I\phi}^2 + \sigma_{I\psi}^2))}{(n_1 + n_2 + n_3)^2}}, \quad (*6.9)$$

is equivalent to

$$\frac{(n_1(\sigma_{I1}^2 + \sigma_{C1}^2) + n_2(\sigma_{C2}^2 + \sigma_{I2}^2) + n_3(\sigma_{I\phi}^2 + \sigma_{I\psi}^2)) \cdot (n_0 + 2n_1 + 2n_2 + 2n_3)}{2(n_1 + n_2 + n_3)^2} < \sigma_{C0}^2. \quad (*6.12)$$

This is the condition that will lead to Setup 3 being superior to Setup 2 under the first criterion in our evaluation framework.

We start by writing the inequality $\theta_{S3}^* < \theta_{S2}^*$ in full:

$$\begin{aligned} & (z_{1-\frac{\alpha}{2}} - z_{1-\pi_{\min}}) \sqrt{\frac{2(n_1(\sigma_{I1}^2 + \sigma_{C1}^2) + n_2(\sigma_{C2}^2 + \sigma_{I2}^2) + n_3(\sigma_{I\phi}^2 + \sigma_{I\psi}^2))}{(n_1 + n_2 + n_3)^2}} \\ & < (z_{1-\frac{\alpha}{2}} - z_{1-\pi_{\min}}) \sqrt{\frac{2(n_0(2\sigma_{C0}^2) + n_1(\sigma_{I1}^2 + \sigma_{C1}^2) + n_2(\sigma_{C2}^2 + \sigma_{I2}^2) + n_3(\sigma_{I\phi}^2 + \sigma_{I\psi}^2))}{(n_0 + n_1 + n_2 + n_3)^2}}. \end{aligned} \quad (A.2.1)$$

Canceling the $z_{1-\frac{\alpha}{2}} - z_{1-\pi_{\min}}$ and $\sqrt{2}$ terms on both sides, then squaring both sides³ yields

$$\begin{aligned} & \frac{n_1(\sigma_{I1}^2 + \sigma_{C1}^2) + n_2(\sigma_{C2}^2 + \sigma_{I2}^2) + n_3(\sigma_{I\phi}^2 + \sigma_{I\psi}^2)}{(n_1 + n_2 + n_3)^2} \\ & < \frac{n_0(2\sigma_{C0}^2) + n_1(\sigma_{I1}^2 + \sigma_{C1}^2) + n_2(\sigma_{C2}^2 + \sigma_{I2}^2) + n_3(\sigma_{I\phi}^2 + \sigma_{I\psi}^2)}{(n_0 + n_1 + n_2 + n_3)^2}. \end{aligned} \quad (A.2.1a)$$

We then write $\xi = n_1(\sigma_{I1}^2 + \sigma_{C1}^2) + n_2(\sigma_{C2}^2 + \sigma_{I2}^2) + n_3(\sigma_{I\phi}^2 + \sigma_{I\psi}^2)$ and move the ξ term on the

³ The operation maintains the *iff* relation as the radicand is positive.

RHS to the LHS:

$$\xi \left(\frac{1}{(n_1 + n_2 + n_3)^2} - \frac{1}{(n_0 + n_1 + n_2 + n_3)^2} \right) < \frac{n_0(2\sigma_{C0}^2)}{(n_0 + n_1 + n_2 + n_3)^2}. \quad (\text{A.2.1b})$$

As the partial fractions on the LHS can be consolidated as

$$\begin{aligned} \frac{1}{(n_1 + n_2 + n_3)^2} - \frac{1}{(n_0 + n_1 + n_2 + n_3)^2} &= \frac{(n_0 + n_1 + n_2 + n_3)^2 - (n_1 + n_2 + n_3)^2}{(n_1 + n_2 + n_3)^2(n_0 + n_1 + n_2 + n_3)^2} \\ &= \frac{(n_0 + 2n_1 + 2n_2 + 2n_3)n_0}{(n_1 + n_2 + n_3)^2(n_0 + n_1 + n_2 + n_3)^2}, \end{aligned}$$

where the second step utilizes the identity $a^2 - b^2 = (a + b)(a - b)$, we can write Inequality (A.2.1b) as

$$\xi \left(\frac{(n_0 + 2n_1 + 2n_2 + 2n_3)n_0}{(n_1 + n_2 + n_3)^2(n_0 + n_1 + n_2 + n_3)^2} \right) < \frac{n_0(2\sigma_{C0}^2)}{(n_0 + n_1 + n_2 + n_3)^2}. \quad (\text{A.2.1c})$$

We finally cancel the n_0 and $(n_0 + n_1 + n_2 + n_3)^2$ terms on both sides, move the factor of two to the LHS, and write ξ in its full form to arrive at

$$\frac{(n_1(\sigma_{I1}^2 + \sigma_{C1}^2) + n_2(\sigma_{C2}^2 + \sigma_{I2}^2) + n_3(\sigma_{I\phi}^2 + \sigma_{I\psi}^2)) \cdot (n_0 + 2n_1 + 2n_2 + 2n_3)}{2(n_1 + n_2 + n_3)^2} < \sigma_{C0}^2,$$

which is identical to Inequality (6.12).

A.2.2 The second criterion

The master inequality We also show that the inequality $\Delta_{S3} - \Delta_{S2} > \theta_{S3}^* - \theta_{S2}^*$, where

$$\Delta_{S2} = \frac{n_1(\mu_{C1} - \mu_{I1}) + n_2(\mu_{I2} - \mu_{C2}) + n_3(\mu_{I\psi} - \mu_{I\phi})}{n_0 + n_1 + n_2 + n_3}, \quad (*6.6)$$

$$\Delta_{S3} = \frac{n_1(\mu_{C1} - \mu_{I1}) + n_2(\mu_{I2} - \mu_{C2}) + n_3(\mu_{I\psi} - \mu_{I\phi})}{n_1 + n_2 + n_3}, \quad (*6.8)$$

$$\theta_{S2}^* = (z_{1-\frac{\alpha}{2}} - z_{1-\pi_{\min}}) \sqrt{\frac{2(n_0(2\sigma_{C0}^2) + n_1(\sigma_{I1}^2 + \sigma_{C1}^2) + n_2(\sigma_{C2}^2 + \sigma_{I2}^2) + n_3(\sigma_{I\phi}^2 + \sigma_{I\psi}^2))}{(n_0 + n_1 + n_2 + n_3)^2}}, \quad (*6.7)$$

$$\theta_{S3}^* = (z_{1-\frac{\alpha}{2}} - z_{1-\pi_{\min}}) \sqrt{\frac{2(n_1(\sigma_{I1}^2 + \sigma_{C1}^2) + n_2(\sigma_{C2}^2 + \sigma_{I2}^2) + n_3(\sigma_{I\phi}^2 + \sigma_{I\psi}^2))}{(n_1 + n_2 + n_3)^2}}, \quad (*6.9)$$

is equivalent to

$$\frac{n_1 + n_2 + n_3}{n_0} \sqrt{2n_0\sigma_{C0}^2 + \xi} > \frac{n_0 + n_1 + n_2 + n_3}{n_0} \sqrt{\xi} - \frac{\eta}{\sqrt{2}z}, \quad (*6.15)$$

where

$$\begin{aligned} \eta &= n_1(\mu_{C1} - \mu_{I1}) + n_2(\mu_{I2} - \mu_{C2}) + n_3(\mu_{I\psi} - \mu_{I\phi}), \\ \xi &= n_1(\sigma_{C1}^2 + \sigma_{I1}^2) + n_2(\sigma_{I2}^2 + \sigma_{C2}^2) + n_3(\sigma_{I\psi}^2 + \sigma_{I\phi}^2), \text{ and} \\ z &= z_{1-\frac{\alpha}{2}} - z_{1-\pi_{\min}}. \end{aligned} \quad (*6.14)$$

We start by writing $\Delta_{S3} - \Delta_{S2} > \theta_{S3}^* - \theta_{S2}^*$ in terms of η , ξ , and z as shown above:

$$\frac{\eta}{n_1 + n_2 + n_3} - \frac{\eta}{n_0 + n_1 + n_2 + n_3} > z \sqrt{\frac{2\xi}{(n_1 + n_2 + n_3)^2}} - z \sqrt{\frac{2(n_0(2\sigma_{C0}^2) + \xi)}{(n_0 + n_1 + n_2 + n_3)^2}}. \quad (\text{A.2.2})$$

Pulling out the common factors on each side we have

$$\eta \left(\frac{1}{n_1 + n_2 + n_3} - \frac{1}{n_0 + n_1 + n_2 + n_3} \right) > \sqrt{2}z \left(\frac{\sqrt{\xi}}{n_1 + n_2 + n_3} - \frac{\sqrt{2n_0\sigma_{C0}^2 + \xi}}{n_0 + n_1 + n_2 + n_3} \right). \quad (\text{A.2.2a})$$

Writing the partial fraction on the LHS of Inequality (A.2.2a) as a composite fraction we have

$$\eta \left(\frac{n_0}{(n_1 + n_2 + n_3)(n_0 + n_1 + n_2 + n_3)} \right) > \sqrt{2}z \left(\frac{\sqrt{\xi}}{n_1 + n_2 + n_3} - \frac{\sqrt{2n_0\sigma_{C0}^2 + \xi}}{n_0 + n_1 + n_2 + n_3} \right). \quad (\text{A.2.2b})$$

We then move the composite fraction to the RHS and the $\sqrt{2}z$ term to the LHS:

$$\frac{\eta}{\sqrt{2}z} > \frac{(n_1 + n_2 + n_3) \cdot (n_0 + n_1 + n_2 + n_3)}{n_0} \left(\frac{\sqrt{\xi}}{n_1 + n_2 + n_3} - \frac{\sqrt{2n_0\sigma_{C0}^2 + \xi}}{n_0 + n_1 + n_2 + n_3} \right), \quad (\text{A.2.2c})$$

and expand the brackets, canceling terms that appear on both the denominator and the numerator of the resultant fractions in the RHS:

$$\frac{\eta}{\sqrt{2}z} > \frac{n_0 + n_1 + n_2 + n_3}{n_0} \sqrt{\xi} - \frac{n_1 + n_2 + n_3}{n_0} \sqrt{2n_0\sigma_{C0}^2 + \xi}. \quad (\text{A.2.2d})$$

Finally, we swap the position of the leftmost term with that of the rightmost term to arrive at

$$\frac{n_1 + n_2 + n_3}{n_0} \sqrt{2n_0\sigma_{C0}^2 + \xi} > \frac{n_0 + n_1 + n_2 + n_3}{n_0} \sqrt{\xi} - \frac{\eta}{\sqrt{2z}},$$

which is identical to Inequality (6.15).

The trivial case: $\text{RHS} \leq 0$ We observe that the LHS of Inequality (6.15) is always positive, and hence the inequality trivially holds if the RHS is non-positive. Here we show $\text{RHS} \leq 0$ is equivalent to

$$\frac{n_0 + n_1 + n_2 + n_3}{n_0} \theta_{S3}^* \leq \Delta_{S3}, \quad (*6.16)$$

where

$$\Delta_{S3} = \frac{n_1(\mu_{C1} - \mu_{I1}) + n_2(\mu_{I2} - \mu_{C2}) + n_3(\mu_{I\psi} - \mu_{I\phi})}{n_1 + n_2 + n_3}, \quad (*6.8)$$

$$\theta_{S3}^* = (z_{1-\frac{\alpha}{2}} - z_{1-\pi_{\min}}) \sqrt{\frac{2(n_1(\sigma_{I1}^2 + \sigma_{C1}^2) + n_2(\sigma_{C2}^2 + \sigma_{I2}^2) + n_3(\sigma_{I\phi}^2 + \sigma_{I\psi}^2))}{(n_1 + n_2 + n_3)^2}}. \quad (*6.9)$$

This can be done by writing $\text{RHS} \leq 0$ in full:

$$\frac{n_0 + n_1 + n_2 + n_3}{n_0} \sqrt{\xi} - \frac{\eta}{\sqrt{2z}} \leq 0, \quad (\text{A.2.3})$$

and moving the second term on the LHS to the RHS:

$$\frac{n_0 + n_1 + n_2 + n_3}{n_0} \sqrt{\xi} \leq \frac{\eta}{\sqrt{2z}}. \quad (\text{A.2.3a})$$

We then add a factor of $\sqrt{2z}/(n_1 + n_2 + n_3)$ on both sides to get

$$\frac{n_0 + n_1 + n_2 + n_3}{n_0} \frac{\sqrt{\xi} \sqrt{2z}}{n_1 + n_2 + n_3} \leq \frac{\eta}{n_1 + n_2 + n_3}. \quad (\text{A.2.3b})$$

Noting from Equations (6.8) and (6.9) that

$$\Delta_{S3} = \frac{\eta}{n_1 + n_2 + n_3} \text{ and } \theta_{S3}^* = \frac{\sqrt{2} \cdot z \cdot \sqrt{\xi}}{n_1 + n_2 + n_3},$$

we finally replace the terms in Inequality (A.2.3b) with Δ_{S3} and θ_{S3}^* to arrive at

$$\frac{n_0 + n_1 + n_2 + n_3}{n_0} \theta_{S3}^* \leq \Delta_{S3},$$

which is identical to Inequality (6.16).

The non-trivial case: RHS > 0 We finally tackle the case where the RHS of the master inequality (Inequality (6.15)) is greater than zero. We show in this non-trivial case, the master inequality

$$\frac{n_1 + n_2 + n_3}{n_0} \sqrt{2n_0\sigma_{C0}^2 + \xi} > \frac{n_0 + n_1 + n_2 + n_3}{n_0} \sqrt{\xi} - \frac{\eta}{\sqrt{2}z} \quad (*6.15)$$

is equivalent to

$$\frac{2\sigma_{C0}^2}{n_0} > \frac{\left(\theta_{S3}^* - \Delta_{S3} + \frac{n_1+n_2+n_3}{n_0}\theta_{S3}^*\right)^2 - \left(\frac{n_1+n_2+n_3}{n_0}\theta_{S3}^*\right)^2}{2z^2}, \quad (*6.17)$$

where

$$\Delta_{S3} = \frac{n_1(\mu_{C1} - \mu_{I1}) + n_2(\mu_{I2} - \mu_{C2}) + n_3(\mu_{I\psi} - \mu_{I\phi})}{n_1 + n_2 + n_3}, \quad (*6.8)$$

$$\theta_{S3}^* = (z_{1-\frac{\alpha}{2}} - z_{1-\pi_{\min}}) \sqrt{\frac{2(n_1(\sigma_{I1}^2 + \sigma_{C1}^2) + n_2(\sigma_{C2}^2 + \sigma_{I2}^2) + n_3(\sigma_{I\phi}^2 + \sigma_{I\psi}^2))}{(n_1 + n_2 + n_3)^2}}, \quad (*6.9)$$

$$\eta = n_1(\mu_{C1} - \mu_{I1}) + n_2(\mu_{I2} - \mu_{C2}) + n_3(\mu_{I\psi} - \mu_{I\phi}),$$

$$\xi = n_1(\sigma_{C1}^2 + \sigma_{I1}^2) + n_2(\sigma_{I2}^2 + \sigma_{C2}^2) + n_3(\sigma_{I\psi}^2 + \sigma_{I\phi}^2), \text{ and}$$

$$z = z_{1-\frac{\alpha}{2}} - z_{1-\pi_{\min}}. \quad (*6.14)$$

We first multiply both sides of Inequality (6.15) with the fraction $\frac{n_0\sqrt{2}z}{n_1+n_2+n_3}$ to get

$$\frac{n_1 + n_2 + n_3}{n_0} \sqrt{2n_0\sigma_{C0}^2 + \xi} \frac{n_0\sqrt{2}z}{n_1 + n_2 + n_3} > \left(\frac{n_0 + n_1 + n_2 + n_3}{n_0} \sqrt{\xi} - \frac{\eta}{\sqrt{2}z} \right) \frac{n_0\sqrt{2}z}{n_1 + n_2 + n_3}. \quad (\text{A.2.4})$$

Cancelling terms on both sides of the fractions we have

$$\sqrt{2n_0\sigma_{C0}^2 + \xi}\sqrt{2}z > (n_0 + n_1 + n_2 + n_3)\frac{\sqrt{\xi}\sqrt{2}z}{n_1 + n_2 + n_3} - n_0\frac{\eta}{n_1 + n_2 + n_3}. \quad (\text{A.2.4a})$$

Again noting the identities for Δ_{S3} and θ_{S3}^* , i.e.,

$$\Delta_{S3} = \frac{\eta}{n_1 + n_2 + n_3} \text{ and } \theta_{S3}^* = \frac{\sqrt{2} \cdot z \cdot \sqrt{\xi}}{n_1 + n_2 + n_3},$$

we can replace the fractions on the RHS to obtain

$$\sqrt{2n_0\sigma_{C0}^2 + \xi}\sqrt{2}z > (n_0 + n_1 + n_2 + n_3)\theta_{S3}^* - n_0\Delta_{S3}. \quad (\text{A.2.4b})$$

We then square both sides of Inequality (A.2.4b) and move the $2z^2$ term to the RHS:

$$2n_0\sigma_{C0}^2 + \xi > \frac{((n_0 + n_1 + n_2 + n_3)\theta_{S3}^* - n_0\Delta_{S3})^2}{2z^2}. \quad (\text{A.2.4c})$$

Note the squaring still allows the implication to go both ways as both sides of Inequality (A.2.4c) are positive. Based on the identity for θ_{S3}^* , we observe ξ can also be written as

$$\xi = \frac{(n_1 + n_2 + n_3)^2(\theta_{S3}^*)^2}{2z^2}. \quad (\text{A.2.4d})$$

Thus, we can group all terms with a $2z^2$ denominator by moving ξ in Inequality (A.2.4c) to the RHS and substituting Equation (A.2.4d) into the resultant inequality:

$$2n_0\sigma_{C0}^2 > \frac{((n_0 + n_1 + n_2 + n_3)\theta_{S3}^* - n_0\Delta_{S3})^2 - ((n_1 + n_2 + n_3)\theta_{S3}^*)^2}{2z^2}. \quad (\text{A.2.4e})$$

We finally normalize the inequality to one with unit Δ_{S3} and θ_{S3}^* to enable effective comparison. We divide both sides of Inequality (A.2.4e) by n_0^2 :

$$\frac{2\sigma_{C0}^2}{n_0} > \frac{\left(\frac{n_0+n_1+n_2+n_3}{n_0}\theta_{S3}^* - \Delta_{S3}\right)^2 - \left(\frac{n_1+n_2+n_3}{n_0}\theta_{S3}^*\right)^2}{2z^2}, \quad (\text{A.2.4f})$$

and split the coefficient of θ_{S3}^* in the first squared term into an integer (1) and a fractional

$(\frac{n_1+n_2+n_3}{n_0})$ part to arrive at

$$\frac{2\sigma_{C0}^2}{n_0} > \frac{\left(\theta_{S3}^* - \Delta_{S3} + \frac{n_1+n_2+n_3}{n_0}\theta_{S3}^*\right)^2 - \left(\frac{n_1+n_2+n_3}{n_0}\theta_{S3}^*\right)^2}{2z^2},$$

which is identical to Inequality (6.17).

A.3 Calculations Related to Dual Control in Personalisation Strategy Experiment Designs

Similar to the appendix above, this appendix expands the exposition in Section 6.3.3 by providing the detailed algebraic manipulations of the inequalities involved. In the section, we determine, under the context of designing experiments for personalisation strategies, the sample size required for a dual control experiment setup (aka Setup 4) to emerge superior to a simpler A/B test setup (Setup 3 in this case).

We recap the mathematical arguments presented in the section. Under the evaluation framework introduced in Chapter 6, we can establish superiority of an experiment setup via two evaluation criteria. Setup 4 is superior to Setup 3 if at least one of the following inequalities hold:

$$\text{First criterion: } \Delta_{S4} > \Delta_{S3} \text{ and } \theta_{S4}^* < \theta_{S3}^*,$$

$$\text{Second criterion: } \Delta_{S4} - \Delta_{S3} > \theta_{S4}^* - \theta_{S3}^*,$$

where Δ_{S3} , θ_{S3}^* , Δ_{S4} , and θ_{S4}^* are that stated in Equations (6.8), (6.9), (6.10), and (6.11) respectively – we will restate the equations below where appropriate.

We showed in the main text (Section 6.3.3) that $\theta_{S4}^* > \theta_{S3}^*$ always holds, and hence Setup 4 will never be superior to Setup 3 under the first criterion. We thus focus on the second criterion $\Delta_{S4} - \Delta_{S3} > \theta_{S4}^* - \theta_{S3}^*$ and show

1. The criterion $\Delta_{S4} - \Delta_{S3} > \theta_{S4}^* - \theta_{S3}^*$ is equivalent to Inequality (6.20).

By applying various simplifying assumptions on the group response variances (i.e. the σ^2 -terms) and the group sample sizes (i.e. the n -terms), we also show:

2. Assuming only the σ^2 -terms are similar in magnitude, we can express the RHS of Inequality (6.20) as Expression (6.21);
3. Assuming only the n -terms are similar in magnitude, we can express the LHS of Inequality (6.20) as Expression (6.22); and

4. Assuming both the σ^2 -terms and the n -terms are similar in magnitude, Inequality (6.20) is approximately equivalent to Inequality (6.23).

We will restate all Inequalities and Expressions in full below where appropriate. Note the appendix only deals with the detailed algebraic manipulation. We refer readers back to Section 6.3.3 in the main text for the interpretations of and the insights from the resultant inequalities.

A.3.1 The master inequality

We first show the criterion $\Delta_{S4} - \Delta_{S3} > \theta_{S4}^* - \theta_{S3}^*$, where

$$\Delta_{S3} = \frac{n_1(\mu_{C1} - \mu_{I1}) + n_2(\mu_{I2} - \mu_{C2}) + n_3(\mu_{I\psi} - \mu_{I\phi})}{n_1 + n_2 + n_3}, \quad (*6.8)$$

$$\Delta_{S4} = \frac{n_2(\mu_{I2} - \mu_{C2}) + n_3(\mu_{I\psi} - \mu_{C3})}{n_2 + n_3} - \frac{n_1(\mu_{I1} - \mu_{C1}) + n_3(\mu_{I\phi} - \mu_{C3})}{n_1 + n_3}, \quad (*6.10)$$

$$\theta_{S3}^* = (z_{1-\frac{\alpha}{2}} - z_{1-\pi_{\min}}) \sqrt{\frac{2(n_1(\sigma_{I1}^2 + \sigma_{C1}^2) + n_2(\sigma_{C2}^2 + \sigma_{I2}^2) + n_3(\sigma_{I\phi}^2 + \sigma_{I\psi}^2))}{(n_1 + n_2 + n_3)^2}}, \quad (*6.9)$$

$$\theta_{S4}^* = 2 \cdot (z_{1-\frac{\alpha}{2}} - z_{1-\pi_{\min}}) \sqrt{\frac{n_1(\sigma_{C1}^2 + \sigma_{I1}^2) + n_3(\sigma_{C3}^2 + \sigma_{I\phi}^2)}{(n_1 + n_3)^2} + \frac{n_2(\sigma_{C2}^2 + \sigma_{I2}^2) + n_3(\sigma_{C3}^2 + \sigma_{I\psi}^2)}{(n_2 + n_3)^2}}, \quad (*6.11)$$

is equivalent to

$$\begin{aligned} & \frac{n_1 \frac{n_2(\mu_{I2} - \mu_{C2}) + n_3(\mu_{I\psi} - \mu_{C3})}{n_2 + n_3} - n_2 \frac{n_1(\mu_{I1} - \mu_{C1}) + n_3(\mu_{I\phi} - \mu_{C3})}{n_1 + n_3}}{\sqrt{n_1(\sigma_{C1}^2 + \sigma_{I1}^2) + n_2(\sigma_{C2}^2 + \sigma_{I2}^2) + n_3(\sigma_{I\phi}^2 + \sigma_{I\psi}^2)}} \\ & > \sqrt{2}z \left[\sqrt{2 \cdot \frac{\left(1 + \frac{n_2}{n_1 + n_3}\right)^2 [n_1(\sigma_{C1}^2 + \sigma_{I1}^2) + n_3(\sigma_{C3}^2 + \sigma_{I\phi}^2)] + \left(1 + \frac{n_1}{n_2 + n_3}\right)^2 [n_2(\sigma_{C2}^2 + \sigma_{I2}^2) + n_3(\sigma_{C3}^2 + \sigma_{I\psi}^2)]}{n_1(\sigma_{C1}^2 + \sigma_{I1}^2) + n_2(\sigma_{C2}^2 + \sigma_{I2}^2) + n_3(\sigma_{I\phi}^2 + \sigma_{I\psi}^2)}} - 1 \right], \quad (*6.20) \end{aligned}$$

where $z = z_{1-\frac{\alpha}{2}} - z_{1-\pi_{\min}}$. There are many terms involved in the inequality, and hence we first simplify the LHS and RHS independently, and combine them in the final step.

We start by writing the LHS (i.e., $\Delta_{S4} - \Delta_{S3}$) in full using Equations (6.8) and (6.10):

$$\frac{n_2(\mu_{I2} - \mu_{C2}) + n_3(\mu_{I\psi} - \mu_{C3})}{n_2 + n_3} - \frac{n_1(\mu_{I1} - \mu_{C1}) + n_3(\mu_{I\phi} - \mu_{C3})}{n_1 + n_3} - \frac{n_1(\mu_{C1} - \mu_{I1}) + n_2(\mu_{I2} - \mu_{C2}) + n_3(\mu_{I\psi} - \mu_{I\phi})}{n_1 + n_2 + n_3}. \quad (\text{A.3.1})$$

The expression can be rewritten in terms of multiplicative products between the n -terms and the (difference between) μ -terms:

$$n_1(\mu_{I1} - \mu_{C1}) \left[-\frac{1}{n_1 + n_3} + \frac{1}{n_1 + n_2 + n_3} \right] + n_2(\mu_{I2} - \mu_{C2}) \left[\frac{1}{n_2 + n_3} - \frac{1}{n_1 + n_2 + n_3} \right] + n_3\mu_{I\psi} \left[\frac{1}{n_2 + n_3} - \frac{1}{n_1 + n_2 + n_3} \right] + n_3\mu_{I\phi} \left[-\frac{1}{n_1 + n_3} + \frac{1}{n_1 + n_2 + n_3} \right] + n_3\mu_{C3} \left[-\frac{1}{n_2 + n_3} + \frac{1}{n_1 + n_3} \right]. \quad (\text{A.3.1a})$$

We then extract a $\frac{1}{n_1 + n_2 + n_3}$ term from Expression (A.3.1a):

$$\frac{1}{n_1 + n_2 + n_3} \left[n_1(\mu_{I1} - \mu_{C1}) \left(-\frac{n_1 + n_2 + n_3}{n_1 + n_3} + 1 \right) + n_2(\mu_{I2} - \mu_{C2}) \left(\frac{n_1 + n_2 + n_3}{n_2 + n_3} - 1 \right) + n_3\mu_{I\psi} \left(\frac{n_1 + n_2 + n_3}{n_2 + n_3} - 1 \right) + n_3\mu_{I\phi} \left(-\frac{n_1 + n_2 + n_3}{n_1 + n_3} + 1 \right) + n_3\mu_{C3} \left(-\frac{n_1 + n_2 + n_3}{n_2 + n_3} + \frac{n_1 + n_2 + n_3}{n_1 + n_3} \right) \right]. \quad (\text{A.3.1b})$$

This allows us to perform some cancellation with the RHS, which also has a $\frac{1}{n_1 + n_2 + n_3}$ term, in the final step. Noting

$$\frac{n_1 + n_2 + n_3}{n_1 + n_3} = 1 + \frac{n_2}{n_1 + n_3} \quad \text{and} \quad \frac{n_1 + n_2 + n_3}{n_2 + n_3} = 1 + \frac{n_1}{n_2 + n_3},$$

we can write the Expression (A.3.1b) as

$$\frac{1}{n_1 + n_2 + n_3} \left[n_1(\mu_{I1} - \mu_{C1}) \left(-\frac{n_2}{n_1 + n_3} \right) + n_2(\mu_{I2} - \mu_{C2}) \left(\frac{n_1}{n_2 + n_3} \right) + n_3\mu_{I\psi} \left(\frac{n_1}{n_2 + n_3} \right) + n_3\mu_{I\phi} \left(-\frac{n_2}{n_1 + n_3} \right) + n_3\mu_{C3} \left(-1 - \frac{n_1}{n_2 + n_3} + 1 + \frac{n_2}{n_1 + n_3} \right) \right], \quad (\text{A.3.1c})$$

and group the $\frac{n_1}{n_2 + n_3}$ and $\frac{n_2}{n_1 + n_3}$ terms to arrive at

$$\frac{1}{n_1 + n_2 + n_3} \left[\frac{n_1}{n_2 + n_3} (n_2(\mu_{I2} - \mu_{C2}) + n_3(\mu_{I\psi} - \mu_{C3})) - \frac{n_2}{n_1 + n_3} (n_1(\mu_{I1} - \mu_{C1}) + n_3(\mu_{I\phi} - \mu_{C3})) \right]. \quad (\text{A.3.1d})$$

We also write the RHS (i.e., $\theta_{S4}^* - \theta_{S3}^*$) in full using Equations (6.9) and (6.11):

$$2z \sqrt{\frac{n_1(\sigma_{C1}^2 + \sigma_{I1}^2) + n_3(\sigma_{C3}^2 + \sigma_{I\phi}^2)}{(n_1 + n_3)^2} + \frac{n_2(\sigma_{C2}^2 + \sigma_{I2}^2) + n_3(\sigma_{C3}^2 + \sigma_{I\psi}^2)}{(n_2 + n_3)^2}} - \sqrt{2}z \sqrt{\frac{n_1(\sigma_{I1}^2 + \sigma_{C1}^2) + n_2(\sigma_{C2}^2 + \sigma_{I2}^2) + n_3(\sigma_{I\phi}^2 + \sigma_{I\psi}^2)}{(n_1 + n_2 + n_3)^2}}, \quad (\text{A.3.2})$$

where $z = z_{1-\frac{\alpha}{2}} - z_{1-\pi_{\min}}$. We then extract a $\frac{\sqrt{2}z}{n_1+n_2+n_3}$ term from Expression (A.3.2) to obtain

$$\frac{\sqrt{2}z}{n_1 + n_2 + n_3} \left[\sqrt{2} \sqrt{\frac{\left(\frac{n_1+n_2+n_3}{n_1+n_3}\right)^2 [n_1(\sigma_{C1}^2 + \sigma_{I1}^2) + n_3(\sigma_{C3}^2 + \sigma_{I\phi}^2)] + \left(\frac{n_1+n_2+n_3}{n_2+n_3}\right)^2 [n_2(\sigma_{C2}^2 + \sigma_{I2}^2) + n_3(\sigma_{C3}^2 + \sigma_{I\psi}^2)]}{\sqrt{n_1(\sigma_{C1}^2 + \sigma_{I1}^2) + n_2(\sigma_{C2}^2 + \sigma_{I2}^2) + n_3(\sigma_{I\phi}^2 + \sigma_{I\psi}^2)}}} - \right] \quad (\text{A.3.2a})$$

where $\frac{n_1+n_2+n_3}{n_2+n_3}$ and $\frac{n_1+n_2+n_3}{n_1+n_3}$ can also be written as $1 + \frac{n_1}{n_2+n_3}$ and $1 + \frac{n_2}{n_1+n_3}$, respectively.

We finally combine both sides of the inequality by taking Expressions (A.3.1d) and (A.3.2a):

$$\begin{aligned} & \frac{1}{n_1 + n_2 + n_3} \left[\frac{n_1}{n_2 + n_3} (n_2(\mu_{I2} - \mu_{C2}) + n_3(\mu_{I\psi} - \mu_{C3})) - \frac{n_2}{n_1 + n_3} (n_1(\mu_{I1} - \mu_{C1}) + n_3(\mu_{I\phi} - \mu_{C3})) \right] \\ & > \frac{\sqrt{2}z}{n_1 + n_2 + n_3} \left[\sqrt{2} \sqrt{\frac{\left(1 + \frac{n_2}{n_1+n_3}\right)^2 [n_1(\sigma_{C1}^2 + \sigma_{I1}^2) + n_3(\sigma_{C3}^2 + \sigma_{I\phi}^2)] + \left(1 + \frac{n_1}{n_2+n_3}\right)^2 [n_2(\sigma_{C2}^2 + \sigma_{I2}^2) + n_3(\sigma_{C3}^2 + \sigma_{I\psi}^2)]}{\sqrt{n_1(\sigma_{C1}^2 + \sigma_{I1}^2) + n_2(\sigma_{C2}^2 + \sigma_{I2}^2) + n_3(\sigma_{I\phi}^2 + \sigma_{I\psi}^2)}}} - \right] \quad (\text{A.3.3}) \end{aligned}$$

Canceling the common $\frac{1}{n_1+n_2+n_3}$ terms on both sides, and dividing both sides by

$\sqrt{n_1(\sigma_{C1}^2 + \sigma_{I1}^2) + n_2(\sigma_{C2}^2 + \sigma_{I2}^2) + n_3(\sigma_{I\phi}^2 + \sigma_{I\psi}^2)}$ leads us to

$$\frac{n_1 \frac{n_2(\mu_{I2} - \mu_{C2}) + n_3(\mu_{I\psi} - \mu_{C3})}{n_2 + n_3} - n_2 \frac{n_1(\mu_{I1} - \mu_{C1}) + n_3(\mu_{I\phi} - \mu_{C3})}{n_1 + n_3}}{\sqrt{n_1(\sigma_{C1}^2 + \sigma_{I1}^2) + n_2(\sigma_{C2}^2 + \sigma_{I2}^2) + n_3(\sigma_{I\phi}^2 + \sigma_{I\psi}^2)}} > \sqrt{2}z \left[\sqrt{2 \cdot \frac{\left(1 + \frac{n_2}{n_1 + n_3}\right)^2 [n_1(\sigma_{C1}^2 + \sigma_{I1}^2) + n_3(\sigma_{C3}^2 + \sigma_{I\phi}^2)] + \left(1 + \frac{n_1}{n_2 + n_3}\right)^2 [n_2(\sigma_{C2}^2 + \sigma_{I2}^2) + n_3(\sigma_{C3}^2 + \sigma_{I\psi}^2)]}{n_1(\sigma_{C1}^2 + \sigma_{I1}^2) + n_2(\sigma_{C2}^2 + \sigma_{I2}^2) + n_3(\sigma_{I\phi}^2 + \sigma_{I\psi}^2)} - 1} \right],$$

which is identical to Inequality (6.20).

A.3.2 Simplifying RHS by assuming similar response variances

We now simplify the σ^2 -terms in Inequality (6.20) by assuming that they are similar in magnitude, i.e.,

$$\sigma_{C1}^2 \approx \sigma_{I1}^2 \approx \dots \approx \sigma_{I\psi}^2 \approx \sigma_G^2,$$

and show the RHS of the inequality, i.e.,

$$\sqrt{2}z \left[\sqrt{2 \cdot \frac{\left(1 + \frac{n_2}{n_1 + n_3}\right)^2 [n_1(\sigma_{C1}^2 + \sigma_{I1}^2) + n_3(\sigma_{C3}^2 + \sigma_{I\phi}^2)] + \left(1 + \frac{n_1}{n_2 + n_3}\right)^2 [n_2(\sigma_{C2}^2 + \sigma_{I2}^2) + n_3(\sigma_{C3}^2 + \sigma_{I\psi}^2)]}{n_1(\sigma_{C1}^2 + \sigma_{I1}^2) + n_2(\sigma_{C2}^2 + \sigma_{I2}^2) + n_3(\sigma_{I\phi}^2 + \sigma_{I\psi}^2)} - 1} \right],$$

is approximately equal to

$$\sqrt{2}z \left[\sqrt{\frac{n_1 + n_2 + n_3}{n_1 + n_3} + \frac{n_1 + n_2 + n_3}{n_2 + n_3} - 1} \right]. \quad (*6.21)$$

It is safe to apply the simplifying assumption as we know from the evaluation framework specification that there are three classes of parameters in the inequality: the user group sizes (n), the mean responses (μ), and the response variances (σ^2). Among these three classes of parameters, only the user group sizes have the potential to scale in any practical settings, and thus we can effectively treat the means and variances as constants below.

We begin by substituting σ_G^2 into Inequality (6.20):

$$\begin{aligned} & \frac{n_1 \frac{n_2(\mu_{I2}-\mu_{C2})+n_3(\mu_{I\psi}-\mu_{C3})}{n_2+n_3} - n_2 \frac{n_1(\mu_{I1}-\mu_{C1})+n_3(\mu_{I\phi}-\mu_{C3})}{n_1+n_3}}{\sqrt{n_1(2\sigma_G^2) + n_2(2\sigma_G^2) + n_3(2\sigma_G^2)}} \\ & > \sqrt{2}z \left[\sqrt{2 \cdot \frac{\left(1 + \frac{n_2}{n_1+n_3}\right)^2 [n_1(2\sigma_G^2) + n_3(2\sigma_G^2)] + \left(1 + \frac{n_1}{n_2+n_3}\right)^2 [n_2(2\sigma_G^2) + n_3(2\sigma_G^2)]}{n_1(2\sigma_G^2) + n_2(2\sigma_G^2) + n_3(2\sigma_G^2)}} - 1 \right]. \end{aligned} \quad (\text{A.3.4})$$

Moving the common $2\sigma_G^2$ terms out and canceling the common terms in the RHS fraction we have

$$\begin{aligned} & \frac{n_1 \frac{n_2(\mu_{I2}-\mu_{C2})+n_3(\mu_{I\psi}-\mu_{C3})}{n_2+n_3} - n_2 \frac{n_1(\mu_{I1}-\mu_{C1})+n_3(\mu_{I\phi}-\mu_{C3})}{n_1+n_3}}{\sqrt{2\sigma_G^2(n_1 + n_2 + n_3)}} > \\ & \sqrt{2}z \left[\sqrt{2 \cdot \frac{\left(1 + \frac{n_2}{n_1+n_3}\right)^2 (n_1 + n_3) + \left(1 + \frac{n_1}{n_2+n_3}\right)^2 (n_2 + n_3)}{n_1 + n_2 + n_3}} - 1 \right]. \end{aligned} \quad (\text{A.3.4a})$$

We can already see the LHS of Inequality (A.3.4a) scales along $O(\sqrt{n})$ – we will demonstrate this result in greater detail in the next subsection.

Focusing on the RHS of the inequality, we express the squared terms as rational fractions and divide each term in the numerator by the denominator to obtain

$$\sqrt{2}z \left[\sqrt{2 \left[\left(\frac{n_1 + n_2 + n_3}{n_1 + n_3} \right)^2 \frac{n_1 + n_3}{n_1 + n_2 + n_3} + \left(\frac{n_1 + n_2 + n_3}{n_2 + n_3} \right)^2 \frac{n_2 + n_3}{n_1 + n_2 + n_3} \right]} - 1 \right]. \quad (\text{A.3.5})$$

Canceling the common $n_1 + n_2 + n_3$ terms leads to that presented in Expression (6.21):

$$\sqrt{2}z \left[\sqrt{\frac{n_1 + n_2 + n_3}{n_1 + n_3} + \frac{n_1 + n_2 + n_3}{n_2 + n_3}} - 1 \right].$$

A.3.3 Simplifying LHS by assuming similar sample sizes

We demonstrate the scaling relation between the LHS of Inequality (6.20) and the number of users in each group by simplifying the n -terms (assuming $n_1 \approx n_2 \approx n_3 \approx n$) and showing

the LHS of the inequality, i.e.,

$$\frac{n_1 \frac{n_2(\mu_{I2} - \mu_{C2}) + n_3(\mu_{I\psi} - \mu_{C3})}{n_2 + n_3} - n_2 \frac{n_1(\mu_{I1} - \mu_{C1}) + n_3(\mu_{I\phi} - \mu_{C3})}{n_1 + n_3}}{\sqrt{n_1(\sigma_{C1}^2 + \sigma_{I1}^2) + n_2(\sigma_{C2}^2 + \sigma_{I2}^2) + n_3(\sigma_{I\phi}^2 + \sigma_{I\psi}^2)}}$$

is approximately equal to

$$\frac{\sqrt{n}((\mu_{I2} - \mu_{C2}) - (\mu_{I1} - \mu_{C1}) + \mu_{I\psi} - \mu_{I\phi})}{2\sqrt{\sigma_{C1}^2 + \sigma_{I1}^2 + \sigma_{C2}^2 + \sigma_{I2}^2 + \sigma_{I\phi}^2 + \sigma_{I\psi}^2}}. \quad (*6.22)$$

Note that unlike the previous subsection, we do not make any simplifying assumptions on the σ^2 -terms here. While the relationship (that the LHS of the inequality scales along $O(\sqrt{n})$) is evident by inspecting Inequality (6.20) (or its simplified version as Inequality (A.3.4a)) itself, we believe the simplification allows us to show the relationship more clearly.

We begin by substituting n into Inequality (6.20) to obtain

$$\begin{aligned} & \frac{n \frac{n(\mu_{I2} - \mu_{C2}) + n(\mu_{I\psi} - \mu_{C3})}{n+n} - n \frac{n(\mu_{I1} - \mu_{C1}) + n(\mu_{I\phi} - \mu_{C3})}{n+n}}{\sqrt{n(\sigma_{C1}^2 + \sigma_{I1}^2) + n(\sigma_{C2}^2 + \sigma_{I2}^2) + n(\sigma_{I\phi}^2 + \sigma_{I\psi}^2)}} \\ & > \sqrt{2}z \left[\sqrt{2 \cdot \frac{\left(1 + \frac{n}{n+n}\right)^2 [n(\sigma_{C1}^2 + \sigma_{I1}^2) + n(\sigma_{C3}^2 + \sigma_{I\phi}^2)] + \left(1 + \frac{n}{n+n}\right)^2 [n(\sigma_{C2}^2 + \sigma_{I2}^2) + n(\sigma_{C3}^2 + \sigma_{I\psi}^2)]}{n(\sigma_{C1}^2 + \sigma_{I1}^2) + n(\sigma_{C2}^2 + \sigma_{I2}^2) + n(\sigma_{I\phi}^2 + \sigma_{I\psi}^2)}} - 1 \right]. \end{aligned} \quad (A.3.6)$$

Moving the common n -terms out and canceling them in the fractions where appropriate lead to

$$\begin{aligned} & \frac{\sqrt{n} \frac{1}{2} [((\mu_{I2} - \mu_{C2}) + (\mu_{I\psi} - \mu_{C3})) - ((\mu_{I1} - \mu_{C1}) + (\mu_{I\phi} - \mu_{C3}))]}{\sqrt{\sigma_{C1}^2 + \sigma_{I1}^2 + \sigma_{C2}^2 + \sigma_{I2}^2 + \sigma_{I\phi}^2 + \sigma_{I\psi}^2}} \\ & > \sqrt{2}z \left[\sqrt{2 \cdot \frac{\left(1 + \frac{1}{2}\right)^2 (\sigma_{C1}^2 + \sigma_{I1}^2 + \sigma_{C3}^2 + \sigma_{I\phi}^2) + \left(1 + \frac{1}{2}\right)^2 (\sigma_{C2}^2 + \sigma_{I2}^2 + \sigma_{C3}^2 + \sigma_{I\psi}^2)}{\sigma_{C1}^2 + \sigma_{I1}^2 + \sigma_{C2}^2 + \sigma_{I2}^2 + \sigma_{I\phi}^2 + \sigma_{I\psi}^2}} - 1 \right], \end{aligned} \quad (A.3.6a)$$

where the LHS is identical to Expression (6.22) after canceling the μ_{C3} terms in the numerator and expressing the half in the numerator as a two in the denominator.

It is clear that there are no n -terms left on the RHS of Inequality (A.3.6a), and hence the RHS remains constant as shown in the previous subsection. To set up the inequality to demonstrate the last result – that the number of users required for a dual control setup to emerge superior is large – we further simplify the RHS of the inequality by rearranging the terms in the square root:

$$\begin{aligned} & \frac{\sqrt{n} [((\mu_{I2} - \mu_{C2}) + (\mu_{I\psi} - \mu_{C3})) - ((\mu_{I1} - \mu_{C1}) + (\mu_{I\phi} - \mu_{C3}))]}{2\sqrt{\sigma_{C1}^2 + \sigma_{I1}^2 + \sigma_{C2}^2 + \sigma_{I2}^2 + \sigma_{I\phi}^2 + \sigma_{I\psi}^2}} \\ & > \sqrt{2}z \left[\sqrt{2 \left(\frac{3}{2}\right)^2 \left(1 + \frac{2\sigma_{C3}^2}{\sigma_{C1}^2 + \sigma_{I1}^2 + \sigma_{C2}^2 + \sigma_{I2}^2 + \sigma_{I\phi}^2 + \sigma_{I\psi}^2}\right)} - 1 \right]. \end{aligned} \quad (\text{A.3.6b})$$

A.3.4 Required number of users for dual control to emerge superior

We finally show that while Setup 4 could emerge superior to Setup 3 as the number of users increase, the number of users required is large. We do so by assuming both the σ^2 - and n -terms are similar in magnitude, i.e. $\sigma_{C1}^2 \approx \sigma_{I1}^2 \approx \dots \approx \sigma_{I\psi}^2 \approx \sigma_G^2$ and $n_1 \approx n_2 \approx n_3 \approx n$, and show that Inequality (6.20) is approximately equivalent to

$$n > \left(2\sqrt{12} \left(\sqrt{6} - 1\right) z\right)^2 \frac{\sigma_G^2}{\Delta^2}, \quad (*6.23)$$

where $\Delta = (\mu_{I2} - \mu_{C2}) - (\mu_{I1} - \mu_{C1}) + \mu_{I\psi} - \mu_{I\phi}$ is the actual effect size difference between Setups 4 and 3. Note we are determining when Setup 4 is superior to Setup 3 under the second evaluation criterion – that *the gain in actual effect* is greater than the loss in sensitivity – and thus assume Δ is positive.

The approximate equivalence can be shown by substituting σ_G^2 into Inequality (A.3.6b), which

is Inequality (6.20) with the additional assumption that n -terms are similar in magnitude:⁴

$$\frac{\sqrt{n} [((\mu_{I2} - \mu_{C2}) + (\mu_{I\psi} - \mu_{C3})) - ((\mu_{I1} - \mu_{C1}) + (\mu_{I\phi} - \mu_{C3}))]}{2\sqrt{6\sigma_G^2}} > \sqrt{2}z \left[\sqrt{2\left(\frac{3}{2}\right)^2 \left(1 + \frac{2\sigma_G^2}{6\sigma_G^2}\right)} - 1 \right]. \quad (\text{A.3.7})$$

Noting the expression within the LHS square bracket is equal to Δ , we simplify the expression within the RHS square root, and move every non- n term to the RHS of the inequality to obtain

$$\sqrt{n} > \sqrt{2}z \left[\sqrt{6} - 1 \right] \frac{2\sqrt{6\sigma_G^2}}{\Delta}. \quad (\text{A.3.7a})$$

As all quantities in the inequality are positive, we can square both sides and consolidate the coefficients on the RHS to arrive at

$$n > \left(2\sqrt{12} \left(\sqrt{6} - 1 \right) z \right)^2 \frac{\sigma_G^2}{\Delta^2},$$

which is identical to Inequality (6.23).

⁴ Alternatively, we can substitute n into Inequality (A.3.4a), which is Inequality (6.20) with the additional assumption that the σ^2 -terms are similar in magnitude. Simplifying the resultant inequality would yield the same result.