ABSTRACT

YU, MIAO. Online Testing and Semiparametric Estimation of Complex Treatment Effects. (Under the direction of Wenbin Lu.)

A/B testing plays a critical role in the high-tech industry to guide product development and accelerate innovation. It performs formal null hypothesis statistical testing to determine if there is a significant difference between the two variants on the metric of interest. However, a typical A/B test presents two problems: (i) a fixed-horizon framework inflates the type I error under continuous monitoring; (ii) the homogeneous effects assumption fails to detect the heterogeneous treatment effects among the population. The first two parts of this thesis aim to develop an online test that can address these two problems simultaneously.

In the first part, we propose an online test, named sequential score test (SST), to detect the multi-dimensional heterogeneous treatment effect under a generalized linear model, which is able to control the type I error under continuous monitoring. We provide an online p-value calculation for SST, making it convenient for continuous monitoring, and extend our tests to online multiple testing settings by controlling the false discovery rate. We examine the empirical performance of the proposed tests and compare them with a state-of-art online test, named mSPRT using simulations and a real dataset. The results show that our proposed test controls type I error at any time, has higher detection power, and allows quick inference on online A/B testing.

In the second part, we propose an online test for subgroup treatment effects based on value difference, named SUBTLE. The proposed test allows the experimenters to "peek" the results during the experiment without harming the statistical guarantees. It considers a nonparametric model and aims to test if some subgroup of the population will benefit from the investigative treatment. If the testing result indicates the existence of such a subgroup, a subgroup will be identified using a readily available estimated optimal treatment rule. The empirical performance of our proposed test is examined on simulations and a real dataset. The results show that the SUBTLE has high detection power with controlled type I error at any time, is robust to noise covariates, and achieves early stopping compared with the corresponding fixed-horizon test.

In the last part, we shift the attention to the estimation of time-dependent treatment effects with zero-inflated outcomes. Motivated by freemium mobile game data, we propose a class of multiplicative structural nested mean models for zero-inflated nonnegative outcomes, which flexibly describes the joint effect of a sequence of treatments in the presence of time-varying confounders. The proposed estimator solves a doubly robust estimating equation, where the nuisance functions, propensity score and conditional outcome means given confounders, are estimated parametrically or nonparametrically. To improve the accuracy, we leverage the characteristic of zero-inflated outcomes by estimating the conditional means in two parts, that is, separately modeling the probability of having positive outcomes given confounders and the mean outcome conditional on its being positive and confounders. We show that the proposed estimator is consistent and asymptotically normal as either the sample size or the follow-up

time goes to infinity. Moreover, the typical sandwich formula can be used to estimate the variance of treatment effect estimators consistently, without accounting for the variation due to estimating nuisance functions. Simulation studies and an application to a freemium mobile game dataset are presented to demonstrate the empirical performance of the proposed method and support our theoretical findings.

Online Testing and Semiparametric Estimation of Complex Treatment Effects

by
Miao Yu

A dissertation submitted to the Graduate Faculty of
North Carolina State University
in partial fulfillment of the
requirements for the Degree of
Doctor of Philosophy

Statistics

Raleigh, North Carolina

2021

APPROVED BY:

Shu Yang	Rui Song
Minh Tang	Wenbin Lu Chair of Advisory Committee

DEDICATION

To my parents, Hairong and Yanhe.

BIOGRAPHY

Miao Yu was born in Weifang, China. In 2013, she graduated from Weifang No.1 High School and moved to Jinan to pursue her undergraduate studies at Shandong University. She graduated from Shandong University with a degree of Bachelor of Science in Mathematics in 2017. In the same year, Miao entered the Statistics Ph.D Program at North Carolina State University. There, she worked with Dr. Wenbin Lu on research topics including causal inference and sequential testing.

ACKNOWLEDGEMENTS

I would like to thank my advisor, Dr. Wenbin Lu, for his patient guidance, professional expertise, and consistent support through my Ph.D. study. We have been worked together for three years and meet every week. I am grateful for all his feedback and suggestions regarding my research project. Without his help, it is impossible to complete this dissertation. I would also like to express my graduate to Dr. Shu Yang for collaborating with me on the third project, and Dr. Rui Song for cooperation on the first and second projects. Their great ideas, statistical expertise, and helpful suggestions are essential parts of my research projects. I am also thankful to my committee member, Dr. Minh Tang, for offering valuable feedback. A special thanks go to Dr. Wenbin Lu and Dr. Rui Song for holding group meetings every week which broaden my knowledge and improve my statistical thinking through communicating with many talented students on various research topics.

I would like to thank all the professors and staff from the Statistics Department at NC State University for giving various courses, organizing department activities, and offering technical support, especially Dr. Jacqueline Hughes-Oliver, Dr. Ryan Martin, Dr. Arnab Maity, Dr. Anastasios Tsiatis, Dr. Eric Chi, Dr. Subhashis Ghoshal, Dr. Ana-Maria Staicu, Dr. Sujit Ghosh, Dr. Wenbin Lu, Dr. Donald Martin, Lanakila, Alison, and Terry. The courses I have learned there help me lay a solid foundation for my academic research and future job in the industry.

I would like to thank all my colleagues from IoT Division at SAS Institute, especially Zohreh Asgharzadeh, Arin Chaudhuri, Gul Ege, Kai Shen, Wenhao Hu, Hansi Jiang, Haoyu Wang, Qi Ma, and Ye Mao. The job at SAS Institute is my first internship in the U.S. Their kind guidance, inclusiveness, and encouragement help me quickly get adapted to the job and deliver quality results.

Finally, I would like to express my deepest love to my parents for their unconditional support and to my boyfriend, Zhen Li, for his wonderful company.

TABLE OF CONTENTS

LIST OF TABLES . . . . . **vii**

LIST OF FIGURES . . . . . **viii**

Chapter 1 INTRODUCTION **1**
1.1 Online Testing of Heterogeneous Treatment Effects 1
1.1.1 Background 1
1.1.2 Review of sequential testing 3
1.1.3 Application to online A/B testing 6
1.2 Semiparametric Estimation of Time-Dependent Treatment Effects 9
1.2.1 Background 9
1.2.2 Review of structural nested mean models 10
1.3 Outline 13

Chapter 2 A NEW FRAMEWORK FOR ONLINE TESTING OF HETEROGENEOUS TREATMENT EFFECTS **14**
2.1 Introduction 14
2.2 Preliminaries 16
2.2.1 Fixed-horizon testing 16
2.2.2 Sequential testing 17
2.2.3 Heterogeneous treatment effects 18
2.3 A New Framework of Sequential Testing 19
2.3.1 Sequential score test 19
2.3.2 Validity of sequential score test 21
2.4 Multiple Testing 22
2.5 Experiments 23
2.5.1 Simulation study 23
2.5.2 Real data application 26
2.6 Conclusions 27

Chapter 3 ONLINE TESTING OF SUBGROUP TREATMENT EFFECTS BASED ON VALUE DIFFERENCE **28**
3.1 Introduction 28
3.2 Related Work 30
3.2.1 Mixture sequential probability ratio test 30
3.2.2 Sequential score test 31
3.3 Subgroup Treatment Effects Test Based on Value Difference 31
3.3.1 Problem setup 31
3.3.2 Algorithm and implementation 32
3.3.3 Validity 35
3.4 Simulated Experiments 36
3.4.1 Type I error & power 36
3.4.2 Noise covariates 38
3.4.3 Stopping time 38
3.5 Real Data Experiments 40
3.6 Conclusion 41

Chapter 4 MULTIPLICATIVE STRUCTURAL NESTED MEAN MODEL FOR ZERO-INFLATED OUTCOMES **42**
4.1 Introduction 42
4.2 Multiplicative Structural Nested Mean Model 44
4.3 Main Methodology 46
4.3.1 Identification and estimation 46
4.3.2 Asymptotic distribution and variance estimation 48
4.4 Simulation Study 53
4.5 Application to Freemium Mobile Game Data 55
4.6 Concluding Remarks 58

BIBLIOGRAPHY **61**

APPENDICES **67**
Appendix A SUPPLEMENT TO CHAPTER 2 68
A.1 Proof of Main Results 68
A.1.1 Proof of Lemma 2.3.1 68
A.1.2 Proof of Remark 2.3.1 70
A.2 Additional Simulation Results 72
Appendix B SUPPLEMENT TO CHAPTER 3 75
B.1 Proof of Main Results 75
B.1.1 Proof of Theorem 3.3.1 75
B.1.2 Proof of Proposition 3.3.1 76
B.2 Additional Simulation Results 77
Appendix C SUPPLEMENT TO CHAPTER 4 79
C.1 Proof of Main Results 79
C.1.1 Proof of Theorem 4.3.1 79
C.1.2 Proof of Theorem 4.3.2 84
C.2 Proof of Auxiliary Results 84
C.2.1 Proof of Lemma C.1.1 84
C.2.2 Proof of Lemma C.1.2 86
C.3 Additional Simulation Results 86

LIST OF TABLES

Table 2.1 Estimated type I error of SST and mSPRT in A/B testing on data with heterogeneous treatment effects . . . 24
Table 2.2 Estimated power of SST and mSPRT in A/B testing on data with heterogeneous treatment effects . . . 25
Table 2.3 Estimated false discovery rate and true positive rate of SST and mSPRT in multiple testing on data with heterogeneous treatment effects . . . 26
Table 2.4 Estimated heterogeneous treatment effect and average treatment effect in the Yahoo dataset . . . 27

Table 3.1 The covariates, $\mu(\mathbf{X})$, and $\theta(\mathbf{X})$ in Models I-IV . . . 37
Table 3.2 Estimated type I error and power of SUBTLE and SST with batch size 20 . . . 38
Table 3.3 Estimated type I error and power of SUBTLE and SST with batch size 40 . . . 38

Table 4.1 Simulation results with same directional treatment effect and $T = 60$ under individual randomization . . . 55
Table 4.2 Simulation results with opposite directional treatment effect and $T = 60$ under individual randomization . . . 56
Table 4.3 Simulation results with same directional treatment effect and $n = 400$ under cluster randomization . . . 57
Table 4.4 Simulation results with opposite directional treatment effect and $n = 400$ under cluster randomization . . . 57

Table A.1 Estimated type I error of SST and mSPRT in A/B testing on data with homogeneous treatment effects . . . 72
Table A.2 Estimated power of SST and mSPRT in A/B testing on data with homogeneous treatment effects . . . 73
Table A.3 Estimated false discovery rate and true positive rate of SST and mSPRT in multiple testing on data with homogeneous treatment effects . . . 74

Table B.1 Estimated type I error and power of SUBTLE with varying mixture density variance 78

Table C.1 Simulation results with same directional treatment effect and $n = 100$ under individual randomization . . . 87
Table C.2 Simulation results with opposite directional treatment effect and $n = 100$ under individual randomization . . . 88

LIST OF FIGURES

Figure 3.1 Estimated power of SUBTLE and SST with data generated from Model I as the number of noise covariates increases 39

Figure 3.2 Histogram of stopping time of SUBTLE and SST with data generated from Model V 39

Figure 3.3 An optimal treatment rule in the Yahoo dataset by a decision tree 40

Figure 4.1 Illustration of challenges due to time-varying confounding: $L_{0,t}$ is a confounder for $A_{0,t}$ and $Y_{0,t}$, which should be controlled for; however, because $L_{0,t}$ is an intermediate variable from $A_{0,t-1}$ to $Y_{0,t}$, if we do control for $L_{0,t}$, then the causal effect of $A_{0,t-1}$ to $Y_{0,t}$ will be blocked. 45

Figure 4.2 Results of the multiplicative structural nested mean model in Freemium Mobile Game Data 59

CHAPTER

1

INTRODUCTION

1.1 Online Testing of Heterogeneous Treatment Effects

1.1.1 Background

Sequential testing is a large class of hypothesis testing, which has been widely used in clinical trials [Arm60; GL83] since the last century. In contrast to the classic fixed-horizon framework where significance is only checked after all samples have been collected, sequential testing continuously checks for significance at every new sample and stops the test as soon as a significant result is detected. This field was first introduced by Wald [Wal45], who proposed the *sequential probability ratio test* (SPRT) to test a simple null hypothesis versus a simple alternative hypothesis with the likelihood ratio as the test statistics, and then was extended to composite hypotheses by much following literature [Sch62; Arm69; Cox63; Rob70; Lai88]. See Lai [Lai01] for a thorough survey on the development of sequential testing. Recently, sequential testing starts to shed light on online *A/B testing*, due to its ability to quickly detect the difference while keeping the false positive rate under control.

Randomized controlled experiment, nowadays has played a critical role in the web-facing industry to guide product development and accelerate innovation [Koh09]. It tests user experience by randomly exposing users to different *variants*, i.e. versions of a product. The difference between variants ranges from user-interface elements, e.g. the color of a caption, to backend algorithms, e.g. the ranking algorithms of a search engine. A/B testing refers to those experiments where only two variants are involved: *control* (A), the existing variant, and *treatment* (B), a new variant being evaluated. A metric of interest often called the *overall evaluation criterion* (OEC) in A/B testing, is chosen as the outcome that is predictive of long-term goals but measurable in the short-term, such as the conversion rate, revenue, unit purchases, or

some weighted combinations. Then, a null hypothesis statistical test is performed to evaluate whether there is a statistically significant difference between the two variants on metrics of interest. This scientific design helps to control for the external variations and thus establish the causality between the variants and the outcome.

Even though the theory is quite simple, the implementation and deployment of A/B testing in practice require a lot of effort. Many papers from the industry shared their lessons [Tan10; Koh11; Koh12; Koh13]. One of the common problems in A/B testing is the inflated type I error due to "peeking" [Goo14]. Most A/B tests employ a fixed-horizon framework, whose validity requires that the sample size is fixed and should be determined before the experiment starts. However, the shareholders in practice want to "peek" the experiment and find the significance as quickly as possible to avoid large (i) time cost: an A/B test may take several weeks or months to collect the determined size of samples; and (ii) opportunity cost: if the new version is actually damaging profit, the users who have been assigned to it will be stuck in a bad experience for a long time [Ju19]. In addition, the sample size calculation depends on the *minimum detectable effect* (MDE). A smaller MDE will result in a larger sample size, which requires a long waiting time. Most users lack good prior knowledge of the trade-off between detection ability and waiting time, and may want to adjust them after peeking early at results. The behaviors of continuously monitoring the p-value and concluding the experiment prematurely will favorably bias getting significant results and lead to very high false positive probabilities, well in excess of the nominal significance level α [Sim11]. As an extreme example in Pekelis et al. [Pek15], stopping the first time that the p-value is less than α actually has a type I error of 1.

Sequential testing, on the other hand, allows dynamic monitoring of the experiment while still having the type I error in control. It generally gives a significant decrease in the required sample size compared to fixed-horizon testing, especially when the true effect is large, and thus can terminate a bad experiment much quickly. These properties make sequential testing become a desirable tool in A/B testing. Starting from 2014, sequential testing was brought to A/B testing [Øy14; Joh15; Mil15]. A notion of *always valid p-value* was proposed to facilitate the usage of sequential testing [Joh17]. It also serves as a tool for converting fixed-horizon multiple testing procedures to a sequential version. Later, Malek et al. [Mal17] showed that if the original multiple testing procedure has a type I error guarantee in a certain family (including false discovery rate and family-wise error rate), then the sequential conversion inherits an analogous guarantee, which laid a foundation for *Bonferroni correction* [Mil66] and *Benjamini-Hochberg* [Ben95] procedure working in a sequential setting.

However, A/B testing mainly focuses on homogeneous treatment effects while it is widely known that the treatment effects are varying across subpopulations. Testing the heterogeneous treatment effects will distinguish subpopulations that may benefit from a particular treatment from those who may not, and thereby will help companies make profits by promoting new products to those target subpopulations. Nevertheless, most works on heterogeneous treatment effects are investigating in estimation and summarization [Den16; Gri17] rather than testing. It motivates us to develop an online testing method for heterogeneous treatment effects. In the next two sections, we will give a review of sequential testing and its application to online A/B testing. In Chapters 2 and 3, we will propose two sequential tests for

heterogeneous treatment effects.

1.1.2 Review of sequential testing

Sequential testing is a statistical test procedure where at any stage of the experiment, one needs to make a decision from (i) rejecting the null hypothesis, (ii) accepting the null hypothesis, or (iii) continuing the sampling. An essential property of it is that the required sample size is not predetermined, but a random variable depending on the data.

The sequential probability ratio test (SPRT) [Wal45] is the first sequential test and lays out the foundation for the development of sequential testing. Let $X_1, X_2, \cdots$ be i.i.d random variables with a probability density function f. The SPRT aims to test a simple null hypothesis against a simple alternative hypothesis:

$$\mathrm{H}_0 : f = f_0 \quad vs. \quad \mathrm{H}_1 : f = f_1. \tag{1.1}$$

Let $R_n = \prod_{i=1}^{n} \{f_1(X_i)/f_0(X_i)\}$ be the likelihood ratio.The SPRT stops the experiment at the sample size

$$N = \inf\{n \geq 1 : R_n \geq C \text{ or } R_n \leq D\},$$

where $C > 1 > D > 0$ depending on the predetermined type I error control $\alpha = P_{H_0}(R_N \geq C)$ and type II error control $\beta = P_{H_1}(R_N \leq D)$. When stopping occurs, we reject H_0 or H_1 according to $R_N \geq C$ or $R_N \leq D$. Wald [Wal45] proved that C and D are approximately equal to $(1-\beta)/\alpha$ and $\beta/(1-\alpha)$. It is shown that among all the tests on hypotheses (1.1) whose error probabilities satisfy

$$P_{H_0}(\text{Reject } \mathrm{H}_0) \leq \alpha \text{ and } P_{H_1}(\text{Reject } \mathrm{H}_1) \leq \beta,$$

the SPRT minimizes the expected sample size N under both null H_0 and alternative H_1 hypotheses [Wal48]. Several sequential tests with corresponding restrictions and properties were developed, extending the SPRT to composite hypothesis testing, including the *mixture sequential probability ratio test* [Rob70], *sequential generalized likelihood ratio test* [Lai88; Lai94], and *repeated significance test* [Arm69].

The mixture sequential probability ratio test (mSPRT) [Rob70] supposes that the i.i.d random variables $X_1, X_2, \cdots$ have a probability density function $f_\theta(x)$ induced by parameter θ, and test

$$\mathrm{H}_0 : \theta = \theta_0 \quad vs. \quad \mathrm{H}_1 : \theta \neq \theta_0.$$

Its test statistics Λ_n^π is a mixture of likelihood ratios as below:

$$\Lambda_n^\pi = \int_\Theta \left\{ \prod_{i=1}^{n} \frac{f_\theta(X_i)}{f_{\theta_0}(X_i)} \right\} \pi(\theta) d\theta, \tag{1.2}$$

with a mixture density $\pi(\cdot)$ over the parameter space Θ. The mSPRT stops the sampling at stage

$$N = \inf\{n \geq 1 : \Lambda_n^\pi \geq \alpha^{-1}\}$$

and rejects the null hypothesis H_0 in favor of H_1. If no such time exists, they continue the sampling indefinitely and accept H_0. The likelihood ratio under H_0 is a nonnegative martingale with an initial value equal to one, and so is the mixture of such likelihood ratios. It is easy to show that the type I error is controlled at α by an application of *Markov's inequality* and *optional stopping theorem* [Gri01]: $P_{H_0}(\Lambda_n^\pi \geq \alpha^{-1}) \leq \mathbb{E}_{H_0}(\Lambda_n^\pi)/\alpha^{-1} = \mathbb{E}_{H_0}(\Lambda_0^\pi)/\alpha^{-1} = \alpha$. The choice of mixture distribution $\pi(\theta)$ does not impact the validity but may make a difference to the expected sample size [Joh15]. It is shown that mSPRT is almost optimal for data from an exponential family of distributions, with respect to certain criteria of optimality [Pol78]. Besides, mSPRT is a test of power one [Rob74], which means that any small deviation from θ_0 can be detected as long as waiting long enough.

The sequential generalized likelihood ratio test refers to a class of tests that use the generalized likelihood ratio as the test statistic. It is derived to test the following hypotheses

$$H_0 : \theta \leq \theta_0 \quad vs. \quad H_1 : \theta \geq \theta_1 (> \theta_0). \tag{1.3}$$

This pair of hypotheses can also be tested by Wald's SPRT in the same way as it tests $H_0 : \theta = \theta_0$ vs. $H_1 : \theta = \theta_1$. The type I and type II errors are still controlled with $\mathbb{E}_{\theta_0}(T)$ and $\mathbb{E}_{\theta_1}(T)$ minimized as before. However, the expected sample size at some $\theta \in \Theta$ of SPRT can be extremely large. Lorden [Lor76] proposed a 2-SPRT that can minimize $\mathbb{E}_{\theta^*}(T)$ for some $\theta^* \in \Theta$ with stopping sample size

$$N = \inf\left\{ n \geq 1 : \prod_{i=1}^{n} \frac{f_{\theta^*}(X_i)}{f_{\theta_0}(X_i)} \geq C_0 \text{ or } \prod_{i=1}^{n} \frac{f_{\theta^*}(X_i)}{f_{\theta_1}(X_i)} \geq C_1 \right\}.$$

When the sampling stops, it rejects H_0 or H_1 according to $\prod_{i=1}^{N} \left\{ f_{\theta^*}(X_i)/f_{\theta_0}(X_i) \right\} \geq C_0$ or $\prod_{i=1}^{N} \left\{ f_{\theta^*}(X_i)/f_{\theta_1}(X_i) \right\} \geq C_1$. A desirable choice of θ^* is the θ that maximizes $\mathbb{E}_\theta(T)$. Weiss [Wei62] showed that in the case of testing the mean from a normal distribution where the error probabilities α, β at values θ_0, θ_1 are equal, the problem can be reduced to minimizing the expected sample size at $\theta^* = (\theta_0 + \theta_1)/2$.

Another ideal choice of θ^* should be the true value of θ. Since the true value is unknown, the maximum likelihood estimator $\hat{\theta}_n$ serves as a good alternative. Schwarz [Sch62] combined the Bayesian approach with decision theory, which put a prior distribution on θ, assigned a cost c to each observation, and considered 0-1 loss to punish the wrong decision for testing hypotheses (1.3) in the exponential family. They derived a sequential test as an asymptotic solution to this Bayes problem as $c \to 0$, which stops at

$$N = \inf\left[n \geq 1 : \max\left\{ \prod_{i=1}^{n} \frac{f_{\hat{\theta}_n}(X_i)}{f_{\theta_0}(X_i)}, \prod_{i=1}^{n} \frac{f_{\hat{\theta}_n}(X_i)}{f_{\theta_1}(X_i)} \right\} \geq c^{-1} \right]$$

and accepts H_1 or H_0 depending on which of $\prod_{i=1}^{N} f_{\theta_1}(X_i)$ and $\prod_{i=1}^{N} f_{\theta_0}(X_i)$ is larger.

Due to variation of the accuracy of the maximum likelihood estimator $\hat{\theta}_n$ with n, Lai [Lai88]

considered a time-varying boundary $g(c\,n)$ with the following stopping rule:

$$N = \inf\left[n \geq 1 : \max\left\{ \sum_{i=1}^{n} \log \frac{f_{\hat{\theta}_n}(X_i)}{f_{\theta_0}(X_i)}, \sum_{i=1}^{n} \log \frac{f_{\hat{\theta}_n}(X_i)}{f_{\theta_1}(X_i)} \right\} \geq g(c\,n) \right],$$

where g is a nonnegative function on $(0, \infty)$ with $g(t) \sim |\log t|$ as $t \to 0$. The decision of accepting H_0 or H_1 depends on $\hat{\theta}_N \leq \tilde{\theta}$ or $\hat{\theta}_N > \tilde{\theta}$, where $\tilde{\theta}$ satisfies $\mathbb{E}_{\tilde{\theta}}\left[\log\left\{f_{\tilde{\theta}}(X_1)/f_{\theta_0}(X_1)\right\}\right] = \mathbb{E}_{\tilde{\theta}}\left[\log\left\{f_{\tilde{\theta}}(X_1)/f_{\theta_1}(X_1)\right\}\right]$. It can also by applied to test

$$\mathrm{H}_0 : \theta \leq \theta_0 \quad vs. \quad \mathrm{H}_1 : \theta > \theta_0$$

by letting $\theta_1 = \theta_0$ and $\tilde{\theta} = \theta_0$. The stopping rule is simplified to

$$N = \inf\left\{ n \geq 1 : \sum_{i=1}^{n} \log \frac{f_{\hat{\theta}_n}(X_i)}{f_{\theta_0}(X_i)} \geq g(c\,n) \right\}.$$

It can be shown that this test is asymptotically optimal from both frequentist and Bayesian perspectives. This idea was generalized to multi-parameter exponential families by Lai & Zhang [Lai94].

It is known that performing repeated tests at different stages of accumulating data will cause a type I error substantially greater than the nominal value. The repeated significance tests (RST) [Arm69] were proposed to control the type I error under repeated checks, whose idea is to calculate the true significance level of repeated tests through simulation or numerical integration. They considered sequential outcomes of three distributions: binomial, normal, and exponential. Depending on the distribution of outcomes, they estimated the probability of finding significant results before n observations under the null hypothesis, $n = 1, 2, \cdots$, with the decision boundary determined by a one-stop test of significance level α'. Then given a significance level α, the maximum number of observations M and nominal level α' can be obtained by inverse interpolation. Take data from a normal distribution $N(\theta, \sigma^2)$ for example, the RST of hypotheses

$$\mathrm{H}_0 : \theta = 0 \quad vs. \quad \mathrm{H}_1 : \theta \neq 0$$

stops the sampling at

$$N = \inf\left\{ n \leq M : |S_n| \geq k\sigma\sqrt{n} \right\}, \tag{1.4}$$

where $S_n = \sum_{i=1}^{n} X_i$ and k is the $(1-\alpha'/2)$ quantile of standard normal distribution. If $N < M$, or $N = M$ and $|S_M| \geq k\sigma\sqrt{M}$, it rejects the null hypothesis H_0.

Pocock [Poc77] extended this idea to group sequential design, where a significance test is carried out after each group of observations. The stopping rule (1.4) is used with X_i now representing an approximately normally distributed statistic of data in the ith group, M being the maximum number of groups of equal size m, and σ being replaced by $\sigma/\sqrt{m}$ denoting the standard deviation of X_i. For a given significance level α and M, the value of α' does not depend on m, but the power $(1-\beta)$ of detecting δ significantly different from 0 is a function of $\sqrt{m}\delta/\sigma$, which illustrates a trade-off between power and sample size. Thus, given α, M, σ, power $(1-\beta)$, and minimum detectable effect δ, m can

be obtained by inverse interpolation of numerical values of $\beta(\delta, \sigma, \alpha, M)$. Other kinds of RST include O'Brien & Fleming [O'B79] and Wang & Tsiatis [Wan87].

1.1.3 Application to online A/B testing

There are two data streams in A/B testing: the control and the treatment. Let $\{Y_i^{(0)}\}_{i=1}^{\infty}$ and $\{Y_i^{(1)}\}_{i=1}^{\infty}$ denote the sequence of outcomes in the control and treatment groups, respectively. A typical A/B test assumes homogeneous treatment effects among the population and aims to test:

$$\mathrm{H}_0 : \theta := \mu_B - \mu_A = 0 \quad \text{vs.} \quad \mathrm{H}_1 : \theta \neq 0, \tag{1.5}$$

or

$$\mathrm{H}_0' : \theta := \mu_B - \mu_A \leq 0 \quad \text{vs.} \quad \mathrm{H}_1' : \theta > 0, \tag{1.6}$$

where μ_A and μ_B denote the population means of the control group and treatment group.

Even though sequential testing has gained great popularity in clinical trials decades ago, its advantage in online A/B testing has not been recognized until recently Johari et al. [Joh15] brought sequential testing to A/B testing. They proposed a notion of *always valid p-value* for sequential testing with definition as below, facilitating the later development of online testing.

Definition 1.1.1 *A sequence of p-values $(p_n)_{n=1}^{\infty}$ is always valid if it satisfies the property that $\forall s \in [0,1]$, $P_{H_0}(p_T \leq s) \leq s$ for any given (possibly infinite) stopping time T.*

Let $(T(\alpha), \delta(\alpha))$ denote a decision rule parameterized by significance level α, where T is a stopping time indicating the sample size at which the test is ended, and δ is a binary indicator for rejection decision. Johari et al. [Joh15] showed that an always valid p-value process $(p_n)_{n=1}^{\infty}$ and a decision rule $(T(\alpha), \delta(\alpha))$ can be derived from each other through:

$$p_n = \inf\{\alpha : T(\alpha) \leq n, \delta(\alpha) = 1\} \tag{1.7}$$

and

$$T(\alpha) = \inf\{n : p_n \leq \alpha\}, \quad \delta(\alpha) = 1\{T(\alpha) < \infty\}.$$

Static allocation is a common way of treatment assignment in A/B testing, where users are assigned to control and treatment groups with a prefixed probability (e.g. 50% vs. 50%), and the placement cannot be adjusted during the experiment based on what is observed. In the simplest case, we suppose that observations arrive in pairs. Then the sequential tests can be directly applied to $Z_i = Y_i^{(1)} - Y_i^{(0)}$, $i = 1, 2, \cdots$. The mixture sequential probability ratio test was brought to A/B testing by Johari et al. [Joh15; Joh17] to test the hypotheses (1.5). They restricted their data model to the two most common cases in practice: normal distribution with means μ_A and μ_B and equal known variance σ^2; and Bernoulli distribution with success probabilities μ_A and μ_B. For normal data, $Z_i = Y_i^{(1)} - Y_i^{(0)} \sim N(\theta, 2\sigma^2)$. By choosing a normal mixture density $\pi(\theta) \sim N(0, \tau^2)$, the test statistic Λ_n^{π} (1.2) at sample size n (each

group has n samples) now has a closed form

$$\Lambda_n^\pi = \sqrt{\frac{2\sigma^2}{2\sigma^2 + n\tau^2}} \exp\left\{ \frac{n^2\tau^2(\bar{Y}_n^{(1)} - \bar{Y}_n^{(0)})^2}{4\sigma^2(2\sigma^2 + n\tau^2)} \right\},$$

where $\bar{Y}_n^{(j)} = n^{-1}\sum_{i=1}^n Y_i^{(j)}$ for $j = 0, 1$. However, the distribution of Z_i is complicated for binary data and a conjugate mixture distribution does not exist, making it challenging to compute the integration. Thus they considered the likelihood ratio of the data pair $\{Y_i^{(0)}, Y_i^{(1)}\}$ with parameters $\theta = \mu_B - \mu_A$ and $\bar{\mu} = (\mu_B + \mu_A)/2$ after reparameterization, and then reduced it to the likelihood ratio based on the sufficient statistics $\bar{Y}_n^{(1)} - \bar{Y}_n^{(0)}$ for θ by applying *Fisher's factorization theorem*. By the *central limit theory*, $\bar{Y}_n^{(1)} - \bar{Y}_n^{(0)}$ asymptotically follows a normal distribution $N(\theta, V_n/n)$, where $V_n = \mu_A(1-\mu_A) + \mu_B(1-\mu_B)$. After replacing V_n with its estimator $\hat{V}_n = \bar{Y}_n^{(0)}(1 - \bar{Y}_n^{(0)}) + \bar{Y}_n^{(1)}(1 - \bar{Y}_n^{(1)})$ and choosing a mixture density $\pi(\theta) \sim N(0, \tau^2)$, the test statistic can be shown to have a closed form

$$\Lambda_n^\pi = \sqrt{\frac{\hat{V}_n}{\hat{V}_n + n\tau^2}} \exp\left\{ \frac{n^2\tau^2(\bar{Y}_n^{(1)} - \bar{Y}_n^{(0)})^2}{2\hat{V}_n(\hat{V}_n + n\tau^2)} \right\}.$$

The decision rule of mSPRT stays the same

$$T(\alpha) = \inf\{n : \Lambda_n^\pi \geq \alpha^{-1}\}, \quad \delta(\alpha) = 1\{T(\alpha) < \infty\}$$

with an always valid p-value nonincreasing in n derived from it by (1.7):

$$p_n = \min\left\{ \frac{1}{\max_{m \leq n} \Lambda_m^\pi}, 1 \right\}.$$

However, this decision rule implies that if the true effect is zero or negligible, one can only wait essentially infinitely to get the right conclusion which is not practical in reality. Johari et al. [Joh15] proposed a *failure time* M so that we would accept the null hypothesis if we could not reject it before the sample size M. They proved that for small significance level α, the mSPRT truncated at M can achieve similar asymptotic power to the uniformly most powerful (UMP) fixed-horizon test with average smaller sample size. They also showed that for normal data and a true normal prior $G(\cdot) \sim N(0, \tau_0^2)$ of θ, a rough match between the mixture density $\pi(\cdot)$ and the prior density $G(\cdot)$ will minimize the expected sample size over $\theta \sim G(\cdot)$.

Abhishek & Mannor [Abh17] considered the case where the metric of interest is a complicated scalar function of several metrics rather than some simple averages, e.g. the ratio of the expected number of successful queries per user session to the expected number of total queries per user session. In such cases, the commonly used exponential families such as normal or Bernoulli distribution are not suitable. They combined the ideas from bootstrap and mSPRT, and proposed a nonparametric sequential test for such complex metrics. They demonstrated that their method works for diverse types of metrics and is robust to misspecification in the distribution of generating data.

Sometimes more than one treatment is compared against a baseline variant (i.e. the control) in an

experiment and each user is randomly allocated to one of these variants. Such tests are called *multivariate testing*. Johari et al. [Joh17] treated each comparison as a two-variant sub-experiment and computed an always valid p-value for each of them. Then the multivariate testing transforms into a sequential multiple comparisons problem. They proved that the two well-studied multiple testing techniques, Bonferroni correction and Benjamini-Hochberg procedure, can also be applied to always valid p-values to control the family-wise error rate and false discovery rate respectively at any time.

A shortcoming of static allocation is that if one variant is better early on in the experiment, we would bear an opportunity cost by having users suffer from a bad experience in a sub-optimal variant. Adaptive allocation is a way that allows users to have higher probabilities of being assigned to a better variant based on previous observations. It may require a *multi-armed bandit* (MAB) policy: a policy to allocate users between competing variants which maximizes the expected reward, i.e. the average outcomes, where each variant's value is only partially known at the time of allocation and may become better understood as more users are allocated to that variant. MAB policy optimally trades off the exploration of all variants against exploitation of the better performing variant.

Let $m(t)$, $n(t)$ represent the number of observations in each group among the first t total observations. Johari et al. [Joh15] proved that under an arbitrary allocation rule in an A/B test, the test statistics of mSPRT based on the joint likelihood of $\mathbf{Y}^{(0)}_{1:m(t)}$ and $\mathbf{Y}^{(1)}_{1:n(t)}$

$$\Lambda^{\pi}_{t} = \int_{\Theta} \frac{f\left(\mathbf{Y}^{(0)}_{1:m(t)}, \mathbf{Y}^{(1)}_{1:n(t)}; \theta, \bar{\mu}\right)}{f\left(\mathbf{Y}^{(0)}_{1:m(t)}, \mathbf{Y}^{(1)}_{1:n(t)}; 0, \bar{\mu}\right)} \pi(\theta) d\theta$$

is still a martingale, where $\theta = \mu_B - \mu_A$ and $\bar{\mu} = (\mu_B + \mu_A)/2$ are parameters after reparameterization. This property applies to multi-armed bandits allocation because the arm pulled at $t+1$ is conditionally independent of the value observed, given the first t observations. The type I error is therefore controlled under adaptive allocation. However, the test statistic requires that the parameter $\bar{\mu}$ is known. Despite we can estimate it by $\{\bar{Y}^{(0)}_{m(t)} + \bar{Y}^{(1)}_{n(t)}\}/2$, the type I error will not be controlled at α anymore.

Ju et al. [Ju19] modified the Girshick test [Gir46], a variant of SPRT, to test the hypothesis (1.6) on binary data. The new test is able to control the type I error and power under both static allocation and adaptive allocation such as *Thompson sampling*. They chose fixed values $\mu_0 < \mu_1$ and transformed hypotheses (1.6) into a pair of simple hypotheses

$$\mathrm{H}_0 : \mu_A = \mu_0, \mu_B = \mu_1, \quad \mathrm{H}_1 : \mu_A = \mu_1, \mu_B = \mu_0,$$

where μ_A and μ_B are the success probabilities of each group. Then they estimated the log-likelihood ratio by

$$\widehat{\log \mathrm{LR}}_t = (-\Delta) \cdot \frac{2mn}{t} \cdot \left\{ \frac{1}{n} \sum_{i=1}^{n} Y_i^{(1)} - \frac{1}{m} \sum_{i=1}^{m} Y_i^{(0)} \right\}$$

for static allocation and by

$$\widehat{\log \mathrm{LR}}_t = (-\Delta) \cdot \sqrt{mn} \cdot \left\{ \frac{1}{n} \sum_{i=1}^{n} Y_i^{(1)} - \frac{1}{m} \sum_{i=1}^{m} Y_i^{(0)} \right\}$$

for Thompson sampling, where m and n are the numbers of observations in each group among the first t total observations, and $\Delta = \log\left[\{(1-\mu_1)\mu_0\}/\{(1-\mu_0)\mu_1\}\right]$. The experiment is terminated at

$$T = \inf\left\{ t \geq 1 : \widehat{\log \mathrm{LR}}_t \geq \log C \text{ or } \widehat{\log \mathrm{LR}}_t \leq \log D \right\}$$

and accepts H_1 or H_0 according to $\widehat{\log \mathrm{LR}}_T \geq \log C$ or $\widehat{\log \mathrm{LR}}_T \leq \log D$. By setting $C = (1-\beta)/\alpha$ and $D = \beta/(1-\alpha)$, their simulation showed that the test under static allocation can control type I and type II error at α and β, while there is some minor inflation in probability errors under Thompson sampling.

Google Analytics also used Thompson sampling as their allocation strategy for the multivariate test on their Content Experiment platform [Sco13; Lu19]. It supposes that there are K variants we want to compare and each has an average reward $\mu_1, \mu_2, \cdots, \mu_K$. The goal of the experiment is to find the best variant while simultaneously achieving the highest average outcomes in doing so. Each μ_k has a prior distribution and the posterior distribution is sequentially updated as data is collected. Thompson sampling assigns a user to arm k ($k = 1, 2, \cdots, K$) based on the probability of the kth arm being the optimal one conditional on previous data. The experiment is terminated when $(1-\alpha)$ of the samples in a Monte Carlo simulation have remaining values less than 1% of the winning arm's value. It showed that this method requires much smaller samples than the fixed-horizon test and has less opportunity cost compared to static allocation. However, the simulations showed that this method has the false positive rate of finding the best variant higher than the significance level.

1.2 Semiparametric Estimation of Time-Dependent Treatment Effects

1.2.1 Background

Mobile games are getting increasingly popular nowadays, becoming the most significant segment of the video game industry beyond console and PC games. Based on a new industry study by Golden Casino News [Gre19], mobile games make up 60% of revenue for the global video game market with 1.36 billion users worldwide and this number will continue to grow. A common monetization strategy for mobile games is the freemium business model [And09]. It provides free downloads and basic gameplay to attract customers and then sells premium-priced add-on components to generate revenue. It is reported that over 90% of mobile games begin as free, and over 90% of the profits from mobile games come from games that began as free [App20]. In freemium mobile games, developers often offer promotions to players to retain users and stimulate consumption, such as sales on the add-on components, more in-game rewards, and holiday promotions. Understanding the effects of a sequence of promotion decisions on daily engagement (i.e. purchase) for heterogeneous users helps the game managers make improved/personalized promotion strategies.

Freemium mobile game data has its unique characteristics, which lead to several challenges in statistical inference. First, the outcomes (i.e. daily engagement) are nonnegative with a clump of observations at zero. It is common to remain free users and never engage in the premium part of freemium games. Moreover, the positive daily engagement is skewed to the right. A 2014 study of freemium mobile games found that 50% of mobile gaming revenue came from the top 10% of players making purchases, which only accounts for 0.15% of total players [Swr14]. Common distributional assumptions such as the normal or gamma are no longer appropriate for freemium mobile game data. Second, players are followed over a period of time, during which a sequence of promotion decisions (promotion or no promotion, referred to as treatments) are implemented. We are interested in estimating the joint effects of a sequence of treatments on the outcomes rather than a one-time treatment. However, there exist time-varying confounders with the following characteristics: they are (i) associated with the outcomes, (ii) affected by earlier treatment, and (iii) predictive of subsequent treatments. In the presence of time-varying confounders, standard regression methods, whether or not adjusting for time-varying confounders, are inappropriate for estimating the causal effects of a sequence of treatments [Rob09]. Lastly, in practice, the promotion assignment over time can be personalized or uniform. We refer to the treatment assignment that allocates the same treatment to all users at a given time as *cluster randomization*, while the personalized treatment assignment is *individual randomization*. In cluster randomization, since all users receive the same treatment at a given time, it increases the difficulties in estimation of treatment effects due to the convergence of the propensity score estimate reliant solely on the number of time points, and thus the standard asymptotic framework requiring only the sample size go to infinity is not sufficient to derive the large sample results.

Exiting works for estimating the causal effect of one-time treatment with semicontinuous outcomes with excessive zeros include the *Two-part model* [Dua83], the *Burden-of-Illness model* [Cha94], and the *Tobit model* or its variants [Tob58; Pow86; Kee19; Che20]. However, it is considerably more challenging to estimate the causal effects of a sequence of treatments in the presence of time-varying confounding. *Structural nested mean models* [Rob94a] have been proposed to overcome this challenge by modeling treatment effects sequentially over time, and G-estimation can be used to disentangle treatment effects from confounding effects. See [Van14] for a review. However, existing structural nested mean models cannot directly handle zero-inflated outcomes, and the standard asymptotic regime for G-estimation requires the sample size to increase to infinity and does not apply to the case when the number of follow-up times increases while the sample size can be fixed. In the next section, we will give a review of structural nested mean models. In Chapter 4, we will propose a class of *multiplicative structural nested mean models* for the joint effects of a sequence of treatments on zero-inflated nonnegative outcomes.

1.2.2 Review of structural nested mean models

Structural nested mean models [Rob94a; Rob00] (SNMM) and the associated G-estimation are proposed for modeling and estimating the joint effects of a sequence of treatments in the presence of time-varying confounders.

We start with the setting of point treatment. Let Y, A, L denote the observed outcome, treatment, and

covariates, respectively, and $Y^{(a)}$ denote the potential outcome that would have been seen had the subject received treatment a. Structural nested mean models parameterize the average treatment effects of a on subjects with covariates l as

$$\mathbb{E}(Y^{(a)}|A=a,L=l)-\mathbb{E}(Y^{(0)}|A=a,L=l)=r(l,a;\theta_0), \tag{1.8}$$

where θ_0 is the true unknown finite-dimensional parameter and $r(l,a;\theta)$ is a known function which is smooth in θ and satisfies $r(l,0;\theta)=0$ for all l and θ.

Consistency assumption and no unmeasured confounders assumption are made to estimate the parameter θ: (i) the observed outcomes is equal to the potential outcome under the actual treatment received, i.e. $Y=Y^{(a)}$; (ii) the covariates L is sufficient to adjust for the confounding of the association between the treatment A and the potential outcome $Y^{(0)}$, i.e. $A \perp\!\!\!\perp Y^{(0)} \mid L$. Then the SNMM constructs a variable

$$H(\theta) \coloneqq Y - r(l,a;\theta),$$

which mimics the potential outcome $Y^{(0)}$ that would have been seen had the treatment been removed. It can be shown that under the consistency assumption, we have

$$\mathbb{E}\{H(\theta_0)|A,L\}=\mathbb{E}\{Y^{(0)}|A,L\}.$$

The no unmeasured confounders assumption together with the SNMM (1.8) implies that

$$\mathbb{E}\{H(\theta_0)|A,L\}=\mathbb{E}\{H(\theta_0)|L\}.$$

Thus the estimation of θ_0 can be based on solving the following estimating equations:

$$0=\sum_{i=1}^{n}[d(A_i,L_i)-\mathbb{E}\{d(A_i,L_i)|L_i\}]\cdot[H_i(\theta)-\mathbb{E}\{H_i(\theta)|L_i\}], \tag{1.9}$$

where $d(A,L)$ is an arbitrary function of the dimension of θ. When the variation of $H(\theta_0)$ given A,L is constant, an efficient estimator of θ_0 can be obtained by choosing

$$d(A,L)=\mathbb{E}\left\{\frac{\partial H(\theta)}{\partial \theta}\Big|_{\theta=\theta_0}\Big|A,L\right\}.$$

under the no unmeasured confounders assumption. Evaluating equations (1.9) requires parametric models for the conditional distribution of A and the conditional expectation of $H(\theta_0)$:

$$f(A|L)=f(A|L;\psi_0)$$
$$\mathbb{E}\{H(\theta_0)|L\}=g(H(\theta_0)|L;\phi_0),$$

where $f(A|L;\psi)$ and $g(H(\theta_0)|L;\phi)$ are known functions, smooth in ψ or ϕ, and ψ_0,ϕ_0 are true unknown finite-dimensional parameters. An estimator of θ is obtained by solving the equations (1.9) with ψ_0 and

ϕ_0 replaced by their consistent estimators under their models $f(A|L;\psi)$ and $g(H(\theta_0)|L;\phi)$, and it is called a G-estimator. It can be shown that the G-estimator is consistent if either $f(A|L;\psi)$ or $g(H(\theta_0)|L;\phi)$ is correctly specified.

In the setting of sequential treatments, we suppose that the measurements are collected at T discrete time points. Let Y_t, A_t, L_t denote the observed outcome, treatment, and covariates at time t, $t = 1, \cdots, T$. The observed data are assumed to be ordered as L_1, A_1, Y_1, L_2, etc. We use the overline to denote the history of a variable and the underline to stand for the future of a variable, e.g. $\overline{A}_t = (A_1, \cdots, A_t)$, $\underline{A}_t = (A_t, \cdots, A_T)$. Let $Y_t^{(\overline{a}_t)}$ denote the potential outcome that would have been seen at time t had the subject received the treatment sequences $\overline{a}_t$ through time t. We assume that the treatment after time t cannot affect the outcome at times up to t. Let $(\overline{a}_t, 0)$ denote the sequence of treatments that agrees with $\overline{a}_t$ through time t and is 0 thereafter. The SNMM models the effects of removal of treatment at time t on subsequent average outcomes:

$$\mathbb{E}(\underline{Y}_t^{(\overline{a}_t,0)}|\overline{A}_t = \overline{a}_t, \overline{L}_t = \overline{l}_t) - \mathbb{E}(\underline{Y}_t^{(\overline{a}_{t-1},0)}|\overline{A}_t = \overline{a}_t, \overline{L}_t = \overline{l}_t) = r_t(\overline{l}_t, \overline{a}_t; \theta_0), \quad t = 1, \cdots, T, \tag{1.10}$$

where $r_t(\overline{l}_t, \overline{a}_t; \theta)$ is a known $(T + 1 - t)$-dimensional function, smooth in θ, and satisfying $r_t(\overline{l}_t, \overline{a}_{t-1}, 0; \theta) = 0$ for all $\overline{l}_t, \overline{a}_{t-1}$ and θ.

Similar to the point treatment setting, we make the following two assumptions: (i) $Y_t = Y_t^{(\overline{a}_t)}$ if $\overline{A}_t = \overline{a}_t$ for $t = 1, \cdots, T$; (ii) $A_t \perp\!\!\!\perp \underline{Y}_t^{(\overline{a}_{t-1},0)} \,|\, \overline{L}_t, \overline{A}_{t-1} = \overline{a}_{t-1}$ for $t = 1, \cdots, T$; and construct a $(T+1-t)$-dimensional variable $H_t(\theta)$ with components

$$Y_k - \sum_{l=t}^{k} r_{l,k}(\overline{L}_l, \overline{A}_l; \theta), \quad k = t, \cdots, T,$$

where $r_{l,k}(\overline{L}_l, \overline{A}_l; \theta)$ is a component of $r_l(\overline{L}_l, \overline{A}_l; \theta)$ that parameterizes the effects on Y_k. Under the above assumptions and SNMM (1.10), we have

$$\mathbb{E}\{H_t(\theta_0)|\overline{A}_{t-1} = \overline{a}_{t-1}, A_t, \overline{L}_t\} = \mathbb{E}\{\underline{Y}_t^{(\overline{a}_{t-1},0)}|\overline{A}_{t-1} = \overline{a}_{t-1}, A_t, \overline{L}_t\}, \quad t = 1, \cdots, T,$$

and

$$\mathbb{E}\{H_t(\theta_0)|\overline{A}_t, \overline{L}_t\} = \mathbb{E}\{H_t(\theta_0)|\overline{A}_{t-1}, \overline{L}_t\}, \quad t = 1, \cdots, T.$$

Then the estimating equations are constructed as follows:

$$0 = \sum_{i=1}^{n}\sum_{t=1}^{T}\big[d_t(\overline{A}_{it}, \overline{L}_{it}) - \mathbb{E}\{d_t(\overline{A}_{it}, \overline{L}_{it})|\overline{A}_{i,t-1}, \overline{L}_{it}\}\big] \cdot \big[H_{it}(\theta) - \mathbb{E}\{H_{it}(\theta)|\overline{A}_{i,t-1}, \overline{L}_{it}\}\big],$$

where $d_t(\overline{A}_{it}, \overline{L}_{it}), t = 1, \cdots, T$ is an arbitrary $p \times (T+1-t)$-dimensional function, with p the dimension of θ.The choice of $d_t(\overline{A}_{it}, \overline{L}_{it})$ for an efficient estimator of θ_0 and the doubly robust property of the G-estimator are similar to the setting of point treatment.

1.3 Outline

The rest of this thesis is structured as follows. In Chapter 2, we propose an online test, named *sequential score test* (SST), to detect a multi-dimensional heterogeneous treatment effect under a generalized linear model. We provide an online p-value calculation for SST and extend SST to an online multiple testing setting by controlling the false discovery rate. We examine the empirical performance of the proposed test and compare it with the mixture sequential probability ratio test (mSPRT) using simulations and a real dataset. The results show that the SST controls the type I error at any time, has higher detection power, and allows quick inference on online A/B testing.

In Chapter 3, we propose an online test for subgroup treatment effects based on value difference, named SUBTLE. It aims to test if some subgroup of the population will benefit from the investigative treatment. If the testing result indicates the existence of such a subgroup, a subgroup will be identified using a readily available estimated optimal treatment rule. The validity of the SUBTLE has been proved by both theoretical and simulation results. The experiments show that the SUBTLE has high detection power with controlled type I error at any time, is robust to noise covariates, and achieves early stopping compared with the corresponding fixed-horizon test.

In Chapter 4, we propose a class of multiplicative structural nested mean models for zero-inflated nonnegative outcomes, which flexibly describes the joint effect of a sequence of treatments in the presence of time-varying confounders. The proposed estimator solves a doubly robust estimating equation, where the nuisance functions, propensity score and conditional outcome means given confounders, are estimated parametrically or nonparametrically. To improve the accuracy, we leverage the characteristic of zero-inflated outcomes by estimating the conditional means in two parts, that is, separately modeling the probability of having positive outcomes given confounders and the mean outcome conditional on its being positive and confounders. We show that under the individual randomization of treatment assignment the proposed estimator is consistent and asymptotically normal as either the sample size or the follow-up time goes to infinity. Similar results are also established under cluster randomization as the follow-up time goes to infinity. Moreover, the typical sandwich formula can be used to estimate the variance of treatment effect estimators consistently, without accounting for the variation due to estimating nuisance functions. Simulation studies and an application to a freemium mobile game dataset are presented to demonstrate the empirical performance of the proposed method and support our theoretical findings.

The theorem proofs and additional simulation results of Chapters 2-4 are presented in the Appendices.

CHAPTER

2

A NEW FRAMEWORK FOR ONLINE TESTING OF HETEROGENEOUS TREATMENT EFFECTS

2.1 Introduction

Randomized controlled experiment, also known as *A/B testing*, is widely used in web-facing industry to improve products and technologies in a data-driven manner [Koh09]. Most A/B tests are conducted by performing formal null hypothesis statistical testing (NHST) with the typical null hypothesis $H_0 : \theta := \mu_B - \mu_A = 0$ to determine if the difference of the metric across the two variants is significant or not. The result of NHST is summarized in a *p-value* and the case that the p-value is less than a preset *significance level* α will lead the null hypothesis to be rejected. A valid test is able to get a high power to detect the difference if there is while controlling the *type I error*, i.e. the probability of erroneously rejecting H_0, to be less than α.

However, the validity of NHST requires that the sample size is fixed in advance, which is often violated in practice. In A/B testing practice, a fast-paced product evolution pushes its shareholders to continuously monitor the p-values and draw conclusions prematurely. In fact, stopping experiments in an adaptive manner can favorably bias getting significant results and lead to very high false positive probabilities, well in excess of the nominal significance level [Goo14; Sim11]. As an extreme example in [Pek15], it can be shown that stopping the first time that the p-value is less than α actually has a type I error probability of 1. Yet for all that, this "peeking" behavior is not without reasons. The time cost and

opportunity cost for such *fixed-horizon* hypothesis testing are large [Ju19], so users want to find the true effects and stop the experiments as quickly as possible. Moreover, the sample size calculation requires an estimate of the *minimum detectable effect* (MDE). Most users lack good prior knowledge of the trade-off between high detection ability and short waiting time and may want to adjust them after peeking early at results.

Another problem with A/B testing is that it assumes there is only an *average treatment effect* (ATE) in the population of the experiment. But underlying this average effect may be substantial variation in how particular subgroups respond to treatments: there may be *heterogeneous treatment effects* (HTE) [Gri17]. It might be that the population average effect of a product with a new feature is not significant, but the feature does benefit a lot among particular subgroups of users. In this case, we won't be able to detect those effects and will lose the chance of making profits by promoting new products to those target subpopulations, if only ATE is tested in A/B testing.

To address the continuously monitoring problem, *sequential testing* was first developed by Wald [Wal45], who introduced the *sequential probability ratio test* (SPRT). Sequential testing allows intermediate checks of significance while providing type I error control at any time; see Lai [Lai01] for a survey on sequential testing. Moreover, sequential testing could help decision-makers conclude an experiment earlier with often much fewer samples than fixed-horizon testing [Wal45]. *Mixture sequential probability ratio test* (mSPRT) [Rob70] and *maximized sequential probability test* (MaxSPRT) [Kul11] are two variants of sequential testing that generalized SPRT to a composite hypothesis. Due to the merits of mSPRT that it is a test with power 1 [Rob74] and almost optimal [Pol78] with respect to the expected time to stop, it was brought to A/B testing by Johari et al. [Joh15; Joh17]. They also proposed a notion of *always valid p-value process* (i.e. sequential p-values) in the same papers and used it as a tool for converting fixed-horizon *multiple testing* procedures to sequential versions. Later, Malek et al. [Mal17] showed that if the original multiple testing procedure has a type I error guarantee in a certain family (including false discovery rate and family-wise error rate), then the sequential conversion inherits an analogous guarantee.

However, current online testing procedures, such as mSPRT, are not suitable for testing heterogeneous treatment effects due to two aspects. First, they can not accommodate the nuisance parameters in the baseline effects. Second, they may not be able to control the type I error and may lack power for detecting heterogeneous treatment effects. In this chapter, we propose a new framework for online testing of heterogeneous treatment effects. The proposed test, named SST, is based on the ratio of asymptotic score statistic distributions, which is able to test multi-dimensional parameters. Furthermore, the asymptotic normality of the score functions guarantees an explicit form of the integral, which allows the integration of the ratio to be efficient. At last, we generalize our framework to online multiple testing, which is often the case in industrial practice.

The remainder of this chapter is structured as follows. In Section 2.2, we introduce some preliminary knowledge about fixed-horizon testing and sequential testing. In Section 2.3, we present the proposed new framework for online testing of heterogeneous treatment effects. We extend SST to multiple testing settings in Section 2.4 and conduct experiments in Section 2.5 to compare our framework with the

widely-used mSPRT. Finally, in Section 2.6, we conclude the chapter and present future directions.

2.2 Preliminaries

2.2.1 Fixed-horizon testing

Fixed-horizon testing is the most widely used procedure in industry where the sample size is fixed in advance. It can be broken down into several steps [Leh06]:

Step 1: Determine the desired significance level α, minimum detectable effect (MDE), and power at MDE. It means that the probability to detect the MDE is at least at the value of power, while the probability of rejecting H_0, if it is actually true, is at most α.

Step 2: Calculate/Estimate the minimum sample size n. The sample size n needs to be large enough to achieve the desired power at MDE while controlling type I error at a significance level, but too large sample size will lead to more opportunity cost of waiting for more samples. One needs to balance these two aspects when choosing the sample size.

Step 3: Collect n samples and compute the observed value of an appropriate test statistics Λ_n. The most common test statistics for two-sample tests are z-tests and t-tests, which assume that data are from a normal distribution with known or unknown variance, respectively.

Step 4: Compute a p-value p_n and reject the null hypothesis if $p_n \leq \alpha$. P-value is a random variable to denote the probability of seeing a test statistic as extreme as the observed statistics Λ_n under the null hypothesis, and can be formally defined as

$$p_n = \inf\{\alpha : \Lambda_n \geq k(\alpha)\}, \tag{2.1}$$

where $k(\alpha)$ is a critical value depending on the significance level and the distribution of Λ_n under H_0. The critical value is determined such that, under the null hypothesis H_0, the event $\Lambda_n \geq k(\alpha)$ occurs with a probability no greater than α. Since the p-value was computed assuming a fixed sample size n, we refer to this as a fixed-horizon p-value. Small p-values suggest evidence in support of the alternative hypothesis.

A *decision rule* is a pair (T, δ) representing a testing, where T is a stopping time indicating the sample size at which the test is ended, and δ is a binary indicator for rejection decision. With the definition of fixed-horizon p-value in (2.1), it is obvious to see that (n, δ_1) and (n, δ_2) with $\delta_1 = 1\{p_n \leq \alpha\}$ and $\delta_2 = 1\{\Lambda_n \geq k(\alpha)\}$ are two equivalent decision rules for fixed-horizon testing. That means the decision rule and p-value can be obtained from each other: find p-value from decision rule (n, δ_2) by (2.1), or make the decision (n, δ_1) from p-value. Hence, we can actually stop at step 3 and reject H_0 if $\Lambda_n \geq k(\alpha)$ for some predetermined significance level α. Nonetheless, the decision-making process using p-values is remarkably simple and transparent: one can choose their own significance level and make a valid decision.

2.2.2 Sequential testing

Sequential testing, in contrast to fixed-horizon, is a procedure where the decision of terminating the sampling process at any stage of the experiment depends on the results of the observations previously made. It has gained recent popularity in online A/B testing [Bal15; Joh15] due to its flexibility of continuously monitoring and ending the experiment as soon as significant results are observed.

The decision rules for sequential testing is a nested family of $(T(\alpha), \delta(\alpha))$, parameterized by significance level α. It has the following two properties [Joh17]: first, the type I error is controlled, that is, $P_{H_0}(\delta(\alpha) = 1) \leq \alpha$; second, $T(\alpha)$ is (almost surely) nonincreasing in α while $\delta(\alpha)$ is (almost surely) nondecreasing in α. In other words, less stringent type I error control allows the test to stop sooner and is more likely to lead to rejection.

Similar to fixed-horizon testing, a notion of sequential p-values was also introduced for sequential testing and named *always valid p-value process* by [Joh17]: *a sequence of fixed-horizon p-values* $(p_n)_{n=1}^{\infty}$ *is always valid if it satisfies the property that* $\forall s \in [0,1]$, $P_{H_0}(p_T \leq s) \leq s$ *for any given (possibly infinite) stopping time* T. It allows the user to balance detection power and sample size dynamically as they see fit while still control type I error. Same as the fixed-horizon p-value, the always valid p-value process can be derived from the decision rule for a sequential test, and vice versa. For a given sequential test $(T(\alpha), \delta(\alpha))$,

$$p_n = \inf\{\alpha : T(\alpha) \leq n, \delta(\alpha) = 1\} \tag{2.2}$$

defines an always valid p-value process. For any always valid p-value process $(p_n)_{n=1}^{\infty}$, a sequential test is obtained as follows:

$$T(\alpha) = \inf\{n : p_n \leq \alpha\} \quad \delta(\alpha) = 1\{T(\alpha) < \infty\}.$$

The mixture sequential probability ratio test (mSPRT) [Rob70] is a well-studied family of sequential tests. Its test statistic based on the first n observations Λ_n^{π} is a mixture of likelihood ratios against the null hypothesis, with the mixture density $\pi(\cdot)$ over the space for target parameter θ. The decision rule for mSPRT is as below:

$$T(\alpha) = \inf\{n : \Lambda_n^{\pi} \geq \alpha^{-1}\} \quad \delta(\alpha) = 1\{T(\alpha) < \infty\}.$$

It can be shown that the type I error for mSPRT is well controlled at α by a simple application of *optional stopping theorem* [Gri01], since the likelihood ratio under H_0 is a nonnegative martingale with the initial value equal to one and so is the mixture of such likelihood ratios; see Malek et al. [Mal17] and Pekelis et al. [Pek15] for a detailed proof.

Johari et al. [Joh17], recently, have brought mSPRT to online A/B tests where testing parameters μ_A, μ_B are assumed to be the mean of a Bernoulli or normal distribution, depending on whether the data is binary or continuous. They modified the original mSPRT to make it applicable to industrial A/B tests based on some approximation techniques and empirically showed that the new test has high detection performance with type I error control.

2.2.3 Heterogeneous treatment effects

Up to now, all the online A/B tests we have talked about are focusing only on testing the *average treatment effect* (ATE). However, treatment effects are commonly believed to be varying among individuals, and individual treatment effects may differ in magnitude and even have opposite directions. This is called the *heterogeneous treatment effect* (HTE). Testing HTE could help us identify subpopulations where treatment shows better performance and allow personalized treatment as well.

To give a better insight into the difference between ATE and HTE testing, let's take the generalized linear model (GLM) for example. Let Y_i denote the observed response for individual i and assume they independently follow an exponential family

$$\begin{aligned} &Y_i \overset{ind}{\sim} \text{Exponential Family}(\gamma_i, \phi), \quad i = 1, \cdots, n \\ &f_{Y_i}(y_i|\gamma_i, \phi) = \exp\left\{\frac{y_i\gamma_i - b(\gamma_i)}{a_i(\phi)} + c(y_i, \phi)\right\}, \end{aligned} \tag{2.3}$$

where n denotes the sample size; $a_i(\cdot)$, $b(\cdot)$ and $c(\cdot,\cdot)$ are known functions; γ_i is the *canonical parameter* and ϕ is a typically known *dispersion parameter*. They are related to the mean and variance of the response through:

$$\mu_i = \mathbb{E}(Y_i) = b'(\gamma_i), \qquad \text{Var}(Y_i) = a_i(\phi) \cdot b''(\gamma_i).$$

A link function $g(\cdot)$ provides the relationship between the linear predictor and the mean of the response:

$$g(\mu_i) = \eta_i. \tag{2.4}$$

where the linear predictor η_i has different forms depending on either ATE or HTE setting. There always exists a well-defined *canonical link* (i.e. a link function such that $g(\mu_i) = \gamma_i$) derived from the response's density function. For example, the normal distribution has an identity function $g(\mu_i) = \mu_i$ as the canonical link, the Bernoulli distribution has a logit link $g(\mu_i) = \log\{\mu_i/(1-\mu_i)\}$, and the Poisson distribution has a log link $g(\mu_i) = \log\mu_i$. HTE and ATE testings have different assumptions about the form of the linear predictor. ATE testing assumes that

$$\eta_i = \beta + \theta A_i \tag{2.5}$$

and test $H_0 : \theta = \theta_0$, whereas HTE testing assumes that

$$\eta_i = \boldsymbol{\beta}^T \mathbf{X}_i + (\boldsymbol{\theta}^T \mathbf{X}_i) A_i \tag{2.6}$$

and test $H_0 : \boldsymbol{\theta} = \boldsymbol{\theta}_0$, where $\mathbf{X}_i$ denotes the covariates vector with the first element being 1 indicating the intercept, and A_i denotes the binary treatment. Note that $\boldsymbol{\theta}$ and $\boldsymbol{\beta}$ in HTE testing are both vectors since at least one covariate is considered.

In the case of HTE testing, mSPRT does not work well for the following reasons:

1. The test statistic may not have an explicit form if a conjugate prior $\pi(\cdot)$ for likelihood ratio doesn't exist, as is often the case in HTE testing, e.g. logistic regression. As a result, the computation is

inefficient to implement in a streaming environment.

2. The nuisance parameter $\boldsymbol{\beta}$ in the likelihood function is unknown. Even though it can be replaced by its estimator, the resulting test statistics is no longer a martingale and hence the type I error cannot be controlled. Johari et al. [Joh15] used a sufficient statistic for the nuisance parameter and applied the *central limit theory* to deal with this issue in A/B tests with Bernoulli distribution. However, this technique failed to be extended to the HTE setting.

Therefore, we want to develop a valid online test that can deal with heterogeneous treatment effects.

2.3 A New Framework of Sequential Testing

In this section, we propose a new framework of sequential testing, called *sequential score test* (SST), which is able to test heterogeneous treatment effects while accounting for unknown individual effects. This framework is applicable to independent observations from an exponential family, which includes a large set of commonly used distributions.

Instead of using integrated likelihood ratios as in mSPRT, we consider the integration of the ratios of asymptotic score statistic distributions under the local alternative against the null hypothesis. The proposed method can naturally handle nuisance parameters in testing HTE. In addition, the asymptotic representation of the score statistics under the local alternative and the null hypotheses (established in Lemma 2.3.1) can lead to a martingale structure under the null similarly as for the integrated likelihood ratio statistics, and the resulting test statistic has a closed form for integration, which facilitates the implementation of the proposed testing procedure.

2.3.1 Sequential score test

Suppose we have i.i.d. data $(Y_i, A_i, \mathbf{X}_i)$, where Y, A, $\mathbf{X}$ respectively denote response, binary treatment, and $(p+1)$-dimensional covariates vector including an intercept, respectively. We assume that the distribution of Y_i conditional on $(A_i, \mathbf{X}_i)$ is an exponential family defined in (2.3)-(2.4) with η_i in the form of (2.6), where $\boldsymbol{\theta}$ and $\boldsymbol{\beta}$ denote the heterogeneous treatment effect and baseline effect, respectively. We want to test null hypothesis $\mathrm{H}_0 : \boldsymbol{\theta} = \boldsymbol{\theta}_0$ against the local alternative $\mathrm{H}_1 : \boldsymbol{\theta} = \boldsymbol{\theta}_0 + \boldsymbol{\delta}/\sqrt{n}$ $(\boldsymbol{\delta} \neq \mathbf{0})$.

To introduce the test statistic of SST, let's start with some notations. For ease of exposition, we suppose that each group has n observations and denote the response of the i-th individual in each group with $Y_i^{(a)}, a = 0, 1$. Let $\mathbf{S}^{(1)}_{n,\boldsymbol{\theta}}(\boldsymbol{\beta}, \boldsymbol{\theta})$ denote the score function of $\boldsymbol{\theta}$ for the treatment group $(A = 1)$:

$$\mathbf{S}^{(1)}_{n,\boldsymbol{\theta}}(\boldsymbol{\beta}, \boldsymbol{\theta}) = \sum_{i=1}^{n} \left[\frac{\partial \mu_i^{(1)}(\boldsymbol{\theta}, \boldsymbol{\beta})}{\partial \boldsymbol{\theta}^T} \cdot \frac{\{Y_i^{(1)} - \mu_i^{(1)}(\boldsymbol{\theta}, \boldsymbol{\beta})\}}{a_i(\phi) \cdot V_i^{(1)}(\boldsymbol{\theta}, \boldsymbol{\beta})} \right]$$

and $\mathbf{S}^{(0)}_{n,\boldsymbol{\beta}}(\boldsymbol{\beta})$ denote the score function of $\boldsymbol{\beta}$ for the control group $(A = 0)$:

$$\mathbf{S}^{(0)}_{n,\boldsymbol{\beta}}(\boldsymbol{\beta}) = \sum_{i=1}^{n} \left[\frac{\partial \mu_i^{(0)}(\boldsymbol{\beta})}{\partial \boldsymbol{\beta}^T} \cdot \frac{\{Y_i^{(0)} - \mu_i^{(0)}(\boldsymbol{\beta})\}}{a_i(\phi) \cdot V_i^{(0)}(\boldsymbol{\beta})} \right],$$

where $\mu_i^{(0)}(\boldsymbol{\beta}) = \mathbb{E}(Y_i|A_i = 0, \mathbf{X}_i)$, $\mu_i^{(1)}(\boldsymbol{\theta}, \boldsymbol{\beta}) = \mathbb{E}(Y_i|A_i = 1, \mathbf{X}_i)$, $a_i(\phi) \cdot V_i^{(0)}(\boldsymbol{\beta}) = \text{Var}(Y_i|A_i = 0, \mathbf{X}_i)$ and $a_i(\phi) \cdot V_i^{(1)}(\boldsymbol{\theta}, \boldsymbol{\beta}) = \text{Var}(Y_i|A_i = 1, \mathbf{X}_i)$. For simplicity, let's assume $a_i(\phi) = a(\phi)$ for all i and $a(\phi)$ is known.

Consider the following estimated average score $\bar{\mathbf{S}}_n$ for the treatment group (A=1) under $\text{H}_0 : \boldsymbol{\theta} = \boldsymbol{\theta}_0$:

$$\bar{\mathbf{S}}_n \coloneqq \frac{1}{n} \mathbf{S}_{n,\boldsymbol{\theta}}^{(1)}(\hat{\boldsymbol{\beta}}_n, \boldsymbol{\theta}_0), \tag{2.7}$$

where $\hat{\boldsymbol{\beta}}_n$ is the *maximum likelihood estimator* of $\boldsymbol{\beta}$ calculated based on data from the control group ($A = 0$). Our proposed SST uses a mixture of asymptotic probability ratios of $\bar{\mathbf{S}}_n$ under local alternative hypothesis to that under the null hypothesis as the test statistic $\tilde{\Lambda}_n^{\pi}$, which is defined as below:

$$\tilde{\Lambda}_n^{\pi} = \int \frac{\boldsymbol{\psi}_{\left(\bar{\mathbf{I}}_n^{(1)}(\hat{\boldsymbol{\beta}}_n)(\boldsymbol{\theta}-\boldsymbol{\theta}_0), \frac{\boldsymbol{\Sigma}_n(\hat{\boldsymbol{\beta}}_n)}{n}\right)}(\bar{\mathbf{S}}_n)}{\boldsymbol{\psi}_{\left(\mathbf{0}, \frac{\boldsymbol{\Sigma}_n(\hat{\boldsymbol{\beta}}_n)}{n}\right)}(\bar{\mathbf{S}}_n)} \pi(\boldsymbol{\theta}) d\boldsymbol{\theta}, \tag{2.8}$$

where

- $\boldsymbol{\psi}_{(\boldsymbol{\mu}, \boldsymbol{\Sigma})}(\cdot)$ denotes the probability density function of a multivariate normal distribution with mean $\boldsymbol{\mu}$ and variance $\boldsymbol{\Sigma}$,
- $\boldsymbol{\Sigma}_n(\boldsymbol{\beta}) = \bar{\mathbf{I}}_n^{(1)}(\boldsymbol{\beta}) + \bar{\mathbf{I}}_n^{(1)}(\boldsymbol{\beta}) \left\{\bar{\mathbf{I}}_n^{(0)}(\boldsymbol{\beta})\right\}^{-1} \bar{\mathbf{I}}_n^{(1)}(\boldsymbol{\beta})$,
- $\bar{\mathbf{I}}_n^{(1)}(\boldsymbol{\beta}) = -\frac{1}{n} \frac{\partial \mathbf{S}_{n,\boldsymbol{\theta}}^{(1)}(\boldsymbol{\beta}, \boldsymbol{\theta}_0)}{\partial \boldsymbol{\beta}} = \frac{1}{n} \sum_{i=1}^{n} \left\{ \frac{\frac{\partial \mu_i^{(1)}(\boldsymbol{\theta}, \boldsymbol{\beta})}{\partial \boldsymbol{\beta}^T} \cdot \frac{\partial \mu_i^{(1)}(\boldsymbol{\theta}, \boldsymbol{\beta})}{\partial \boldsymbol{\theta}}}{a(\phi) \cdot V_i^{(1)}(\boldsymbol{\theta}, \boldsymbol{\beta})} \right\} \Bigg|_{\boldsymbol{\theta} = \boldsymbol{\theta}_0}$,
- $\bar{\mathbf{I}}_n^{(0)}(\boldsymbol{\beta}) = -\frac{1}{n} \frac{\partial \mathbf{S}_{n,\boldsymbol{\beta}}^{(0)}(\boldsymbol{\beta})}{\partial \boldsymbol{\beta}} = \frac{1}{n} \sum_{i=1}^{n} \left\{ \frac{\frac{\partial \mu_i^{(0)}(\boldsymbol{\beta})}{\partial \boldsymbol{\beta}^T} \cdot \frac{\partial \mu_i^{(0)}(\boldsymbol{\beta})}{\partial \boldsymbol{\beta}}}{a(\phi) \cdot V_i^{(0)}(\boldsymbol{\beta})} \right\}$,
- $\pi(\cdot)$ is a mixture distribution of true effects $\boldsymbol{\theta}$ over the parameter space. It is assumed to be positive everywhere. For ease of computation, we often choose $\boldsymbol{\theta} \sim \mathbf{MVN}(\boldsymbol{\theta}_0, \tau^2 \mathbf{I})$, where $\mathbf{I}$ denotes the $(p+1) \times (p+1)$ identity matrix and τ is chosen based on historical data.

Intuitively, a large value of $\tilde{\Lambda}_n^{\pi}$ represents the evidence against H_0 in favor of a mixture of alternatives $\boldsymbol{\theta} \neq \boldsymbol{\theta}_0$, weighted by $\boldsymbol{\theta} \sim \pi(\cdot)$. The decision rule for SST is quite simple and is shown in (2.9). That is, given a significance level α, the test stops and rejects the null hypothesis at the first time that $\tilde{\Lambda}_n^{\pi} \geq \alpha^{-1}$; if no such time exists, it accepts the null hypothesis.

$$T(\alpha) = \inf\{n : \tilde{\Lambda}_n^{\pi} \geq \alpha^{-1}\} \quad \delta(\alpha) = 1\{T(\alpha) < \infty\}. \tag{2.9}$$

The corresponding sequential (always valid) p-value at sample size n, by definition of (2.2), is the reciprocal of the maximum value of $\tilde{\Lambda}_n^{\pi}$ up to n:

$$p_n = \frac{1}{\max_{m \leq n} \tilde{\Lambda}_m^{\pi}}. \tag{2.10}$$

It is obvious to see that the sequential p-value is monotonically nonincreasing in n and $p_{T(\alpha)} = \alpha$.

2.3.2 Validity of sequential score test

The intuition of $\tilde{\Lambda}_n^{\pi}$ being the appropriate test statistic comes from representing the mixture asymptotic probability ratios of $\bar{\mathbf{S}}_n$. In this section, we will give the asymptotic distribution of $\bar{\mathbf{S}}_n$ under the null hypothesis and the local alternative hypothesis, respectively. Meanwhile, we will offer some insights to demonstrate the approximate validity of SST, that is, the type I error is controlled at a large sample size.

The following lemma provides the asymptotic distributions of $\bar{\mathbf{S}}_n$ with proof shown in Appendix A.1.1.

Lemma 2.3.1 *For the generalized linear model in (2.3)-(2.4)(2.6) and $\bar{\mathbf{S}}_n$ in (2.7), define the information matrix for each group as below:*

$$\mathbf{I}^{(0)}(\boldsymbol{\beta}) \coloneqq \mathbb{E}_{(\mathbf{X},Y)}\left\{\bar{\mathbf{I}}_n^{(0)}(\boldsymbol{\beta})\right\} = \mathbb{E}_{(\mathbf{X},Y)}\left\{\frac{\frac{\partial \mu_1^{(0)}(\boldsymbol{\beta})}{\partial \boldsymbol{\beta}^T}\cdot\frac{\partial \mu_1^{(0)}(\boldsymbol{\beta})}{\partial \boldsymbol{\beta}}}{a(\phi)\cdot V_1^{(0)}(\boldsymbol{\beta})}\right\}$$

$$\mathbf{I}^{(1)}(\boldsymbol{\beta}) \coloneqq \mathbb{E}_{(\mathbf{X},Y)}\left\{\bar{\mathbf{I}}_n^{(1)}(\boldsymbol{\beta})\right\} = \mathbb{E}_{(\mathbf{X},Y)}\left\{\frac{\frac{\partial \mu_1^{(1)}(\boldsymbol{\theta},\boldsymbol{\beta})}{\partial \boldsymbol{\beta}^T}\cdot\frac{\partial \mu_1^{(1)}(\boldsymbol{\theta},\boldsymbol{\beta})}{\partial \boldsymbol{\theta}}}{a(\phi)\cdot V_1^{(1)}(\boldsymbol{\theta},\boldsymbol{\beta})}\right\}\Bigg|_{\boldsymbol{\theta}=\boldsymbol{\theta}_0}.$$

Then, under the null hypothesis $H_0 : \boldsymbol{\theta} = \boldsymbol{\theta}_0$,

$$\sqrt{n}\bar{\mathbf{S}}_n \xrightarrow[H_0]{d} \mathbf{MVN}_{p+1}\left(\mathbf{0}, \boldsymbol{\Sigma}(\boldsymbol{\beta})\right),$$

whereas under the local alternative $H_1 : \boldsymbol{\theta} = \boldsymbol{\theta}_0 + \frac{\boldsymbol{\delta}}{\sqrt{n}}$,

$$\sqrt{n}\left\{\bar{\mathbf{S}}_n - \mathbf{I}^{(1)}(\boldsymbol{\beta})(\boldsymbol{\theta}-\boldsymbol{\theta}_0)\right\} \xrightarrow[H_1]{d} \mathbf{MVN}_{p+1}\left(\mathbf{0}, \boldsymbol{\Sigma}(\boldsymbol{\beta})\right),$$

where $\boldsymbol{\Sigma}(\boldsymbol{\beta}) = \mathbf{I}^{(1)}(\boldsymbol{\beta}) + \mathbf{I}^{(1)}(\boldsymbol{\beta})\left\{\mathbf{I}^{(0)}(\boldsymbol{\beta})\right\}^{-1}\mathbf{I}^{(1)}(\boldsymbol{\beta})$.

By Lemma 2.3.1, the asymptotic probability ratio of $\bar{\mathbf{S}}_n$ under the local alternative $\mathrm{H}_1 : \boldsymbol{\theta} = \boldsymbol{\theta}_0 + \boldsymbol{\delta}/\sqrt{n}$ against the null hypothesis $\mathrm{H}_0 : \boldsymbol{\theta} = \boldsymbol{\theta}_0$ can be represented as

$$\lambda_n = \frac{\boldsymbol{\psi}_{\left(\mathbf{I}^{(1)}(\boldsymbol{\beta})(\boldsymbol{\theta}-\boldsymbol{\theta}_0), \frac{\boldsymbol{\Sigma}(\boldsymbol{\beta})}{n}\right)}(\bar{\mathbf{S}}_n)}{\boldsymbol{\psi}_{\left(\mathbf{0}, \frac{\boldsymbol{\Sigma}(\boldsymbol{\beta})}{n}\right)}(\bar{\mathbf{S}}_n)}. \tag{2.11}$$

Different from the likelihood ratio, λ_n is not an exact martingale, but we can show that the approximate martingale property does hold when the sample size n is large enough. See the following remark for mathematical expression; the proof can be found in Appendix A.1.2.

Remark 2.3.1 *For the generalized linear model in (2.3)-(2.4)(2.6) and λ_n defined by (2.11), let $\mathscr{F}_n$*

denote the filtration that contains historical information as below:

$$\mathscr{F}_n = \{(\mathbf{X}_i^{(j)}, Y_i^{(j)}), i = 1, \cdots, n;\ j = 0, 1\}.$$

Then, under the null hypothesis $H_0 : \boldsymbol{\theta} = \boldsymbol{\theta}_0$, $\mathbb{E}(\lambda_{n+1}|\mathscr{F}_n)$ *is approximately equal to* $\lambda_n \cdot \exp\{o_P(1)\}$.

For practical purposes, we usually replace λ_n with its following empirical version

$$\tilde{\lambda}_n = \frac{\boldsymbol{\psi}_{\left(\bar{\mathbf{I}}_n^{(1)}(\hat{\boldsymbol{\beta}}_n)(\boldsymbol{\theta}-\boldsymbol{\theta}_0), \frac{\boldsymbol{\Sigma}_n(\hat{\boldsymbol{\beta}}_n)}{n}\right)}(\bar{\mathbf{S}}_n)}{\boldsymbol{\psi}_{\left(\mathbf{0}, \frac{\boldsymbol{\Sigma}_n(\hat{\boldsymbol{\beta}}_n)}{n}\right)}(\bar{\mathbf{S}}_n)},$$

which is exactly the main term in the definition (2.8) of $\tilde{\Lambda}_n^{\pi}$. The empirical ratio $\tilde{\lambda}_n$ shares the same martingale property as λ_n when the sample size is large enough.

Similar to mSPRT in Section 2.2.2, if we can show that the ratio $\tilde{\lambda}_n$ is a martingale under the null hypothesis $\mathrm{H}_0 : \boldsymbol{\theta} = \boldsymbol{\theta}_0$, the type I error control for SST follows immediately by applying optional stopping theorem and the fact that *a mixture of martingales is also a martingale*. Clearly, as a result of asymptotic distribution and empirical replacement, exact martingale cannot be proved for $\tilde{\lambda}_n$. But with approximate martingale property in Remark 2.3.1, the decision rule (2.9) for SST approximately controls type I error at small α where the large sample size is necessary to reject H_0.

2.4 Multiple Testing

The SST framework can also be applied to multiple testing, where more than one treatment variation is compared against a baseline variation, or more than one metric is of interest between two variations. The main problem in multiple comparisons is that the probability to find at least one statistically significant effect across a set of tests, even when in fact there is nothing going on, increases with the number of comparisons [Hsu96].

In fixed-horizon, *Bonferroni correction* [Mil66] and *Benjamini-Hochberg* [Ben95] are two well-studied methods designed to address this issue. The Bonferroni correction deals with multiple testing by controlling the *family-wise error rate* (FWER): the probability of making at least one false rejection. Although FWER control provides the safest inference, it is too conservative to offer sufficient detection power. Therefore, the Benjamini-Hochberg procedure is proposed to control the *false discovery rate* (FDR): the expected proportion of rejections that are false. Both procedures take as input the vector of p-values for each comparison and produce a set of rejections.

In sequential test, the always valid p-value defined in (2.2) works as the ordinary p-value in fixed horizon testing. It is trivial to show that Bonferroni or Benjamini-Hochberg procedure applied on a collection of sequential p-values controls FWER or FDR (respectively) in the presence of arbitrary continuous monitoring [Joh17]. The corresponding algorithms for sequential multiple comparisons under SST framework can be summarized in proposition 2.4.1 and 2.4.2.

Proposition 2.4.1 *(Bonferroni correction for SST). For arbitrary stopping time T, compute the corresponding sequential p-values $(p_T^i)_{i=1}^m$ by (2.10) for m comparisons. Then reject hypotheses $(1),...(j)$, where j is the maximal such that $p_T^{(j)} \leq \alpha/m$, and $p_T^{(1)},...,p_T^{(m)}$ are the p-values arranged in an increasing order.*

Proposition 2.4.2 *(Benjamini-Hochberg procedure for SST). For arbitrary stopping time T, compute the corresponding sequential p-values $(p_T^i)_{i=1}^m$ by (2.10) for m comparisons. Then reject hypotheses $(1),...(j)$, where j is the maximal such that*

$$p_T^{(j)} \leq \frac{\alpha j}{m \sum_{r=1}^m 1/r}, \tag{2.12}$$

and $p_T^{(1)},...,p_T^{(m)}$ are the p-values arranged in an increasing order.

Note that the term $\sum_{r=1}^m 1/r$ in (2.12) accounts for the fact that the p-values may be correlated [Ben01].

2.5 Experiments

2.5.1 Simulation study

In this section, we compare our SST with the widely-used mSPRT for both A/B tests (two-variations tests) and multiple tests on simulation data generated from combinations of three generalized linear models (2.3)-(2.4)(2.6) and five types of covariates. The significance level $\alpha = 0.05$, null value of the testing parameter $\boldsymbol{\theta}_0 = (0,0)$ (for 2-dimensional covariates) or $(0,0,0)$ (for 3-dimensional covariates), and true nuisance parameter $\boldsymbol{\beta} = (0,1)$ or $(0,1,-1)$ are fixed for all experiments. Each experiment is repeated 1000 times.

Generalized linear models: We choose three generalized linear models to represent responses in different applications. For binary outcomes, such as clicks, conversions, etc., we use logistic regression (Bernoulli distribution). For real-valued responses like revenue, ordinary linear regression (normal distribution) is a good choice. If the responses are nonnegative integers, a Poisson distribution that corresponds to log regression is appropriate. However, mSPRT did not provide the form of test statistics for Poisson distribution, so we only give our SST result for log regression.

Covariates generation: We consider 5 different distributions for 2 or 3-dimensional ($p = 1 \text{or} 2$) covariates. The first dimension is always 1 to indicate the intercept. The other element of 2-dimensional covariates is generated from a normal distribution $\mathrm{N}(0,1)$, uniform distribution $\mathrm{U}[-1,1]$, and Bernoulli distribution $\mathrm{Ber}(0.5)$, respectively. The last two elements of 3-dimensional covariates are generated either from a multivariate normal with mean $\binom{0}{0}$ and variance $\left(\begin{smallmatrix} 1 & 0.5 \\ 0.5 & 1 \end{smallmatrix}\right)$, or a hybrid distribution with one variable from $\mathrm{N}(0,1)$ and the other one from $\mathrm{U}[-1,1]$ independently.

In A/B testing, data are generated in batches with batch size 200 and then are assigned equally to the control group and treatment group. After each batch, we compute the corresponding test statistic and reject the null hypothesis the first time it exceeds some predetermined threshold. We also set a failure time $N = 10000$, which means that we would accept the null hypothesis if the test statistic does not exceed the threshold before the data are accumulated to $N = 10000$ (for each group). We set the true

value of HTE $\boldsymbol{\theta}$ to be three vectors with different scales, including the null value $\boldsymbol{\theta}_0$. For $\boldsymbol{\theta} = \boldsymbol{\theta}_0$, we estimate the type I error by computing the rejection ratio among 1000 repeated experiments. For the other two vectors, we estimate the power in the same way.

It shows that the sequential score test is able to control the type I error (Table 2.1), and achieve higher detection power (Table 2.2) than mSPRT if there exists a heterogeneous treatment effect. We also find that if there exists an individual effect on the response, that is, $\boldsymbol{\beta} \neq \mathbf{0}$, mSPRT may not be able to control the type I error (Table 2.1). That is because the model assumption for mSPRT given in formula (2.5) cannot handle individual baseline effects (i.e. $\boldsymbol{\beta}$ in (2.6)) and possible treatment-covariates interaction effects (i.e. HTE effects described by $\boldsymbol{\theta}$ in formula (2.6)). Therefore, the mSPRT test cannot adjust baseline covariates and may lead to incorrect type I errors for testing HTE. For example, when $\boldsymbol{\beta} \neq \mathbf{0}$ while $\boldsymbol{\theta} = \mathbf{0}$ in (2.6), mSPRT may reject the null hypothesis due to the outcome difference caused by baseline effects, which may lead to inflated type I errors. On the other hand, when $\boldsymbol{\theta} \neq \mathbf{0}$, the mSPRT may fail to detect the HTE (lose power) or need to wait a long time to reject the null hypothesis since treatment effects may be masked by individual heterogeneity.

Table 2.1 Estimated type I error of SST and mSPRT in A/B testing on data with heterogeneous treatment effects

GLM	$\boldsymbol{\theta}$	$\boldsymbol{\beta}$	Covariates	Type I error (SST)	Type I error (mSPRT)
Logistic Regression	(0,0)	(0,1)	N(0,1)	0.017	0.001
			U[-1,1]	0.019	0.004
			Ber(0.5)	0.016	0.005
	(0,0,0)	(0,1,-1)	MVN	0.023	0.003
			N(0,1)+U[-1,1]	0.026	0.002
Linear Regression	(0,0)	(0,1)	N(0,1)	0.001	0.132
			U[-1,1]	0.003	0.026
			Ber(0.5)	0.005	0.021
	(0,0,0)	(0,1,-1)	MVN	< 0.001	0.136
			N(0,1)+U[-1,1]	< 0.001	0.201
Log Regression	(0,0)	(0,1)	N(0,1)	0.006	NA
			U[-1,1]	0.008	
			Ber(0.5)	0.006	
	(0,0,0)	(0,1,-1)	MVN	0.003	
			N(0,1)+U[-1,1]	0.004	

Our proposed test also works with high-dimensional covariates. We conduct an additional simulation with 21 covariates (p=20) under logistic regression. Except for the first dimension (being 1), the last 20 covariates are independently generated from different distributions. Among these 20 covariates, 7 are generated from a normal distribution with variance 1 and different means between -0.3 to 0.3, 8

Table 2.2 Estimated power of SST and mSPRT in A/B testing on data with heterogeneous treatment effects

GLM	$\boldsymbol{\theta}$	$\boldsymbol{\beta}$	Covariates	Power (SST)	Power (mSPRT)
Logistic Regression	(-0.12,0.12)	(0,1)	N(0,1)	0.730	0.356
			U[-1,1]	0.709	0.514
			Ber(0.5)	0.215	0.079
	(-0.15,0.15)	(0,1)	N(0,1)	0.956	0.655
			U[-1,1]	0.938	0.851
			Ber(0.5)	0.436	0.169
	(-0.12,0.12,-0.12)	(0,1,-1)	MVN	0.544	0.384
			N(0,1)+U[-1,1]	0.559	0.287
	(-0.15,0.15,-0.15)	(0,1,-1)	MVN	0.875	0.636
			N(0,1)+U[-1,1]	0.897	0.564
Linear Regression	(-0.05,0.05)	(0,1)	N(0,1)	0.685	0.607
			U[-1,1]	0.419	0.535
			Ber(0.5)	0.053	0.099
	(-0.08,0.08)	(0,1)	N(0,1)	1	0.943
			U[-1,1]	0.979	0.960
			Ber(0.5)	0.413	0.323
	(-0.05,0.05,-0.05)	(0,1,-1)	MVN	0.400	0.607
			N(0,1)+U[-1,1]	0.602	0.653
	(-0.08,0.08,-0.08)	(0,1,-1)	MVN	0.994	0.955
			N(0,1)+U[-1,1]	0.999	0.943
Log Regression	(-0.05,0.05)	(0,1)	N(0,1)	0.499	NA
			U[-1,1]	0.132	
			Ber(0.5)	0.026	
	(-0.08,0.08)	(0,1)	N(0,1)	0.996	
			U[-1,1]	0.758	
			Ber(0.5)	0.282	
	(-0.05,0.05,-0.05)	(0,1,-1)	MVN	0.242	
			N(0,1)+U[-1,1]	0.726	
	(-0.08,0.08,-0.08)	(0,1,-1)	MVN	0.971	
			N(0,1)+U[-1,1]	1	

covariates are from a uniform distribution with mean 0 and upper limit between 0.3 to 1, and the last 5 covariates are generated from a binomial distribution with the success probability between 0.1 to 0.5. The individual baseline effect $\boldsymbol{\beta}$ has two non-zero components. The simulation result shows that our SST still can control type I error under the null (type I error is 0.025, which is less than the significance level α), and has reasonable power (i.e. when HTE effects $\boldsymbol{\theta}$ has three non-zero components with the value of 0.2, the power is 0.731; and when $\boldsymbol{\theta}$ has three non-zero components with the value of 0.3, the power is 1).

In multiple testing, we consider $m = 64$ hypotheses, among which 3/4 are true null hypotheses ($\boldsymbol{\theta} = (0,0)$ or $(0,0,0)$) and the remaining 1/4 true alternatives are equally placed at $\boldsymbol{\theta} = (-B, B)$ or

$(-B, B, -B)$, where $B = 0.1, 0.2, 0.3, 0.4$, respectively. For each comparison, we wait until the data are accumulated to $N = 10000$ (for each group) and compute the sequential p-value p_N according to (2.10). After applying the Benjamini-Hochberg procedure, we get the rejections from which we can estimate the *false discovery rate* (FDR) and *true positive rate* (TPR) (i.e. *recall*). The true positive rate, defined as the proportion of correct rejections in true alternatives, is a metric for detection power in multiple testing. Same as A/B testing, the results in Table 2.3 show that SST applied on multiple testing achieves higher true positive rates than mSPRT while maintaining the false discovery rate in control.

Table 2.3 Estimated false discovery rate and true positive rate of SST and mSPRT in multiple testing on data with heterogeneous treatment effects

GLM	Covariates	FDR (SST)	FDR (mSPRT)	TPR (SST)	TPR (mSPRT)
Logistic Regression	N(0,1)	0.0119	0.0008	0.8038	0.7191
	U[-1,1]	0.0059	0.0009	0.7957	0.7662
	Ber(0.5)	0.0067	0.0009	0.6501	0.4761
	MVN	0.0114	0.0009	0.7664	0.7171
	N(0,1)+U[-1,1]	0.0148	0.0007	0.7775	0.6944
Linear Regression	N(0,1)	0.0004	0.1983	1	0.3787
	U[-1,1]	0.0009	0.0687	0.9994	0.2868
	Ber(0.5)	0.0011	0.3332	0.8725	0.0504
	MVN	0.0024	0.2013	0.9999	0.3748
	N(0,1)+U[-1,1]	0.0005	0.2708	1	0.4092
Log Regression	N(0,1)	0.0011	NA	1	NA
	U[-1,1]	0.0025		0.9750	
	Ber(0.5)	0.0019		0.8396	
	MVN	0.0010		0.9997	
	N(0,1)+U[-1,1]	0.0011		1	

2.5.2 Real data application

We also compare SST with mSPRT on the Yahoo dataset which contains user click events on articles over 10 days. Each event has a timestamp, a unique article id, a binary click indicator, and five user features which are between 0 and 1 and sum to 1 for each user (we only use the last four features). We treat each article as different treatment variations, click actions as the binary responses. Our goal is to test if there is any article effect on user click behaviors with (SST) or without (mSPRT) accounting for the user features.

We first conduct an A/A test to show the validity of SST on click events with the most popular article (id=109510) on the date May 1st, 2009, by randomly assigning fake treatment indicators to them. Then we conduct an A/B test on events with the two most popular articles (id=109510 and 109520) on the

date May 1st, 2009. With every 200 events (from both articles) coming in a time sequence, we compute the corresponding test statistics. As soon as the statistics exceed the predetermined critical value $(1/\alpha)$, we stop and reject the null hypothesis. If all the data are used up, we accept the null hypothesis. The experiment shows that both SST and mSPRT accept the null hypothesis for the A/A test, indicating type I errors are well controlled for both tests under the considered hypotheses. For the A/B test, SST needs $n = 19600$ events to get a rejection conclusion while mSPRT needs $n = 67600$. It means that we are able to discover the difference early by accounting for the covariates. We also provide estimated HTE ($\boldsymbol{\theta}$ in (2.6)) and ATE (θ in (2.5)) by fitting logistic regression.

Table 2.4 Estimated heterogeneous treatment effect and average treatment effect in the Yahoo dataset

Control article id	Treatment article id	HTE ($\boldsymbol{\theta}$)	ATE (θ)
109510	109520	(-0.401, -0.091, -0.068, 0.661, -0.178)	-0.179

For multiple testing, we choose 10 articles and do pairwise comparisons. Hence, there are $m = 45$ comparisons in total. We compute p_T for each pair with $T = 20000$ from each article and then apply the Benjamini-Hochberg procedure. Among the 45 pair comparisons, we reject 43 with SST and 23 with mSPRT.

2.6 Conclusions

We propose a new framework of online testing based on the probability ratio of a score function. It is able to test a multi-dimensional heterogeneous treatment effect while accounting for the unknown individual effect. The asymptotic normality of the score function guarantees an explicit form for the test statistic, greatly improving the computation efficiency. We provide an online p-value for SST and extend SST to online multiple testing. We validate our testing procedure by both theoretical proof and empirical results. We also compare SST with a state-of-art online test named mSPRT on simulated data and real data. The results show that our proposed test controls type I error at any time, has higher detection power, and allows quick inference on online A/B testing.

There is still some interesting work we may do in the future. The decision rule of our test implies that we can only get rejection conclusions unless we wait essentially indefinitely, which is impossible in practice. This necessitates truncating SST at a failure time and admitting an inclusive result if we ever reach it, which may diminish the power more or less. How to choose the failure time to balance waiting time and power remains a problem.

CHAPTER

3

ONLINE TESTING OF SUBGROUP TREATMENT EFFECTS BASED ON VALUE DIFFERENCE

3.1 Introduction

Online A/B testing, as a kind of randomized control experiment, is widely used in the high-tech industry to assess the value of ideas in a scientific manner [Koh09]. It randomly exposes users to one of the two variants: *control* (A), the currently-used version, or *treatment* (B), a new version being evaluated, and collects the metric of interest, such as conversion rate, revenue, etc. Then, a null hypothesis statistical test is performed to evaluate whether there is a statistically significant difference between the two variants on the metric of interest. This scientific design helps to control for the external variations and thus establish the causality between the variants and the outcome. However, the current A/B testing has its limitations in terms of framework and model assumptions.

First of all, most A/B tests employ a *fixed-horizon* framework, whose validity requires that the sample size should be fixed and determined before the experiment starts. However, experimenters, driven by a fast-paced product evolution in practice, often "peek" the experiment and hope to find the significance as quickly as possible to avoid large (i) time cost: an A/B test may take prohibitively long time to collect the predetermined size of samples; and (ii) opportunity cost: the users who have been assigned to a suboptimal variant will be stuck in a bad experience for a long time [Ju19]. The behaviors of continuously monitoring and concluding the experiment prematurely will be favorably biased towards

getting significant results and lead to very high false positive probabilities, well in excess of the nominal significance level α [Sim11; Goo14]. Another limitation of A/B tests is that they assume homogeneous treatment effects among the population and mainly focus on testing the average treatment effect. However, it is common that treatment effects vary across subpopulations. Testing the subgroup treatment effects will help decision-makers distinguish the subpopulation that may benefit from a particular treatment from those who may not, and thereby guide companies' marketing strategies in promoting new products.

The first problem can be addressed by applying the sequential testing framework. *Sequential testing*, in contrast to the classic fixed-horizon test, is a statistical testing procedure that continuously checks for significance at every new sample and stops the test as soon as a significant result is detected while controlling the type I error at any time. It generally gives a significant decrease in the required sample size compared to the fixed-horizon test with the same type I error and type II error control [Wal45], and thus is able to end an experiment much earlier. This field was first introduced by Wald [Wal45], who proposed the *sequential probability ratio test* (SPRT) for simple hypotheses using *likelihood ratio* as the test statistics, and then was extended to composite hypotheses by much following literature [Sch62; Cox63; Arm69; Rob70; Lai88]. A thorough review was given in Lai [Lai01]. However, the advantage of sequential testing in online A/B testing has not been recognized until recently Johari et al. [Joh15] brought mSPRT, a variant of SPRT to A/B tests.

The second problem shows a demand for a test on subgroup treatment effects. Although sequential testing is rapidly developing in online A/B tests, little work focuses on subgroup treatment effect testing. Yu et al. [Yu20] proposed a *sequential score test* (SST) based on the score statistics under a *generalized linear model*, which aims to test if there is any difference between treatment and control groups among any subjects. However, this test is based on a restrictive parametric assumption on treatment-covariates interaction and can't be used to test the subgroup treatment effects.

In this chapter, we consider a flexible model and propose a sequential test for *SUBgroup Treatment effects based on vaLuE difference* (SUBTLE), which aims to test if some group of the population would benefit from the investigative treatment. Our method does not require to specify any parametric form of covariate-specific treatment effects. If the null hypothesis is rejected, a beneficial subgroup can be easily obtained based on the estimated optimal treatment rule.

The remainder of this chapter is structured as follows. In Section 3.2, we review the idea of the mSPRT and SST, and discuss how they are related to our test. Then in Section 3.3, we introduce our proposed method SUBTLE and provide the theoretical guarantee for its validity. We conduct simulations in Section 3.4 and real data experiments in Section 3.5 to demonstrate the validity, detection power, robustness, and efficiency of our proposed test. Finally, in Section 3.6, we conclude the chapter and present future directions.

3.2 Related Work

3.2.1 Mixture sequential probability ratio test

The *mixture sequential probability ratio test* (mSPRT) [Rob70] supposes that the independent and identically distributed (i.i.d.) random variables $Y_1, Y_2, \cdots$ have a probability density function $f_\theta(x)$ induced by parameter θ, and aims to test

$$\mathrm{H}_0 : \theta = \theta_0 \quad vs. \quad \mathrm{H}_1 : \theta \neq \theta_0.$$

Its test statistics Λ_n^π at sample size n is a mixture of likelihood ratios weighted by a mixture density $\pi(\cdot)$ over the parameter space Θ:

$$\Lambda_n^\pi = \int_\Theta \left\{ \prod_{i=1}^{n} \frac{f_\theta(Y_i)}{f_{\theta_0}(Y_i)} \right\} \pi(\theta) d\theta.$$

The mSPRT stops the sampling at stage

$$N = \inf\{n \geq 1 : \Lambda_n^\pi \geq 1/\alpha\} \tag{3.1}$$

and rejects the null hypothesis H_0 in favor of H_1. If no such time exists, it continues the sampling indefinitely and accepts the H_0. Since the likelihood ratio under H_0 is a nonnegative martingale with the initial value equal to 1, and so is the mixture of such likelihood ratios Λ_n^π, the type I error of mSPRT can be proved to be always controlled at α by an application of *Markov's inequality* and *optional stopping theorem*: $P_{H_0}(\Lambda_n^\pi \geq \alpha^{-1}) \leq \mathbb{E}_{H_0}(\Lambda_n^\pi)/\alpha^{-1} = \mathbb{E}_{H_0}(\Lambda_0^\pi)/\alpha^{-1} = \alpha$. Besides, mSPRT is a test of power one [Rob74], which means that any small deviation from θ_0 can be detected as long as waiting long enough. It is also shown that mSPRT is almost optimal for data from an exponential family of distributions, with respect to the expected stopping time [Pol78].

The mSPRT was brought to A/B tests by Johari et al. [Joh15; Joh17], who assume that the observations in control and treatment groups arrive in pairs $(Y_i^{(0)}, Y_i^{(1)})$, $i = 1, 2, \cdots$. They restricted their data model to the two most common cases in practice: normal distribution and Bernoulli distribution, with μ_A and μ_B denoting the mean for the control and treatment groups, respectively. They test the hypotheses as below:

$$\mathrm{H}_0 : \theta := \mu_\mathrm{B} - \mu_\mathrm{A} = 0 \quad vs. \quad \mathrm{H}_1 : \theta \neq 0,$$

by directly applying mSPRT to the distribution of the differences $Z_i = Y_i^{(1)} - Y_i^{(0)}$ (normal), or the joint distribution of data pairs $(Y_i^{(0)}, Y_i^{(1)})$ (Bernoulli), $i = 1, 2, \cdots$. After making some approximations to the likelihood ratio and choosing a normal mixture density $\pi(\theta) \sim \mathrm{N}(0, \tau^2)$, the test statistic Λ_n^π is able to have a closed form for both normal and Bernoulli observations.

However, the mSPRT does not work well in testing heterogeneous treatment effects due to the complexity of the likelihood induced by individual covariates. Specifically, a conjugate prior $\pi(\cdot)$ for the likelihood ratio may not exist anymore, making the computation for the test statistic challenging. The unknown baseline covariates effect also increases the difficulty in constructing and approximating the

likelihood ratio [Yu20].

3.2.2 Sequential score test

The sequential score test (SST) [Yu20] assumes a *generalized linear model* with a link function $g(\cdot)$ for the outcome Y:

$$g(\mathbb{E}(Y|A,\mathbf{X})) = \boldsymbol{\mu}^T\mathbf{X} + (\boldsymbol{\theta}^T\mathbf{X})A, \tag{3.2}$$

where A and $\mathbf{X}$ respectively denote the binary treatment indicator and user covariates vector. It tests the multi-dimensional treatment-covariates interaction effect:

$$\mathrm{H}_0 : \boldsymbol{\theta} = \mathbf{0} \quad vs. \quad \mathrm{H}_1 : \boldsymbol{\theta} \neq \mathbf{0}, \tag{3.3}$$

while accounting for the linear baseline covariates effect $\boldsymbol{\mu}^T\mathbf{X}$. For the test statistics Λ_n^π, instead of using a mixture of likelihood ratios as mSPRT, SST employs a mixture of asymptotic probability ratios of the score statistics. Since the probability ratio has the same martingale structure as the likelihood ratio, the type I error can still be controlled with the same decision rule as mSPRT (3.1). The asymptotic normality of the score statistics also guarantees a closed form of Λ_n^π with a multivariate normal mixture density $\pi(\cdot)$. However, the considered parametric model (3.2) can only be used to test if there is a linear covariate-treatment interaction effect and may fail to detect the existence of a subgroup with enhanced treatment effects. In addition, the subgroup estimated based on the index $\boldsymbol{\theta}^T\mathbf{X}$ may be biased if the assumed linear model (3.2) is misspecified. Therefore in this chapter, we propose a subgroup treatment effect test, which is able to test the existence of a beneficial subgroup and does not require specifying the form of treatment effects.

3.3 Subgroup Treatment Effects Test Based on Value Difference

3.3.1 Problem setup

Suppose we have i.i.d. data $\mathbf{O}_i = \{Y_i, A_i, \mathbf{X}_i\}$, $i = 1, 2, \cdots$, where Y_i, A_i, $\mathbf{X}_i$ respectively denote the observed outcome, binary treatment indicator, and user covariates vector. Here, we consider a flexible generalized linear model:

$$g(\mathbb{E}(Y_i|A_i,\mathbf{X}_i)) = \mu(\mathbf{X}_i) + \theta(\mathbf{X}_i)A_i, \tag{3.4}$$

where baseline covariates effect $\mu(\cdot)$ and treatment-covariates interaction effect $\theta(\cdot)$ are completely unspecified functions, and $g(\cdot)$ is a prespecified link function. For example, we use the identity link $g(\mu) = \mu$ for normal responses and the logit link $g(\mu) = \log\{\mu/(1-\mu)\}$ for binary outcomes.

Assuming Y is coded such that larger values indicate a better outcome, we consider the following test of subgroup treatment effects:

$$\begin{aligned} &\mathrm{H}_0 : \forall \mathbf{x} \in \mathscr{X},\ \theta(\mathbf{x}) \leq 0 \quad vs. \\ &\mathrm{H}_1 : \exists \mathscr{X}_0 \subset \mathscr{X} \text{ such that } \theta(\mathbf{x}) > 0 \text{ for all } \mathbf{x} \in \mathscr{X}_0, \end{aligned} \tag{3.5}$$

where $\mathscr{X}_0$ is the beneficial subgroup with $P(\mathbf{X} \in \mathscr{X}_0) > 0$. Note that the above subgroup test is very different from the covariate-treatment interaction test considered in (3.3) and is much more challenging due to several aspects. First, both $\mu(\cdot)$ and $\theta(\cdot)$ are nonparametric and need to be estimated. Second, the considered hypotheses (3.5) are moment inequalities that are nonstandard. Third, it allows a nonregular setting, i.e. $P(\theta(\mathbf{X}) = 0) > 0$, which makes associated inference difficult. Here, we propose a test based on the value difference between the optimal treatment rule and a fixed treatment rule.

Let $Y^*(a)$ denote the potential outcome had the subject received treatment a ($a = 0, 1$), and $V(d) = \mathbb{E}_{\{Y^*(a),\mathbf{X}\}}\{Y^*(d(\mathbf{X}))\}$ denote a value function for a fixed treatment rule $d(\mathbf{X})$ which maps the information in $\mathbf{X}$ to treatment $\{0, 1\}$, where the subscript represents the expectation is taken with respect to the joint distribution of $\{Y^*(a), \mathbf{X}\}$. Consider the value difference $\Delta = V(d^{\text{opt}}) - V(0)$ between the optimal treatment rule $d^{\text{opt}} = 1\{\theta(\mathbf{X}) > 0\}$ and the treatment rule that assigns control to everyone $d = 0$, where $1\{\cdot\}$ is an indicator function. If the null hypothesis is true, no one would benefit from the treatment and the optimal treatment rule assigns everyone to control, and therefore the value difference is zero. However, if the alternative hypothesis is true, some people would have higher outcomes being assigned to treatment and thus the value difference is positive. In this way, the testing hypotheses (3.5) can be equivalently transformed into the following pair:

$$\mathrm{H}_0 : \Delta = 0 \quad vs. \quad \mathrm{H}_1 : \Delta > 0. \tag{3.6}$$

We make the following standard causal inference assumptions: (i) consistency, which states that the observed outcome is equal to the potential outcome under the actual treatment received, i.e. $Y = Y^*(1)1(A = 1) + Y^*(0)1(A = 0)$; (ii) no unmeasured confounders, i.e. $Y^*(a) \perp\!\!\!\perp A|\mathbf{X}$ for $a = 0, 1$, which means that the potential outcome is independent of treatment given covariates; (iii) positivity, i.e. $P(A = a|\mathbf{X} = \mathbf{x}) > 0$ for $a = 0, 1$ and all $\mathbf{x} \in \mathscr{X}$ such that $P(\mathbf{X} = \mathbf{x}) > 0$. Under these assumptions, it can be shown that

$$\begin{aligned} V(d) &= \mathbb{E}_{\mathbf{X}}\left[\mathbb{E}\left\{Y^*(d(\mathbf{X}))|\mathbf{X}\right\}\right] \\ &= \mathbb{E}_{\mathbf{X}}[\mathbb{E}\{Y|A = d(\mathbf{X}), \mathbf{X}\}]. \end{aligned}$$

3.3.2 Algorithm and implementation

We estimate the value function of a given treatment rule d by the *augmented inverse probability weighted* (AIPW) estimator [Rob94b; Zha12]:

$$\hat{V}_{\text{AIPW}}(d) = \frac{1}{n}\sum_{i=1}^{n}\left[\frac{Y_i \cdot 1(A_i = d)}{p_{A_i}(\mathbf{X}_i)} - \left\{\frac{1(A_i = d)}{p_{A_i}(\mathbf{X}_i)} - 1\right\} \times \mathbb{E}(Y_i|A_i = d, \mathbf{X}_i)\right],$$

where $p_A(\mathbf{X}) = A \cdot p(\mathbf{X}) + (1 - A) \cdot (1 - p(\mathbf{X}))$ and $p(\mathbf{X}) = P(A = 1|\mathbf{X})$ is the propensity score. This estimator is unbiased, i.e. $\mathbb{E}_{(Y,A,\mathbf{X})}\{\hat{V}_{\text{AIPW}}(d)\} = V(d)$. Moreover, the AIPW estimator has double robustness, that is, the estimator remains consistent if either the estimator of $\mathbb{E}(Y|A = d, \mathbf{X})$ or the estimator of the propensity score $p(\mathbf{X})$ is consistent, which gives much flexibility. Then the value difference Δ is unbiasedly estimated

by

$$D(\mathbf{O}_i;\mu,\theta,p) \coloneqq \left[\frac{Y_i\cdot 1\left(A_i=1\{\theta(\mathbf{X}_i)>0\}\right)}{p_{A_i}(\mathbf{X}_i)} - \left\{\frac{1\left(A_i=1\{\theta(\mathbf{X}_i)>0\}\right)}{p_{A_i}(\mathbf{X}_i)}-1\right\} \times g^{-1}\left(\mu(\mathbf{X}_i)+\theta(\mathbf{X}_i)1\{\theta(\mathbf{X}_i)>0\}\right)\right] - \left[\frac{Y_i\cdot 1(A_i=0)}{1-p(\mathbf{X}_i)} - \left\{\frac{1(A_i=0)}{1-p(\mathbf{X}_i)}-1\right\}\times g^{-1}\left(\mu(\mathbf{X}_i)\right)\right]$$

under our assumed model (3.4), where $g^{-1}(\cdot)$ is the inverse of the link function. That is,

$$\mathbb{E}_{(Y,A,\mathbf{X})}\left\{D(\mathbf{O}_i;\mu,\theta,p)\right\} = \Delta. \tag{3.7}$$

Since $\mu(\cdot)$, $\theta(\cdot)$, and $p(\cdot)$ are usually unknown, we let data come in batches and estimate them based on previous batches of data.

Like SST, the key idea behind SUBTLE is to construct a statistic that has an (asymptotic) normal distribution with different means under the null hypothesis and alternative hypothesis (3.6), and then build the test statistics as a mixture of (asymptotic) probability ratios of it. Define $\hat{\Delta}_k$ as below:

$$\hat{\Delta}_k \coloneqq \frac{\sum_{j=1}^{k}\hat{\sigma}_j^{-1}\bar{D}_j(\mathscr{C}_j;\overline{\mathscr{C}}_{j-1})}{\sum_{j=1}^{k}\hat{\sigma}_j^{-1}}. \tag{3.10}$$

Note that $\hat{\Delta}_k$ is an asymptotic unbiased estimator for Δ by (3.7). R_k (3.8) is a multiplier of $\hat{\Delta}_k$ and thus has the mean related to the value difference Δ. In section 3.3.3 we will show that R_k has an asymptotic normal distribution with the same variance but different means under the null and local alternative hypotheses so that our test statistics Λ_k^{π} (3.9) is a mixture of asymptotic probability ratios of R_k. As we discussed in Section 3.2, the probability ratio has the same martingale structure as the likelihood ratio, so SUBTLE can be shown to control type I error with the decision rule (3.1). Algorithm 1 shows our complete testing procedures.

In step (ii) of Algorithm 1, we estimate $\mu(\cdot)$ and $\theta(\cdot)$ by respectively building a random forest on control observations and treatment observations in previous batches. The propensity score $p(\cdot)$ is estimated by computing the proportion of treatment observations ($A=1$) in previous batches. In step (iv) we estimate σ_k with $\hat{\sigma}_k=\sqrt{s_k^2/m}$, where s_k^2 is the sample variance of $D(\mathbf{O}_i;\hat{\mu}_{k-1},\hat{\theta}_{k-1},\hat{p}_{k-1})$, $\forall \mathbf{O}_i \in \overline{\mathscr{C}}_{k-1}$. Since the value difference is always nonnegative, in step (v) we choose a truncated normal $\pi(\Delta)=2/\sqrt{2\pi\tau^2}\cdot\exp\left\{-\Delta^2/2\tau^2\right\}\cdot 1\{\Delta>0\}$ as the mixture density, where τ^2 is estimated based on historical data. The normality of the mixture density guarantees a closed form for our test statistic:

$$\Lambda_k^{\pi} = 2\left\{\frac{k}{k+\left(\tau\cdot\sum_{j=1}^{k}\hat{\sigma}_j^{-1}\right)^2}\right\}^{1/2} \times \exp\left[\frac{\left(\tau\cdot\sum_{j=1}^{k}\hat{\sigma}_j^{-1}\cdot R_k\right)^2}{2\left\{\left(\tau\cdot\sum_{j=1}^{k}\hat{\sigma}_j^{-1}\right)^2+k\right\}}\right]\times\{1-F(0)\},$$

where $F(\cdot)$ is the cumulative distribution function of a normal distribution with mean

Algorithm 1: SUBTLE

1. Initialize batch index $k=0$, test statistic $\Lambda_k^\pi=0$. Choose a significance level $0<\alpha<1$, a batch size m, an initial batch size l, and a failure time M.
2. Sample l observations to formulate initial batch $\mathscr{C}_0$.

while *True* **do**

(i) k=k+1.

(ii) Let $\overline{\mathscr{C}}_{k-1}=\cup_{j=0}^{k-1}\mathscr{C}_j$. Estimate $\mu(\cdot)$, $\theta(\cdot)$, and $p(\cdot)$ based on data in $\overline{\mathscr{C}}_{k-1}$ to get $\hat{\mu}_{k-1}$, $\hat{\theta}_{k-1}$, and $\hat{p}_{k-1}$.

(iii) Sample another m observations to formulate batch $\mathscr{C}_k$. For each $\mathbf{O}_i\in\mathscr{C}_k$, calculate $D(\mathbf{O}_i;\hat{\mu}_{k-1},\hat{\theta}_{k-1},\hat{p}_{k-1})$. Let

$$\bar{D}_k(\mathscr{C}_k;\overline{\mathscr{C}}_{k-1})=\frac{1}{m}\sum_{\mathbf{O}_i\in\mathscr{C}_k}D(\mathbf{O}_i;\hat{\mu}_{k-1},\hat{\theta}_{k-1},\hat{p}_{k-1}).$$

(iv) Estimate the conditional standard deviation $\sigma_k=sd\left(\bar{D}_k(\mathscr{C}_k;\overline{\mathscr{C}}_{k-1})|\overline{\mathscr{C}}_{k-1}\right)$ based on data in $\overline{\mathscr{C}}_{k-1}$ and denote it as $\hat{\sigma}_k$.

(v) Calculate

$$R_k=\frac{1}{\sqrt{k}}\sum_{j=1}^{k}\hat{\sigma}_j^{-1}\bar{D}_j(\mathscr{C}_j;\overline{\mathscr{C}}_{j-1}) \tag{3.8}$$

and

$$\Lambda_k^\pi=\int\frac{\psi_{\left(\frac{1}{\sqrt{k}}\left(\sum_{j=1}^{k}\hat{\sigma}_j^{-1}\right)\Delta,\,1\right)}(R_k)}{\psi_{(0,1)}(R_k)}\pi(\Delta)d\Delta, \tag{3.9}$$

where $\psi_{(\mu,\sigma^2)}(\cdot)$ denotes the probability density function of a normal distribution with mean μ and variance σ^2.

if $\Lambda_k^\pi>1/\alpha$ *or* $k\times m+l>M$ **then**

break

end

end

if $\Lambda_k^\pi>1/\alpha$ **then**

Reject H_0. Estimate $\theta(\cdot)$ using all the data up to now and identify a subgroup $1\{\hat{\theta}(\mathbf{X})>0\}$.

else

Accept H_0.

end

$(\sqrt{k}\sum_{j=1}^{k}\hat{\sigma}_j^{-1}R_k)/\{(\sum_{j=1}^{k}\hat{\sigma}_j^{-1})^2+k\}$ and variance $k\tau^2/\{\tau^2(\sum_{j=1}^{k}\hat{\sigma}_j^{-1})^2+k\}$. In theory, the choice of mixture density variance τ^2 will not have any effect on the type I error control. Johari et al. [Joh15] proved that an optimal τ^2 in terms of stopping time is the prior variance times a correction for truncating, so we suggest estimating τ^2 based on historical data. The simulation result in Appendix B.2 shows considerable robustness in choosing τ^2. In the last step, we add a failure time M to the decision rule to terminate the test externally and accept the null hypothesis if we ever reach it. If the null hypothesis is rejected, we can employ random forests to estimate $\theta(\cdot)$ based on all the data up to the time that the experiment ends. Then the estimated treatment effect $\hat{\theta}(\mathbf{x})$ naturally gives the beneficial subgroup $\mathscr{X}_0=\{\mathbf{x}:\hat{\theta}(\mathbf{x})>0\}$.

3.3.3 Validity

In this section, we will show that our proposed test SUBTLE is able to control type I error at any time, that is, $P_{H_0}(\Lambda_k^\pi > 1/\alpha) < \alpha$ for any $k \in \mathbb{N}^+$. Theorem 3.3.1 gives the respective asymptotic distributions of R_k under the null and local alternative hypotheses, which demonstrates that the test statistics Λ_k^π is a mixture of asymptotic probability ratios weighted by $\pi(\cdot)$. Proposition 3.3.1 shows that this asymptotic probability ratio has a martingale structure under H_0 when the sample size is large enough. Combining these two results with the demonstration in Section 3.2.1, we can conclude that the type I error of SUBTLE is always controlled at α.

The upcoming theorem relies on the following conditions:

(C1) The number of batches k diverges to infinity as sample size n diverges to infinity.

(C2) Lindeberg-like condition: for all $\epsilon > 0$

$$\frac{1}{k}\sum_{j=1}^{k}\mathbb{E}\left[\left\{\frac{\bar{D}_j(\mathscr{C}_j;\overline{\mathscr{C}}_{j-1})}{\hat{\sigma}_j}\right\}^2 \times 1\left\{\frac{|\bar{D}_j(\mathscr{C}_j;\overline{\mathscr{C}}_{j-1})|}{\hat{\sigma}_j} > \epsilon\sqrt{k}\right\}\middle|\overline{\mathscr{C}}_{j-1}\right] = o_P(1)$$

(C3) $\frac{1}{k}\sum_{j=1}^{k}\frac{\sigma_j^2}{\hat{\sigma}_j^2} \xrightarrow{P} 1.$

(C4)

$$\frac{1}{k}\sum_{j=1}^{k}\hat{\sigma}_j^{-1}\left[\mathbb{E}\left\{\bar{D}_j(\mathscr{C}_j;\overline{\mathscr{C}}_{j-1})|\overline{\mathscr{C}}_{j-1}\right\} - \mathbb{E}\left\{\bar{D}_j(\mathscr{C}_j;\hat{d}_{j-1}^{opt},\mu,\theta,p)|\overline{\mathscr{C}}_{j-1}\right\}\right] = o_P(k^{-1/2}),$$

where $\bar{D}_j(\mathscr{C}_j;\hat{d}_{j-1}^{opt},\mu,\theta,p)$ is $\bar{D}_j(\mathscr{C}_j;\overline{\mathscr{C}}_{j-1})$ with $\hat{\mu}_{j-1},\hat{\theta}_{j-1},\hat{p}_{j-1}$ replaced by the true value μ,θ,p, but $\hat{d}_{j-1}^{opt} = 1\{\hat{\theta}_{j-1}(\mathbf{X}) > 0\}$ unchanged.

(C5) $\frac{1}{k}\sum_{j=1}^{k}\hat{\sigma}_j^{-1}\left[\mathbb{E}\left\{\bar{D}_j(\mathscr{C}_j;\hat{d}_{j-1}^{opt},\mu,\theta,p)|\overline{\mathscr{C}}_{j-1}\right\} - \Delta\right] = o_P(k^{-1/2}).$

Similar conditions are also used in the literature for studying the properties of doubly robust estimators. Their appropriateness was discussed in Section 7 of Luedtke & Van Der Laan [Lue16]. Both (C4) and (C5) rely on the convergence rate of the estimator of μ,θ,p. Wager & Athey [Wag18] showed that under certain constraints on the subsampling rate, random forest predictions converge at the rate $n^{s-1/2}$, where s is chosen to satisfy some conditions. We assume that under this rate, (C4) and (C5) hold.

Theorem 3.3.1 *For $\hat{\Delta}_k$ defined in (3.10), under conditions (C1)-(C5),*

$$\frac{1}{\sqrt{k}}\left(\sum_{j=1}^{k}\hat{\sigma}_j^{-1}\right)\left(\hat{\Delta}_k - \Delta\right) \xrightarrow{d} N(0,1) \quad \text{as } k \to \infty,$$

where $\xrightarrow{d}$ represents convergence in distribution. In particular, as $k \to \infty$, $R_k \xrightarrow[H_0]{d} N(0,1)$ under the

null hypothesis $\Delta = 0$*, while* $R_k - k^{-1/2}\left(\sum_{j=1}^{k} \hat{\sigma}_j^{-1}\right)\Delta \xrightarrow[H_1]{d} N(0,1)$ *under the local alternative* $\Delta = \delta/\sqrt{k}$*, where* $\delta > 0$ *is fixed.*

Proposition 3.3.1 *Let*

$$\lambda_k = \frac{\psi_{\left(\frac{1}{\sqrt{k}}\left(\sum_{j=1}^{k} \hat{\sigma}_j^{-1}\right)\Delta, 1\right)}(R_k)}{\psi_{(0,1)}(R_k)}$$

and $\mathscr{F}_k$ *denote a filtration that contains all the historical information in the first* $(k+1)$ *batches* $\overline{\mathscr{C}}_k$*. Then under the null hypothesis* $H_0 : \Delta = 0$*,* $\mathbb{E}(\lambda_{k+1}|\mathscr{F}_k)$ *is approximately equal to* $\lambda_k \cdot \exp\{o_P(1)\}$*.*

The proofs of the above results are given in Appendix B.1.

3.4 Simulated Experiments

In this section, we evaluate the test SUBTLE on three metrics: type I error, power, and sample size. We first compare SUBTLE with SST in terms of type I error and power under five models in Section 3.4.1. Then in Section 3.4.2, we present the impact of noise covariates on their powers. Finally, in Section 3.4.3, we compare the stopping time of SUBTLE to the required sample size of a fixed-horizon value difference test. The significance level $\alpha = 0.05$, initial batch size $l = 300$, failure time $M = 2300$, and variance of mixture distribution $\tau^2 = 1$ are fixed for both SUBTLE and SST in all simulation settings unless otherwise specified.

3.4.1 Type I error & power

We consider five data generation models in the form of (3.4) with logistic link $g(\cdot)$. Data are generated in batches with batch size $m = 20$ and are randomly assigned to two groups with a fixed propensity score $p(\mathbf{X}) = 0.5$. Each experiment is repeated 1000 times to estimate the type I error and power. For the first four models, we consider

Five covariates :

$$X_1 \overset{iid}{\sim} \text{Ber}(0.5),\ X_2 \overset{iid}{\sim} \text{U}[-1,1],$$
$$X_3, X_4, X_5 \overset{iid}{\sim} \text{N}(0,1).$$

Two baseline effects :

$$\mu_1(\mathbf{X}) = -2 - X_1 + X_3^2,$$
$$\mu_2(\mathbf{X}) = -1.3 + X_1 + 0.5X_2 - X_3^2.$$

Two treatment-covariates interaction effects :

$$\theta_1(\mathbf{X}) = c \cdot 1\{X_1 + 2X_3 > 0\},$$

$$\theta_2(\mathbf{X}) = c \cdot 1\{X_2 > 0 \text{ or } X_5 < -0.5\}.$$

Table 3.1 displays which covariates, $\mu(\mathbf{X})$, and $\theta(\mathbf{X})$ are employed in each model. For model V, we consider the following high-dimensional setting:

$$\begin{aligned}
X_r &\overset{iid}{\sim} \mathrm{N}(0.2r - 0.6, 1),\ r = 1,2,3,4,5 \\
X_r &\overset{iid}{\sim} \mathrm{N}(0.2r - 1.6, 2),\ r = 6,7,8,9,10 \\
X_r &\overset{iid}{\sim} \mathrm{U}[-0.5r + 5, 0.5r - 5],\ r = 11,12,13 \\
X_{14} &\overset{iid}{\sim} \mathrm{U}[-0.5, 1.5],\ X_{15} \overset{iid}{\sim} \mathrm{U}[-1.5, 0.5] \\
X_r &\overset{iid}{\sim} \mathrm{Ber}(0.2r - 3.1),\ r = 16,17,18,19,20 \\
\mu(\mathbf{X}) &= -0.8 + X_{18} + 0.5X_{12} - X_3^2 \\
\theta(\mathbf{X}) &= c \cdot 1\{(X_{14} > -0.1)\&(X_{20} = 1)\},
\end{aligned}$$

where c varies among $\{-1, 0, 0.6, 0.8, 1\}$ indicating the intensity of the value difference. When $c = -1$ and 0, the null hypothesis is true and the type I error is estimated; while when $c = 0.6, 0.8, 1$, the alternative is true and the power is estimated.

Table 3.1 The covariates, $\mu(\mathbf{X})$, and $\theta(\mathbf{X})$ in Models I-IV

Model	Input covariates	$\mu(\mathbf{X})$	$\theta(\mathbf{X})$
I	X_1, X_3	$\mu_1(\mathbf{X})$	$\theta_1(\mathbf{X})$
II	X_1, X_2, X_3, X_4, X_5	$\mu_2(\mathbf{X})$	$\theta_2(\mathbf{X})$
III	X_1, X_2, X_3, X_4, X_5	$\mu_1(\mathbf{X})$	$\theta_2(\mathbf{X})$
IV	X_1, X_2, X_3, X_4, X_5	$\mu_2(\mathbf{X})$	$\theta_1(\mathbf{X})$

Table 3.2 shows that the SUBTLE is able to control type I error and achieve competing detection power, especially under high-dimensional setting (Model V); however, SST couldn't control type I error especially when $c = -1$. This can be explained by two things: (i) the linearity of the model (3.2) is violated; (ii) SST is testing if there is a difference between treatment and control groups among any subjects, instead of the existence of a beneficial subgroup. Specifically, SST is testing if the least false parameter $\boldsymbol{\theta}^*$, to which the MLE of $\boldsymbol{\theta}$ under model misspecification converges, is zero or not. We also perform experiments with batch size $m = 40$, and the results (Table 3.3) do not have much difference.

Table 3.2 Estimated type I error and power of SUBTLE and SST with batch size 20

Model	I		II		III		IV		V	
c	SUBTLE	SST	SUBTLE	SST	SUBTLE	SST	SUBTLE	SST	SUBTLE	SST
-1	0.009	0.695	0.002	0.589	0.003	0.224	0.004	0.411	0.002	0.008
0	0.015	0.134	0.010	0.023	0.006	0.095	0.010	0.023	0.006	0.038
0.6	0.323	0.564	0.491	0.513	0.269	0.389	0.424	0.425	0.559	0.170
0.8	0.623	0.845	0.878	0.900	0.719	0.723	0.822	0.824	0.925	0.390
1	0.911	0.974	0.988	0.996	0.952	0.943	0.985	0.982	0.997	0.742

Table 3.3 Estimated type I error and power of SUBTLE and SST with batch size 40

Model	I		II		III		IV		V	
c	SUBTLE	SST	SUBTLE	SST	SUBTLE	SST	SUBTLE	SST	SUBTLE	SST
-1	0.003	0.662	0.000	0.588	0.000	0.219	0.000	0.368	0.002	0.006
0	0.012	0.126	0.002	0.023	0.003	0.077	0.002	0.023	0.006	0.034
0.6	0.297	0.552	0.465	0.549	0.326	0.397	0.414	0.451	0.585	0.216
0.8	0.633	0.837	0.868	0.896	0.680	0.703	0.826	0.816	0.931	0.373
1	0.901	0.969	0.993	0.995	0.947	0.933	0.985	0.978	0.999	0.715

3.4.2 Noise covariates

It is common in practice that a large number of covariates are incorporated in the experiment whereas the actual outcome only depends on a few of them. Some covariates do not have any effect on the response, like X_4 in Model II, III, IV, and we call them *noise covariates*. In the following simulation, we explore the impact of noise covariates on the detection power of SUBTLE and SST. We choose Model I with $c = 0.8$ as the base model, and at each time add three noise covariates which are respectively from normal $\mathrm{N}(0,1)$, uniform $\mathrm{U}[-1,1]$, and Bernoulli $\mathrm{Ber}(0.5)$ distributions. The batch size is set to $m = 40$ for computation efficiency. Figure 3.1 shows that SST has continuously decreasing powers as the number of noise covariates increases, while the power of SUBTLE is more robust to the noise covariates.

3.4.3 Stopping time

A key feature of sequential testing is that it has an expected smaller sample size than fixed-horizon testing. For comparison, we consider a fixed-horizon version of SUBTLE, which leverages Theorem 3.3.1 and rejects the null hypothesis $\mathrm{H}_0 : \Delta = 0$ when $R_k > Z_\alpha$ for some predetermined k, where Z_α denotes the $(1-\alpha)$ quantile of standard normal distribution. We assume $\sigma^{-1} = \lim_{k\to\infty} \frac{1}{k}\sum_{j=1}^{k} \hat{\sigma}_j^{-1}$, then the required number of batches k can be calculated as $k = \sigma^2(Z_\alpha + Z_{1-\text{power}})^2/\Delta^2$, and thus the required sample size is $n = k \times m + l$. The true value difference Δ can be directly estimated from data generated under the true model with two treatment rules, while σ^2 is estimated by the sample variance of $\hat{\Delta}_{k'}$ times k' for

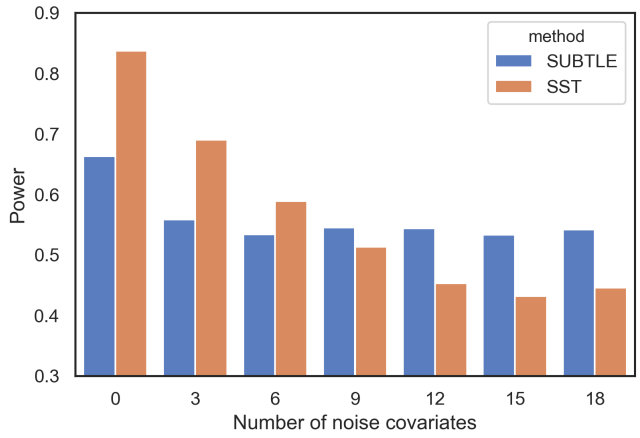

Figure 3.1 Estimated power of SUBTLE and SST with data generated from Model I as the number of noise covariates increases

some fixed large k'. Here, we choose Model V with $c = 1$ and batch size $m = 20$. The stopping sample size of our sequential SUBTLE over 1000 replicates is shown in Figure 3.2, and the dashed vertical line indicates the required sample size for the fixed-horizon SUBTLE with the same power of 0.997 (seen from Table 3.2) under the same setting. We can find that most of the time our sequential SUBTLE arrives at the decision early than the fixed-horizon version, but occasionally it can take longer. The distribution of the stopping time for sequential SUBTLE is right-skewed, which is in line with the findings in Johari et al. [Joh15] and Ju et al. [Ju19].

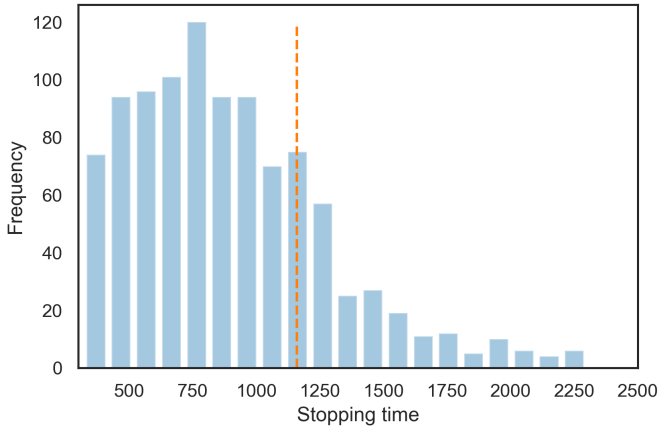

Figure 3.2 Histogram of stopping time of SUBTLE and SST with data generated from Model V

3.5 Real Data Experiments

We use the Yahoo dataset to examine the performance of our SUBTLE, which contains user click events on articles over 10 days. Each event has a timestamp, a unique article id (variant), a binary click indicator (response), and four independent user features (covariates). We choose two articles (id=109520 and 109510) with the highest click-through rates as control and treatment, respectively. We set the significance level $\alpha = 0.05$, initial batch size and batch size $l = m = 200$, and the failure time $M = 50000$.

To demonstrate the false positive control of our method, we conduct an A/A test and a permutation test. For the A/A test, we only use data from article 109510 and randomly generate fake treatment indicators. Our method accepts the null hypothesis. For the permutation test, we use combined data from articles 109510 and 109520 and permute their response 1000 times while leaving the treatment indicator and covariates unchanged. The estimated false positive rate is below the significance level.

Then we test if there is any subgroup of users who would have a higher click-through rate on article 109510. In this experiment, SUBTLE rejects the null hypothesis with a sample size of 12400. We identify the beneficial subgroup $1\{\theta(\mathbf{X}) > 0\}$ by estimating $\hat{\theta}(\mathbf{X})$ with random forest on the first 12400 observations. To get a structured optimal treatment rule, we then build a classification tree on the same 12400 samples with a random forest estimator $1\{\hat{\theta}(\mathbf{X}) > 0\}$ as true labels. The resulting decision tree (Figure 3.3) suggests that the users in the subgroup defined by $\{X_3 < 0.7094 \text{ or } (X_3 \geq 0.7094, X_1 \geq 0.0318 \text{ and } X_4 < 0.0003)\}$ have a higher click-through rate on article 109510 compared with article 109520.

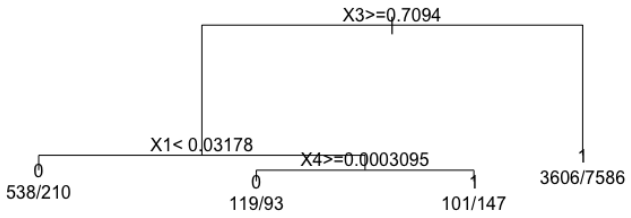

Figure 3.3 An optimal treatment rule in the Yahoo dataset by a decision tree

We then use the 50000 samples after the first 12400 samples as test data and compute the difference of click-through rates between articles 109510 and 109520 on the test data (overall treatment effect), and the same difference in the subgroup of the test data (subgroup treatment effect). We found that the subgroup treatment effect of 0.009 is larger than the overall treatment effect of 0.006, which shows that the identified subgroup has enhanced treatment effects than the overall population.

We further compute the *inverse probability weighted* (IPW) estimator $n^{-1}\sum_{i=1}^{n}[Y_i \cdot 1\{A_i = d(\mathbf{X}_i)\}/p_{A_i}(\mathbf{X}_i)]$ of the values of two treatment rules using the test data: a treatment rule $d_1(\mathbf{X}) = 0$

that assigns everyone to article 109520 and the optimal treatment rule $d_2(\mathbf{X}) = 1\{\hat{\theta}(\mathbf{X}) > 0\}$ estimated by random forest. Their IPW estimates are respectively 0.043 and 0.049, which suggests that the estimated optimal treatment rule is better than the fixed rule that assigns all users to article 109520 in terms of click-through rate. It also implies that there exists a subgroup of the population that has a higher click-through rate on article 109510 compared with article 109520.

3.6 Conclusion

In this chapter, we propose SUBTLE, which is able to sequentially test if some subgroup of the population will benefit from the investigative treatment. If the null hypothesis is rejected, a beneficial subgroup can be easily identified based on the estimated optimal treatment rule. The validity of the test has been proved by both theoretical and simulation results. The experiments also show that SUBTLE has high detection power especially under high-dimensional settings, is robust to noise covariates, and allows quick inference most of the time compared with fixed-horizon testing.

Same as mSPRT and SST, the rejection condition of SUBTLE may never be reached in some cases, especially when the true effect size is negligible. Thus, a failure time is needed to terminate the test externally and accept the null hypothesis if we ever reach it. How to choose a failure time to balance waiting time and power need to be studied in the future. Another future direction is the application of our test under adaptive allocation, where users will have higher probabilities of being assigned to a beneficial variant based on previous observations. However, the validity may not be guaranteed anymore under adaptive allocation and more theoretical investigations are needed.

CHAPTER

4

MULTIPLICATIVE STRUCTURAL NESTED MEAN MODEL FOR ZERO-INFLATED OUTCOMES

4.1 Introduction

Due to the ubiquitous presence of smartphones and the transition of casual gaming to mobile devices, mobile games (referred to as the video games played on an Internet-enabled mobile device such as tablets, phones, etc.) are getting increasingly popular nowadays [McD17; Ban19]. Based on a new industry study by Golden Casino News [Gre19], mobile games make up 60% of revenue for the global video game market, becoming the most significant segment of the video game industry beyond console and PC games. It is reported that mobile games comprise 33% of app downloads, 74% of consumer spending, and 10% of time spent on apps [Kap19]. In fact, there are 1.36 billion users playing mobile games worldwide and this number will continue to grow [Gre19].

A common monetization strategy for mobile games is the freemium business model [And09], which provides free downloads and basic gameplay to attract customers, and then offers an option to pay for premium content such as in-game currency, extra content, or customization. The free part helps to increase the size of the user base quickly [Bou19] while the premium part generates revenue. Industry reports show that over 90% of mobile games begin as free, and over 90% of the profits from mobile games come from games that began as free [Ban19; App20]. Apple also found that the in-app purchases of freemium mobile games account for about 70-80% of the iOS revenue each year [Jac15]. To retain

users and stimulate consumption, developers offer promotions to players from time to time, such as sales on the add-on components, more in-game rewards, and holiday promotions. Understanding the effects of a sequence of promotion decisions on daily engagement (i.e. purchase) for heterogeneous users helps the game managers make improved/personalized promotion strategies.

Freemium mobile game data has its unique characteristics, which lead to several challenges in statistical inference. First, the outcomes, i.e. daily engagement, are zero-inflated. It is common to remain free users and never engage in the premium part of freemium games. Moreover, the positive daily engagement is skewed to the right. A 2014 study of freemium mobile games found that 50% of mobile gaming revenue came from the top 10% of players making purchases, which only accounts for 0.15% of total players [Swr14]. Common distributional assumptions such as the normal or gamma are thereby no longer appropriate for freemium mobile game data.

Second, players are followed over a period of time, during which a sequence of promotion decisions (promotion or no promotion, referred to as treatments) are implemented. We are interested in estimating the joint effects of a sequence of treatments on the outcomes rather than a one-time treatment. However, there exist time-varying confounders with the following characteristics: they are (i) associated with the outcomes, (ii) affected by earlier treatment, and (iii) predictive of subsequent treatments. For example, for estimating the promotion effects on user daily engagement, daily activity time is likely to be a time-varying confounder. This is because players with longer playing time are more likely to purchase in the game, players who received promotions tend to spend more time playing games, and activity time may also be an important factor in making subsequent promotion decisions. In the presence of time-varying confounders, standard regression methods, whether or not adjusting for time-varying confounders, are inappropriate for estimating the causal effects of a sequence of treatments [Rob09]; see Figure 4.1 for an illustration.

Lastly, in practice, the promotion assignment over time can be personalized or uniform. We refer to the treatment assignment that allocates the same treatment to all users at a given time as *cluster randomization*, while the personalized treatment assignment is *individual randomization*. In cluster randomization, since all users receive the same treatment at a given time, it increases the difficulties in estimation of treatment effects due to the convergence of the propensity score estimate reliant solely on the number of time points, and thus the standard asymptotic framework requiring only the sample size go to infinity is not sufficient to derive the large sample results.

Exiting works for estimating the causal effect of one-time treatment with semicontinuous outcomes with excessive zeros include the *Two-part model* [Dua83], the *Burden-of-Illness model* [Cha94], and the *Tobit model* or its variants [Tob58; Pow86; Kee19; Che20]. Estimating the causal effects of a sequence of treatments is considerably more challenging in the presence of time-varying confounding. *Structural nested mean models* [Rob94a] have been proposed to overcome this challenge by modeling treatment effects sequentially over time, and G-estimation can be used to disentangle treatment effects from confounding effects. See [Van14] for a review. However, existing structural nested mean models cannot directly handle zero-inflated outcomes, and the standard asymptotic regime for G-estimation requires the sample size to increase to infinity and does not apply to the case when the number of follow-up times

increases while the sample size can be fixed.

In this chapter, we propose a class of *multiplicative structural nested mean models* for zero-inflated nonnegative outcomes. Our contributions are from several folds. First, our proposed model describes flexibly and concisely the joint effects of a sequence of treatments in the presence of time-varying confounders based on the ratio of conditional means of outcomes, which can naturally accommodate zero-inflated nonnegative outcomes as demonstrated by the example given in the next section. Second, for parameter estimation, we propose a class of doubly robust estimating equations, where the nuisance functions, i.e. propensity score and conditional means of outcomes given confounders, are estimated parametrically or nonparametrically. Moreover, the conditional means of outcomes are estimated in two parts to leverage the characteristic of zero-inflated outcomes. That is, we separately model the probability of having positive outcomes and the mean outcome conditional on its being positive, both parts conditional on confounders. This facilitates the derivation of the doubly robust estimator with the desired theoretical properties. Third, in terms of theory development, we establish the consistency and asymptotic normality of the resulting doubly robust estimator as either the sample size or the follow-up time goes to infinity under the setting for individual randomization of treatment assignment. This task is nontrivial since individual data are dependent over time, and uniform consistency and weak convergence results need to be established for general martingale processes. We also establish similar results for the setting with cluster randomization as the follow-up time goes to infinity. However, the convergence rate is slower than the setting of individual randomization since it is only determined by the number of observations over time regardless of the sample size. Moreover, we show that the typical sandwich formula can be used for consistent variance estimation without accounting for the variation due to the estimation of nuisance functions, making our inference procedure easy to implement.

The rest of the chapter is structured as follows. We introduce the model setup in Section 4.2. In Section 4.3, we discuss the identifying assumptions and the associated estimation and inference procedure. Theoretical results are also presented. Simulation studies and a real data application are conducted to evaluate the empirical performance of our proposed method in Section 4.4 and 4.5, respectively. We conclude the chapter with some discussions in Section 4.6. All the technical proofs are provided in the Appendix C.

4.2 Multiplicative Structural Nested Mean Model

Suppose that measurements are collected at T discrete time points. Let $L_{0,t}$ be the time-varying covariates collected at time point t, $A_{0,t} \in \{0,1\}$ denote the treatment indicator (1 for promotion and 0 for no promotion) at time t, and $Y_{0,t}$ stand for the observed outcome at time t, $t = 1, \cdots, T$. We presume that the observed data are ordered as $L_{0,1}, A_{0,1}, Y_{0,1}, L_{0,2}, \cdots$; thus the covariates and treatments precede the observed outcomes and $Y_{0,t}$ can be a part of $L_{0,t+1}$. We use the overline notation to denote a variable's history and omit the subscript T when $t = T$; e.g. $\overline{A}_{0,t} = (A_{0,1}, \cdots, A_{0,t})$ and $\overline{A}_0 = \overline{A}_{0,T}$. Figure 4.1 illustrates the challenges due to time-varying confounding. Standard regression approaches, whether or not adjusting for confounders, are unable to estimate the causal effects of treatments.

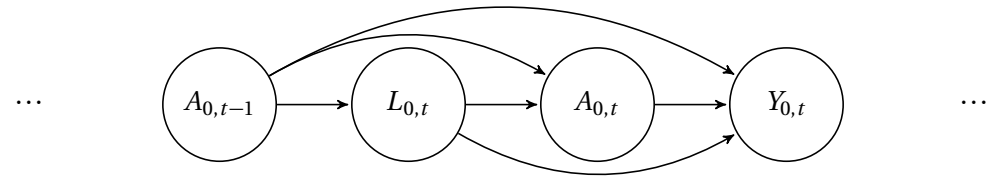

Figure 4.1 Illustration of challenges due to time-varying confounding: $L_{0,t}$ is a confounder for $A_{0,t}$ and $Y_{0,t}$, which should be controlled for; however, because $L_{0,t}$ is an intermediate variable from $A_{0,t-1}$ to $Y_{0,t}$, if we do control for $L_{0,t}$, then the causal effect of $A_{0,t-1}$ to $Y_{0,t}$ will be blocked.

We use the potential outcomes framework to define the causal effect of treatments. Let $Y_{0,t}^{(\overline{a})}$ denote the potential outcome that would be seen at time t, had the subject received the sequence of treatments $\overline{a}$ through time T. In particular, $Y_{0,t}^{(\overline{0})}$ is the potential outcome at time t had the subjects never received treatments. We are interested in estimating the casual effects of a sequence of treatments $\overline{a}$ on a group of users with covariates sequence $\overline{l}$, which can be defined as the ratio of the conditional expectation of the potential outcomes had this group of subjects received $\overline{a}$ and $\overline{0}$, i.e. $\mathbb{E}(Y_{0,t}^{(\overline{a})}|\overline{A}_0=\overline{a},\overline{L}_0=\overline{l})/\mathbb{E}(Y_{0,t}^{(\overline{0})}|\overline{A}_0=\overline{a},\overline{L}_0=\overline{l})$. Here we use the ratio rather than the difference of the conditional mean outcomes to accommodate zero-inflated nonnegative outcomes.

We presume that the treatments and covariates after time t cannot affect the potential outcomes at times up to t, and consider a class of semiparametric multiplicative structural nested mean models in the following form:

$$\frac{\mathbb{E}(Y_{0,t}^{(\overline{a})}|\overline{A}_0=\overline{a},\overline{L}_0=\overline{l})}{\mathbb{E}(Y_{0,t}^{(\overline{0})}|\overline{A}_0=\overline{a},\overline{L}_0=\overline{l})}=\exp\{f_{\theta_0}(\overline{v}_t)\cdot a_t\}, \tag{4.1}$$

where $\overline{v}_t=(\overline{l}_t,\overline{a}_{t-1})$; $\overline{l}_t,\overline{a}_t$ respectively represent the first t elements of the sequences $\overline{l}$ and $\overline{a}$; a_t is the t-th element of the sequence $\overline{a}$; $f_\theta(\cdot)$ is a known function with a p-dimensional vector of parameters $\theta\in\Theta\subseteq\mathbb{R}^p$ with true parameter value θ_0. Typically, the parameterization is chosen to be $f_\theta(\cdot)\equiv 0$ for $\theta=0$, so that $\theta_0=0$ encodes no treatment effects. The proposed multiplicative structural nested mean model is semiparametric in nature because the conditional mean $\mathbb{E}(Y_{0,t}^{(\overline{0})}|\overline{A}_0=\overline{a},\overline{L}_0=\overline{l})$ is completely unspecified.

Next, we use a toy example to illustrate the proposed treatment effect model.

Example 4.2.1 *Use the freemium mobile game data as an illustration. Suppose the potential daily engagement follows a zero-inflated log-normal distribution. Specifically, $Y_{0,t}^{(\overline{a})}$ has a probability p_t to be positive and follows a log-normal distribution Log-normal(ν_t,σ_t^2) and a probability $(1-p_t)$ to be zero, where p_t and ν_t are functions of the covariates and potential treatments up to time t. Suppose p_t follows a log-linear model and ν_t is a linear model. Then, the conditional mean of the potential outcome is*

$$\mathbb{E}(Y_{0,t}^{(\overline{a})}|\overline{A}_0=\overline{a},\overline{L}_0=\overline{l})=p_t\exp(\nu_t+\frac{1}{2}\sigma_t^2)$$

$$= \exp(\beta_p^T \tilde{l}_t + \gamma_p^T \tilde{l}_t \cdot a_t) \cdot \exp(\beta_v^T \tilde{l}_t + \gamma_v^T \tilde{l}_t \cdot a_t + \frac{1}{2}\sigma_t^2),$$

where $\tilde{l}_t = (1, l_t^T)^T$ *and* $\beta_p, \gamma_p, \beta_v, \gamma_v$ *are corresponding coefficients. Thus, the multiplicative structural nested mean model is*

$$\frac{\mathbb{E}(Y_{0,t}^{(\overline{a})} | \overline{A}_0 = \overline{a}, \overline{L}_0 = \overline{l})}{\mathbb{E}(Y_{0,t}^{(\overline{0})} | \overline{A}_0 = \overline{a}, \overline{L}_0 = \overline{l})} = \exp\{(\gamma_p + \gamma_v)^T \tilde{l}_t \cdot a_t\},$$

which satisfies model (4.1).

4.3 Main Methodology

For notational convenience, let $V_{0,t} = (A_{0,t-1}, L_{0,t})$ and $O_{0,t} = (L_{0,t}, A_{0,t}, Y_{0,t})$. Then, a subject's full record can be represented as $\overline{O}_0 = \{(L_{0,t}, A_{0,t}, Y_{0,t})\}_{1 \le t \le T}$. Suppose n subjects are randomly sampled from a population, so that the observed data $\overline{O}_1, \overline{O}_2, \cdots, \overline{O}_n$ are i.i.d. copies of $\overline{O}_0$. We denote the associated probability space by $(\Omega, \mathscr{F}, P)$.

4.3.1 Identification and estimation

If all potential outcomes were observed for each subject, we can directly compare these outcomes to infer the treatment effect; however, the fundamental problem in causal inference is that we can not observe all potential outcomes for a particular subject. For the causal treatment effects to be estimable based on observed data, we require the following assumptions as widely used in the literature.

Assumption 4.3.1 (Consistency) *The observed outcome is equal to the potential outcome under the sequence of actual treatments received; i.e.* $Y_{0,t} = Y_{0,t}^{(\overline{A}_0)}$ *for all* $t \in \{1, \cdots, T\}$.

Assumption 4.3.2 (No unmeasured confounders) $A_{0,t} \perp\!\!\!\perp Y_{0,t}^{(\overline{0})} \,|\, \overline{V}_{0,t}$*, which means that* $A_{0,t}$ *is conditionally independent of* $Y_{0,t}^{(\overline{0})}$ *given* $\overline{V}_{0,t}$.

The consistency assumption links the observed data to the potential outcome, which implicitly makes the stable unit treatment assumption that rules out multiple versions of treatment and interference. The no unmeasured confounders assumption is required for the identification of the causal parameters [Rob92]. It is plausible if, at each time t, the observed histories of covariates $\overline{L}_{0,t}$ and treatments $\overline{A}_{0,t-1}$ capture all predictors of both the potential outcome $Y_{0,t}^{(\overline{0})}$ and the current treatment $A_{0,t}$.

We develop a G-estimator of the causal parameter θ_0 in model (4.1). Define

$$H_{0,t}(\theta_0) = Y_{0,t} \cdot \exp\left\{-f_{\theta_0}(\overline{V}_{0,t}) \cdot A_{0,t}\right\}.$$

Intuitively, $H_{0,t}(\theta_0)$ mimics the potential outcome $Y_{0,t}^{(\overline{0})}$ that would have been seen had the treatment never been implemented. The following two propositions establish the properties of $H_{0,t}(\theta_0)$. In particular, Proposition 4.3.1 shows that with the consistency assumption, the conditional mean of $H_{0,t}(\theta_0)$ is equal to the conditional mean of the potential outcome $Y_{0,t}^{(\overline{0})}$. Proposition 4.3.2 states that with the consistency

and no unmeasured confounders assumptions, $H_{0,t}(\theta_0)$ is conditionally independent of $A_{0,t}$ given $\overline{L}_{0,t}$ and $\overline{A}_{0,t-1}$.

Proposition 4.3.1 *Under Assumption 4.3.1, we have*

$$\mathbb{E}\{H_{0,t}(\theta_0)|A_{0,t}, \overline{V}_{0,t}\} = \mathbb{E}\{Y_{0,t}^{(\overline{0})}|A_{0,t}, \overline{V}_{0,t}\}$$

for all $t \in \{1, \cdots, T\}$.

Proposition 4.3.2 *Under Assumptions 4.3.1 and 4.3.2, we have*

$$\mathbb{E}\{H_{0,t}(\theta_0)|A_{0,t}, \overline{V}_{0,t}\} = \mathbb{E}\{H_{0,t}(\theta_0)|\overline{V}_{0,t}\}$$

for all $t \in \{1, \cdots, T\}$.

Let $h_0(\overline{v}_t) = \mathbb{E}\{H_{0,t}(\theta_0)|\overline{V}_{0,t} = \overline{v}_t\}$ and $\pi_0(\overline{v}_t) = P(A_{0,t} = 1|\overline{V}_{0,t} = \overline{v}_t)$ denote the conditional mean of $H_{0,t}(\theta_0)$ and the propensity score, respectively. For any given propensity score π and conditional mean h, we consider a class of doubly robust estimating functions for θ as below:

$$\psi(\overline{O}_{0,t}; \theta, \pi, h, c) = c(\overline{V}_{0,t})\left\{H_{0,t}(\theta) - h(\overline{V}_{0,t})\right\}\left\{I(A_{0,t} = 1) - \pi(\overline{V}_{0,t})\right\}, \tag{4.2}$$

where $c(\cdot)$ is a p-dimensional function of the covariate and treatment history $\overline{V}_{0,t}$. It can be shown that under Proposition 4.3.2 the expectation of (4.2) is zero for $\theta = \theta_0$, regardless of c as long as $\pi = \pi_0$ or $h = h_0$, which is presented in Proposition 4.3.3.

Proposition 4.3.3 *Under Assumptions 4.3.1 and 4.3.2, and the multiplicative structural nested mean model (4.1), the estimating function*

$$\mathbf{G}(\theta; \pi, h, c) \equiv \frac{1}{nT}\sum_{i=1}^{n}\sum_{t=1}^{T}\psi(\overline{O}_{i,t}; \theta, \pi, h, c)$$

is unbiased when $\pi = \pi_0$ *or* $h = h_0$.

In practice, π_0 and h_0 are unknown and need to be estimated from data. The propensity score model can be estimated parametrically, such as logistic regression, or nonparametrically, such as random forest. In mobile game data applications, a treatment is usually randomly assigned using either individual randomization or cluster randomization, where the propensity score can be estimated by the sample proportion. However, we cannot directly estimate h_0 by regressing $H_{i,t}(\theta_0)$ on $\overline{V}_{0,t}$ because $H_{i,t}(\theta_0)$ depends on the unknown parameters θ_0. Thus, we consider representing h_0 with other estimable conditional means. Define $\mu_{0,0}(\overline{v}_t) = \mathbb{E}(Y_{0,t}|A_{0,t} = 0, \overline{V}_{0,t} = \overline{v}_t)$ and $\mu_{1,0}(\overline{v}_t) = \mathbb{E}(Y_{0,t}|A_{0,t} = 1, \overline{V}_{0,t} = \overline{v}_t)$. Then, h_0 can be represented as

$$h_0(\overline{v}_t) = \mu_{1,0}(\overline{v}_t)\exp\{-f_{\theta_0}(\overline{v}_t)\}\pi_0(\overline{v}_t) + \mu_{0,0}(\overline{v}_t)\{1 - \pi_0(\overline{v}_t)\}.$$

Accordingly, we rewrite our estimating equation as

$$\frac{1}{nT}\sum_{i=1}^{n}\sum_{t=1}^{T}\psi(\overline{O}_{i,t};\theta,\hat{\pi},\hat{\mu}_1,\hat{\mu}_0,c)=\mathbf{0}, \tag{4.3}$$

where $\hat{\pi},\hat{\mu}_1,\hat{\mu}_0$ are the estimators of $\pi_0,\mu_{1,0},\mu_{0,0}$ respectively, and

$$\begin{aligned}\psi(\overline{O}_{i,t};\theta,\hat{\pi},\hat{\mu}_1,\hat{\mu}_0,c)=c(\overline{V}_{i,t})\big[H_{i,t}(\theta)-\hat{\mu}_1(\overline{V}_{i,t})\exp\{-f_\theta(\overline{V}_{i,t})\}\hat{\pi}(\overline{V}_{i,t})-\hat{\mu}_0(\overline{V}_{i,t})\{1-\hat{\pi}(\overline{V}_{i,t})\}\big]\\ \times\big\{I(A_{i,t}=1)-\hat{\pi}(\overline{V}_{i,t})\big\}.\end{aligned}$$

Taking into account that outcomes are zero-inflated, $\mu_{1,0}$ and $\mu_{0,0}$ can be represented as

$$\mu_{a,0}(\overline{v}_t)=P(Y_{0,t}>0|A_{0,t}=a,\overline{V}_{0,t}=\overline{v}_t)\cdot\mathbb{E}(Y_{0,t}|Y_{0,t}>0,A_{0,t}=a,\overline{V}_{0,t}=\overline{v}_t),\ a=0,1. \tag{4.4}$$

To improve the estimation accuracy, we estimate $\mu_{1,0}$ and $\mu_{0,0}$ by modeling the probability part and the mean part on the right-hand side of (4.4) separately, using some nonparametric regression techniques, such as generalized additive models (GAMs).

The choice of function $c(\cdot)$ generally does not affect the consistency of the estimator but may make a difference in the efficiency. Following Robins [Rob94a], it can be shown that under certain conditions, an efficient estimator of θ_0 can be obtained by setting

$$c(\overline{v}_t)=\left.\frac{\mathbb{E}\left(\frac{\partial H_{0,t}(\theta)}{\partial\theta}\middle|A_{0,t}=1,\overline{V}_{0,t}=\overline{v}_t\right)-\mathbb{E}\left(\frac{\partial H_{0,t}(\theta)}{\partial\theta}\middle|A_{0,t}=0,\overline{V}_{0,t}=\overline{v}_t\right)}{\mathrm{Var}(Y_{0,t}^{(\overline{0})}|\overline{V}_{0,t}=\overline{v}_t)}\right|_{\theta=\theta_0}. \tag{4.5}$$

The optimal c function depends on the conditional variance of $Y_{0,t}^{(\overline{0})}$, which may be difficult to estimate well based on observed data. In practice, we can choose a simple function for c, such as $c(\overline{v}_t)=\partial f_\theta(\overline{v}_t)/\partial\theta$. We compare the empirical performance of the estimators constructed based on the simple c function and the optimal c function in our simulations; see Section 4.4.

4.3.2 Asymptotic distribution and variance estimation

Let $\hat{\theta}$ denote the proposed estimator by solving the estimating equation (4.3). Before delving into the theoretical analysis of $\hat{\theta}$, we introduce more notation for simplicity of exposition. Define the triplet function $\eta=(\pi,\mu_1,\mu_0)$. The true value of η is given by $\eta_0=(\pi_0,\mu_{1,0},\mu_{0,0})$, and its estimator is given by $\hat{\eta}=(\hat{\pi},\hat{\mu}_1,\hat{\mu}_0)$. Let $\psi_{i,t}(\theta,\eta)=\psi(\overline{O}_{i,t};\theta,\pi,\mu_1,\mu_0,c)$ and

$$\mathbb{P}g\{\psi(\theta,\eta)\}=\frac{1}{T}\sum_{t=1}^{T}\mathbb{E}\big[g\{\psi_{0,t}(\theta,\eta)\}\big], \tag{4.6}$$

$$\mathbb{P}_n g\{\psi(\theta,\eta)\}=\frac{1}{nT}\sum_{i=1}^{n}\sum_{t=1}^{T}g\{\psi_{i,t}(\theta,\eta)\}, \tag{4.7}$$

where g is any given function or operator of $\psi_{0,t}$, e.g. $g\{\psi_{0,t}(\theta,\eta)\} = \partial\psi_{0,t}(\theta,\eta)/\partial\theta$. In formula (4.6), we assume $T \in \mathbb{N}^+ \cup \{\infty\}$ and define $\frac{1}{T}\sum_{t=1}^{T} = \lim_{T\to\infty}\frac{1}{T}\sum_{t=1}^{T}$ when $T=\infty$. Denote the Euclidean norm by $\|\cdot\|_2$ and the $\mathscr{L}_2(P)$ norm by $\|\cdot\|_{2,P}$ with the definition $\|\eta\|_{2,P}^2 = \frac{1}{T}\sum_{t=1}^{T}\int\|\eta(\overline{V}_{0,t})\|_2^2\, dP(\overline{V}_{0,t})$. Define a function set $\mathscr{G}_{\eta_0} = \{\eta : \|\eta-\eta_0\|_{2,P} < \delta\}$ for some $\delta > 0$ and a Cartesian product $\mathscr{U} = \{(\theta,\eta) : \theta\in\Theta, \eta\in\mathscr{G}_{\eta_0}\}$.

To establish the asymptotic normality of $\hat{\theta}$ under individual randomization, we require the following conditions.

(C1) The solution to $\mathbb{P}\psi(\theta,\eta_0) = \mathbf{0}$ is unique. $\left\|\mathbb{P}\psi(\theta_n,\eta_0)\right\|_2 \to 0$ implies $\|\theta_n-\theta_0\|_2 \to 0$ for any sequence of $\{\theta_n\}\in\Theta$.

(C2) There exists a finite ϵ-net $\mathscr{U}_\epsilon$ of $\mathscr{U}$ for any $\epsilon>0$. In addition, $\mathscr{G}_{\eta_0}$ has uniformly integrable entropy. That is,

$$\int_0^\infty \sup_Q \sqrt{\log N(\epsilon\|F\|_{2,Q}, \mathscr{G}_{\eta_0}, \|\cdot\|_{2,Q})}\, d\epsilon < \infty,$$

where $F:\Omega\to\mathbb{R}^3$ is a square integrable envelop for $\mathscr{G}_{\eta_0}$, and the covering number $N(\epsilon,\mathscr{G}_{\eta_0},\|\cdot\|)$ is the minimum number of balls $\{\eta' : \|\eta'-\eta\| < \epsilon\}$ of radius ϵ needed to cover $\mathscr{G}_{\eta_0}$.

(C3) (i) $|f_\theta(\overline{V}_{0,t})|$ is bounded almost surely for all $\theta\in\Theta$ and $t\in\{1,\cdots,T\}$.

(ii) $|Y_{0,t}|$, $\|c(\overline{V}_{0,t})\|_2$, $\sigma_{0,0}^2(\overline{V}_{0,t})$, $\sigma_{1,0}^2(\overline{V}_{0,t})$ are bounded almost surely for $t\in\{1,\cdots,T\}$, where $\sigma_{a,0}^2(\overline{v}_t) = \mathrm{Var}(Y_{0,t}|A_{0,t}=a, \overline{V}_{0,t}=\overline{v}_t)$, $a=0,1$.

(iii) $\|\eta(\overline{V}_{0,t})\|_2 < b_0$ almost surely for all $\eta\in\mathscr{G}_{\eta_0}$ and $t\in\{1,\cdots,T\}$ for some constant b_0.

(iv) $\left\|\frac{\partial\psi_{0,t}(\theta,\eta)}{\partial\eta^T}\right\|_2 = \left\|(\frac{\partial\psi_{0,t}(\theta,\eta)}{\partial\pi}, \frac{\partial\psi_{0,t}(\theta,\eta)}{\partial\mu_1}, \frac{\partial\psi_{0,t}(\theta,\eta)}{\partial\mu_0})\right\|_2 < b^*$ almost surely for all $\theta\in\Theta$, $\eta\in\mathscr{G}_{\eta_0}$ and $t\in\{1,\cdots,T\}$ for some constant b^*.

(C4) As $nT\to\infty$,

$$\|\hat{\eta}-\eta_0\|_{2,P} \xrightarrow{P} 0 \quad \text{and}$$
$$(\|\hat{\mu}_1-\mu_{1,0}\|_{2,P} + \|\hat{\mu}_0-\mu_{0,0}\|_{2,P} + \|\hat{\pi}-\pi_0\|_{2,P})\cdot\|\hat{\pi}-\pi_0\|_{2,P} = o_P(1/\sqrt{nT}).$$

(C5) For each $\eta\in\mathscr{G}_{\eta_0}$,

$$\frac{1}{nT}\sum_{i=1}^{n}\sum_{t=1}^{T}\mathbb{E}\left\{\mathbf{m}_{i,t}(\eta)\mathbf{m}_{i,t}^T(\eta)|\overline{V}_{i,t}\right\} \xrightarrow{P} \mathbf{\Sigma}(\eta) \text{ as } nT\to\infty, \tag{4.8}$$

where $\mathbf{m}_{i,t}(\eta) = \psi_{i,t}(\theta_0,\eta) - \mathbb{E}\left\{\psi_{i,t}(\theta_0,\eta)|\overline{V}_{i,t}\right\}$ and $\mathbf{\Sigma}(\eta)$ is a constant positive definite matrix for each η.

(C6) For $j=1,\cdots,p$, there exists a constant M such that as $nT\to\infty$

$$P\left(\sup_{\eta,\eta'\in\mathscr{G}_{\eta_0}} \frac{\frac{1}{nT}\sum_{i=1}^{n}\sum_{t=1}^{T}\mathbb{E}\left[\left\{\mathbf{m}_{i,t}^{(j)}(\eta)-\mathbf{m}_{i,t}^{(j)}(\eta')\right\}^2|\overline{V}_{i,t}\right]}{\|\eta-\eta'\|_{2,P}^2} \geq M\right) \to 0,$$

where $\mathbf{m}_{i,t}(\eta)$ is defined in (C5) and $\mathbf{m}_{i,t}^{(j)}(\eta)$ is the j-th component of $\mathbf{m}_{i,t}(\eta)$.

Condition (C1) is the identifiability condition, which is commonly assumed to ensure the consistency of Z-estimators (see Theorem 2.10 of Kosorok [Kos07]). Condition (C2) controls the complexity of the function set $\mathscr{G}_{\eta_0}$ to show the uniform convergence (Lemma C.1.1) and weak convergence (Lemma C.1.2) of a function-indexed martingale difference sequence. Similar conditions are used in Theorem 2.3 of Kosorok [Kos07] and Theorem 2 of Bae et al. [Bae10]. In fact, many important classes of functions, such as VC graph classes, have uniformly integrable entropy. See Section 2.6 of Van Der Vaart & Wellner [VDV96] for details. Condition (C3) is standard, which generally holds when the covariates L, outcomes Y, and parameters are bounded.

Condition (C4) provides the required convergence rates for the estimators of the nuisance functions under individual randomization. The first part helps to establish the consistency of the estimator $\hat{\theta}$, while the second part is used to derive the asymptotic distribution of $\hat{\theta}$. This condition is widely used in the causal inference literature to derive the asymptotic distribution of doubly robust estimators when nuisance functions are estimated parametrically or nonparametrically with proper rates [see e.g. Far21; Kal20]. For example, if π_0 is estimated based on a correctly specified parametric model so that $\left\|\hat{\pi}-\pi_0\right\|_{2,P} = O_P(1/\sqrt{nT})$, then we only need $\hat{\mu}_0$ and $\hat{\mu}_1$ to be consistent, i.e. $\left\|\hat{\mu}_0-\mu_{0,0}\right\|_{2,P} = o_P(1)$ and $\left\|\hat{\mu}_1-\mu_{1,0}\right\|_{2,P} = o_P(1)$. On the other hand, if all the nuisance functions are estimated nonparametrically, we require $\left\|\hat{\pi}-\pi_0\right\|_{2,P} = o_P\{(nT)^{-1/4}\}$, $\left\|\hat{\mu}_0-\mu_{0,0}\right\|_{2,P} = o_P\{(nT)^{-1/4}\}$, and $\left\|\hat{\mu}_1-\mu_{1,0}\right\|_{2,P} = o_P\{(nT)^{-1/4}\}$. Many nonparametric/semiparametric estimators can satisfy the convergence rate of $o_P\{(nT)^{-1/4}\}$, such as single-index models, generalized additive models, and partially linear models [Hor09].

Conditions (C5) and (C6) are imposed to derive the weak convergence (Lemma C.1.2). Their validity is discussed in Remarks 4.3.1 and 4.3.2 below, respectively.

Remark 4.3.1 (Discussion of (C5)) *The left-hand side of (4.8) can be calculated as*

$$\begin{aligned}\frac{1}{nT}\sum_{i=1}^{n}\sum_{t=1}^{T}\Big[c(\overline{V}_{i,t})c^T(\overline{V}_{i,t})\big\{\{1-\pi(\overline{V}_{i,t})\}^2\pi_0(\overline{V}_{i,t})\exp\{-2f_{\theta_0}(\overline{V}_{i,t})\}\sigma^2_{1,0}(\overline{V}_{i,t})\\ +\ \pi^2(\overline{V}_{i,t})\{1-\pi_0(\overline{V}_{i,t})\}\sigma^2_{0,0}(\overline{V}_{i,t})+\big(\pi_0(\overline{V}_{i,t})\{1-\pi_0(\overline{V}_{i,t})\}\\ \times\ \big[\mu_{1,0}(\overline{V}_{i,t})\exp\{-f_{\theta_0}(\overline{V}_{i,t})\}\pi_0(\overline{V}_{i,t})+\mu_{0,0}(\overline{V}_{i,t})\{1-\pi_0(\overline{V}_{i,t})\}\\ -\ \mu_1(\overline{V}_{i,t})\exp\{-f_{\theta_0}(\overline{V}_{i,t})\}\pi(\overline{V}_{i,t})+\mu_0(\overline{V}_{i,t})\{1-\pi(\overline{V}_{i,t})\}\big]^2\big)\big\}\Big].\end{aligned}$$

When $n\to\infty$ and T is finite, (C5) holds under mild conditions by applying the law of large numbers. When $T\to\infty$ and n is finite, letting

$$X_t = \frac{1}{n}\sum_{i=1}^{n}\mathbb{E}\left\{\mathbf{m}_{i,t}(\eta)\mathbf{m}_{i,t}^T(\eta)|\overline{V}_{i,t}\right\},$$

$\{X_t\}_{t\geq 1}$ forms a stochastic process. If $\{X_t\}_{t\geq 1}$ is uniformly integrable and α-mixing, by Theorem 1 of

Andrews [And88], we have

$$\frac{1}{T}\sum_{t=1}^{T} X_t \xrightarrow{P} \frac{1}{T}\sum_{t=1}^{T}\mathbb{E}(X_t).$$

Thus, (C5) also holds under this setting.

Remark 4.3.2 (Discussion of (C6)) *Let $\eta_{i,t} = \eta(\overline{V}_{i,t})$. Then, we have*

$$\begin{aligned}
&\mathbb{E}\left[\left\{\mathbf{m}_{i,t}^{(j)}(\eta_{i,t}) - \mathbf{m}_{i,t}^{(j)}(\eta'_{i,t})\right\}^2 | \overline{V}_{i,t}\right] \\
&= \mathbb{E}\left[\left\{\frac{\partial \mathbf{m}_{i,t}^{(j)}(\eta_{i,t})}{\partial \eta_{i,t}}|_{\eta=\tilde{\eta}}\right\}^T (\eta_{i,t}-\eta'_{i,t})(\eta_{i,t}-\eta'_{i,t})^T \left\{\frac{\partial \mathbf{m}_{i,t}^{(j)}(\eta_{i,t})}{\partial \eta_{i,t}}|_{\eta=\tilde{\eta}}\right\} \Big| \overline{V}_{i,t}\right] \\
&= \mathbb{E}\left[(\eta_{i,t}-\eta'_{i,t})^T \left\{\frac{\partial \mathbf{m}_{i,t}^{(j)}(\eta_{i,t})}{\partial \eta_{i,t}}|_{\eta=\tilde{\eta}}\right\}\left\{\frac{\partial \mathbf{m}_{i,t}^{(j)}(\eta_{i,t})}{\partial \eta_{i,t}}|_{\eta=\tilde{\eta}}\right\}^T (\eta_{i,t}-\eta'_{i,t}) \Big| \overline{V}_{i,t}\right] \\
&= (\eta_{i,t}-\eta'_{i,t})^T \mathbb{E}\left[\left\{\frac{\partial \mathbf{m}_{i,t}^{(j)}(\eta_{i,t})}{\partial \eta_{i,t}}|_{\eta=\tilde{\eta}}\right\}\left\{\frac{\partial \mathbf{m}_{i,t}^{(j)}(\eta_{i,t})}{\partial \eta_{i,t}}|_{\eta=\tilde{\eta}}\right\}^T \Big| \overline{V}_{i,t}\right](\eta_{i,t}-\eta'_{i,t}),
\end{aligned}$$

where $\|\tilde{\eta}-\eta\|_2 < \|\eta'-\eta\|_2$. If the largest eigenvalue of the expectation term in the above quadratic form is uniformly bounded over i, t by some constant M, we can get

$$\begin{aligned}
&\sup_{\eta,\eta'\in\mathcal{G}_{\eta_0}} \frac{\frac{1}{nT}\sum_{i=1}^{n}\sum_{t=1}^{T}\mathbb{E}\left[\left\{\mathbf{m}_{i,t}^{(j)}(\eta_{i,t}) - \mathbf{m}_{i,t}^{(j)}(\eta'_{i,t})\right\}^2 | \overline{V}_{i,t}\right]}{\|\eta-\eta'\|_{2,P}^2} \\
&= \sup_{\eta,\eta'\in\mathcal{G}_{\eta_0}} \frac{\frac{1}{nT}\sum_{i=1}^{n}\sum_{t=1}^{T}(\eta_{i,t}-\eta'_{i,t})^T \mathbb{E}\left[\left\{\frac{\mathbf{m}_{i,t}^{(j)}(\eta_{i,t})}{\partial \eta_{i,t}}|_{\eta=\tilde{\eta}}\right\}\left\{\frac{\mathbf{m}_{i,t}^{(j)}(\eta_{i,t})}{\partial \eta_{i,t}}|_{\eta=\tilde{\eta}}\right\}^T \Big| \overline{V}_{i,t}\right](\eta_{i,t}-\eta'_{i,t})}{\|\eta-\eta'\|_{2,P}^2} \\
&\le M \cdot \sup_{\eta,\eta'\in\mathcal{G}_{\eta_0}} \frac{\frac{1}{nT}\sum_{i=1}^{n}\sum_{t=1}^{T}\|\eta_{i,t}-\eta'_{i,t}\|_2^2}{\|\eta-\eta'\|_{2,P}^2}.
\end{aligned}$$

If we can show

$$\frac{1}{nT}\sum_{i=1}^{n}\sum_{t=1}^{T}\|\eta_{i,t}-\eta'_{i,t}\|_2^2 \xrightarrow{P} \|\eta-\eta'\|_{2,P}^2 \quad \text{as } nT\to\infty,$$

we are able to derive (C6). The proof of the above convergence can be similarly derived as given in Remark 4.3.1.

In cluster randomization, because the convergence rate of the estimator for propensity score depends only on the length of follow-up time, we replace (C4) with (C4') as follows.

(C4') As $T\to\infty$,

$$\|\hat{\eta}-\eta_0\|_{2,P} \xrightarrow{P} 0 \quad \text{and}$$
$$(\|\hat{\mu}_1-\mu_{1,0}\|_{2,P} + \|\hat{\mu}_0-\mu_{0,0}\|_{2,P} + \|\hat{\pi}-\pi_0\|_{2,P})\cdot\|\hat{\pi}-\pi_0\|_{2,P} = o_P(1/\sqrt{T}).$$

Conditions (C5) and (C6) also need to be modified. Specifically, define $\tilde{\mathbf{m}}_t(\eta) = n^{-1}\sum_{i=1}^n \psi_{i,t}(\theta_0,\eta) - \mathbb{E}\left\{n^{-1}\sum_{i=1}^n \psi_{i,t}(\theta_0,\eta)|\overline{V}_{\cdot,t}\right\}$, where $\overline{V}_{\cdot,t} = (\overline{V}_{1,t},\ldots,\overline{V}_{n,t})$. We require the following modified conditions.

(C5') For each $\eta \in \mathcal{G}_{\eta_0}$,

$$\frac{1}{T}\sum_{t=1}^T \mathbb{E}\left\{\tilde{\mathbf{m}}_t(\eta)\tilde{\mathbf{m}}_t^T(\eta)|\overline{V}_{\cdot,t}\right\} \xrightarrow{P} \mathbf{\Sigma}_1(\eta) \text{ as } T\to\infty,$$

where $\mathbf{\Sigma}_1(\eta)$ is a constant positive definite matrix for each η.

(C6') For $j = 1,\cdots,p$, there exists a constant M such that as $T\to\infty$

$$P\left(\sup_{\eta,\eta'\in\mathcal{G}_{\eta_0}} \frac{\frac{1}{T}\sum_{t=1}^T \mathbb{E}\left[\left\{\tilde{\mathbf{m}}_t^{(j)}(\eta) - \tilde{\mathbf{m}}_t^{(j)}(\eta')\right\}^2|\overline{V}_{\cdot,t}\right]}{\|\eta-\eta'\|_{2,P}^2} \geq M\right) \to 0,$$

where $\tilde{\mathbf{m}}_t^{(j)}(\eta)$ is the j-th component of $\tilde{\mathbf{m}}_t(\eta)$.

We establish the bidirectional asymptotics of $\hat{\theta}$ under individual randomization in Theorem 4.3.1 and the asymptotic distribution of $\hat{\theta}$ under cluster randomization in Theorem 4.3.2.

Theorem 4.3.1 (Bidirectional asymptotics) *If the treatment assignment is individual randomization, under conditions (C1)-(C6), as either $n\to\infty$ or $T\to\infty$, we have*

$$\sqrt{nT}(\hat{\theta}-\theta_0) \xrightarrow{d} MVN\left(0, \mathbf{B}\mathbf{\Sigma}\mathbf{B}^T\right),$$

where

$$\mathbf{B} = \left\{\mathbb{P}\dot{\psi}_\theta(\theta_0,\eta_0)\right\}^{-1}, \quad \dot{\psi}_\theta(\theta,\eta) = \frac{\partial\psi(\theta,\eta)}{\partial\theta^T},$$
$$\mathbf{\Sigma} = \mathbf{\Sigma}(\eta_0) = \lim_{nT\to\infty}\frac{1}{nT}\sum_{i=1}^n\sum_{t=1}^T \mathbb{E}\left\{\psi_{i,t}(\theta_0,\eta_0)\psi_{i,t}^T(\theta_0,\eta_0)|\overline{V}_{i,t}\right\}.$$

In addition, $\mathbf{B}$ *and* $\mathbf{\Sigma}$ *can be estimated by*

$$\hat{\mathbf{B}} = \left\{\frac{1}{nT}\sum_{i=1}^n\sum_{t=1}^T \frac{\partial\psi_{i,t}(\theta,\hat{\eta})}{\partial\theta^T}\Big|_{\theta=\hat{\theta}}\right\}^{-1} \quad \text{and} \quad \hat{\mathbf{\Sigma}} = \frac{1}{nT}\sum_{i=1}^n\sum_{t=1}^T \psi_{i,t}(\hat{\theta},\hat{\eta})\psi_{i,t}^T(\hat{\theta},\hat{\eta}).$$

Theorem 4.3.2 *If the treatment assignment is cluster randomization, under conditions (C1)-(C3) and (C4')-(C6'), as $T\to\infty$, we have*

$$\sqrt{T}(\hat{\theta}-\theta_0) \xrightarrow{d} MVN\left(0, \mathbf{B}\mathbf{\Sigma}_1\mathbf{B}^T\right),$$

where $\mathbf{B}$ *is defined as the same as in Theorem 4.3.1 and*

$$\boldsymbol{\Sigma}_1 = \boldsymbol{\Sigma}_1(\eta_0) = \lim_{T\to\infty} \frac{1}{T}\sum_{t=1}^{T} \mathbb{E}\left[\left\{\frac{1}{n}\sum_{i=1}^{n}\psi_{i,t}(\theta_0,\eta_0)\right\}\left\{\frac{1}{n}\sum_{i=1}^{n}\psi_{i,t}(\theta_0,\eta_0)\right\}^T \middle| \overline{V}_{\cdot,t}\right].$$

In addition, $\boldsymbol{\Sigma}_1$ *can be estimated by*

$$\hat{\boldsymbol{\Sigma}}_1 = \frac{1}{T}\sum_{t=1}^{T}\left\{\frac{1}{n}\sum_{i=1}^{n}\psi_{i,t}(\hat{\theta},\hat{\eta})\right\}\left\{\frac{1}{n}\sum_{i=1}^{n}\psi_{i,t}(\hat{\theta},\hat{\eta})\right\}^T.$$

The proofs of Theorems 4.3.1 and 4.3.2 are given in Appendix C.1. Both theorems state that the asymptotic variance of the proposed estimator $\hat{\theta}$ can be consistently estimated by the typical sandwich formula without accounting for the variation of $\hat{\eta}$. This is mainly due to the rate double robustness for $\hat{\eta}$ that satisfies Condition (C4) or (C4').

4.4 Simulation Study

We conduct Monte Carlo simulations to examine the finite sample performance of the proposed estimator. We consider three covariates, which are generated as below:

- $L_{i,t}^{(1)} = \log(t)$.
- $L_{i,t}^{(2)} = \log(K_{i,t}^{(2)}+1)$, where $K_{i,t}^{(2)} \sim (1-q_{i,t})I(K_{i,t}^{(2)}=0) + q_{i,t}\,\text{Log-normal}(w_{i,t}, 0.5^2)$ with $q_{i,t} = \{1+\exp(1.1-0.5l_{i,t-1}^{(2)} - l_{i,t}^{(3)} - 0.5a_{i,t-1})\}^{-1}$ and $w_{i,t} = 0.8l_{i,t-1}^{(2)} + 0.6l_{i,t}^{(3)} + 0.25a_{i,t-1}$.
- $L_{i,t}^{(3)} = L_t^{(3)}$, where $L_t^{(3)} \sim \text{Ber}(0.5)$.

The first covariate is to model the time trend effect of treatment; the second covariate is a log transformation of $K_{i,t}$, which mimics the daily activity time in mobile game data and follows a zero-inflated log-normal distribution; the third covariate gives the daily effect of treatment, such as weekday vs. weekend.

The treatment $A_{i,t}$ is generated from a Bernoulli distribution by two means: (i) individual randomization; (ii) cluster randomization, i.e. $A_{i,t} = A_{j,t}\ \forall i \neq j$. In the first scenario, we consider a constant propensity score $\pi_0 = 0.5$ and a covariates-dependent propensity score $\pi_0(\overline{v}_{i,t}) = \{1+\exp(0.5-0.5l_{i,t}^{(2)} - 0.8l_{i,t}^{(3)})\}^{-1}$. We consider various sample sizes n with a fixed follow-up time T. In the second scenario, we only consider a constant propensity score $\pi_0 = 0.5$, with a fixed n and increasing T.

We generate the potential outcomes $Y_{i,t}^{(\overline{a})}$ from a zero-inflated log-normal distribution

$$(1-p_{i,t}^{(\overline{a})})I(Y_{i,t}^{(\overline{a})}=0) + p_{i,t}^{(\overline{a})}\,\text{Log-normal}(\nu_{i,t}^{(\overline{a})}, 0.5^2),$$

where $p_{i,t}^{(\overline{a})} = \exp(\beta_p^T\tilde{l}_{i,t} + \gamma_p^T\tilde{l}_{i,t}\cdot a_{i,t})$ and $\nu_{i,t}^{(\overline{a})} = \beta_\nu^T\tilde{l}_{i,t} + \gamma_\nu^T\tilde{l}_{i,t}\cdot a_{i,t}$ with $\tilde{l}_{i,t} = (1, l_{i,t}^T)^T$. Here, we choose small β_p, γ_p to ensure $p_{i,t}^{(\overline{a})} < 1$. Then, the true $f_{\theta_0}(\cdot)$ in model (4.1) is $f_{\theta_0}(l_{i,t}) = \theta_0^T\tilde{l}_{i,t}$ with $\theta_0 = \gamma_p + \gamma_\nu$. In practice, the treatment may have opposite effects on the proportion of users with positive

outcomes and the conditional mean of positive outcomes. For example, sales on the add-on components in the game may increase the proportion of players purchasing them, but it also decreases the average cost for the users who pay for them. In our simulation, we consider two scenarios: (i) the treatment effect has the same direction on the probability of having positive outcomes and the conditional mean of positive outcomes, i.e. $\gamma_p \geq 0, \gamma_v \geq 0$ or $\gamma_p \leq 0, \gamma_v \leq 0$, where $\geq$ or $\leq$ stands for the point-wise comparison; (ii) the treatment effect has the opposite direction on the probability of having positive outcomes and the conditional mean of positive outcomes, i.e. $\gamma_p \geq 0, \gamma_v \leq 0$ or $\gamma_p \leq 0, \gamma_v \geq 0$. The settings are described as below:

- same direction: $\beta_p = (-0.8, -0.3, 0, 0)^T$, $\beta_v = (0, 0, 0.3, 0.5)^T$, $\gamma_p = (0.2, 0.15, 0, 0)^T$, $\gamma_v = (0.1, 0, 0.05, 0.08)^T$, $\theta_0 = (0.3, 0.15, 0.05, 0.08)^T$;
- opposite direction: $\beta_p = (-0.6, -0.15, 0, 0)^T$, $\beta_v = (0, 0, 0.3, 0.5)^T$, $\gamma_p = (-0.2, -0.15, 0, 0)^T$, $\gamma_v = (0.1, 0, 0.05, 0.08)^T$, $\theta_0 = (-0.1, -0.15, 0.05, 0.08)^T$.

The observed outcome is $Y_{i,t} = Y_{i,t}^{(\overline{A}_i)}$, $i = 1, \cdots, n$ and $t = 1, \cdots, T$.

The parameters θ_0 are estimated by solving the estimating equation (4.3), where the propensity score π_0 is estimated by fitting a correctly specified logistic regression model in individual randomization and by sample proportion in cluster randomization. The conditional means of outcomes μ_0 and μ_1 are estimated nonparametrically by building a generalized additive model on the probability of being positive and the conditional mean for positive outcomes, respectively. In addition, we consider two choices of function c in the estimating equation: an optimal function $c(\overline{v}_{i,t})$ in (4.5), denoted by Opt. c, and a simple function $c(\overline{v}_{i,t}) = (1, l_{i,t}^{(1)}, l_{i,t}^{(2)}, l_{i,t}^{(3)})^T$, denoted by. Sim. c.

For each simulation setting, we conduct 1000 replicates and report the average (Mean) and standard deviation (SD) of the estimates, the average of the estimated standard error (SE) of the estimator using the sandwich formula given in Theorems 4.3.1 and 4.3.2, and the empirical coverage probability (C.P.%) of the 95% Wald-type confidence interval. Tables 4.1 and 4.2 respectively show the results for the settings with the same and opposite direction of treatment effects under individual randomization. Based on the results, we can see that the proposed estimators are nearly unbiased, the average of estimated standard errors is close to the standard deviation of the estimates, and the empirical coverage probability is close to the nominal level. In addition, the estimators obtained using the simple c function show comparable performance compared with those obtained using the optimal c function, with slightly larger standard deviations. We also considered cases with a fixed $n = 100$ and increasing T from 60 to 1000. The results are similar, which are shown in Appendix C.3.

Next, we consider settings with cluster randomization. Tables 4.3 and 4.4 show the results for the settings with the same and opposite directions of treatment effects, respectively. We observe that the proposed estimator is nearly unbiased even with $T = 60$. However, when $T = 60$, the average of the estimated standard errors is smaller than the standard deviation of the estimators for some parameters, like intercept and time trend effect. As such, the empirical coverage probability is lower than the nominal level for these parameters. But as T increases, the average of the estimated standard errors gets much closer to the standard deviation of estimates and the resulting empirical coverage probability is close

Table 4.1 Simulation results with same directional treatment effect and $T = 60$ under individual randomization

Propensity score	n	Parameter	Opt. c				Sim. c			
			Mean	SD	SE	C.P.	Mean	SD	SE	C.P.
constant	1000	$\theta_1 = 0.30$	0.300	0.053	0.052	94	0.299	0.057	0.056	94
		$\theta_2 = 0.15$	0.150	0.016	0.016	95	0.150	0.017	0.017	95
		$\theta_3 = 0.05$	0.050	0.017	0.017	96	0.050	0.019	0.019	96
		$\theta_4 = 0.08$	0.081	0.035	0.035	95	0.081	0.038	0.037	95
	400	$\theta_1 = 0.30$	0.302	0.087	0.082	94	0.302	0.093	0.089	95
		$\theta_2 = 0.15$	0.150	0.026	0.025	95	0.149	0.028	0.027	95
		$\theta_3 = 0.05$	0.050	0.027	0.027	96	0.051	0.029	0.030	96
		$\theta_4 = 0.08$	0.080	0.057	0.055	94	0.081	0.061	0.059	94
	200	$\theta_1 = 0.30$	0.299	0.116	0.117	95	0.304	0.127	0.126	94
		$\theta_2 = 0.15$	0.150	0.035	0.036	97	0.149	0.038	0.038	96
		$\theta_3 = 0.05$	0.052	0.039	0.039	95	0.052	0.043	0.043	95
		$\theta_4 = 0.08$	0.078	0.079	0.078	95	0.077	0.084	0.083	94
covariates dependent	1000	$\theta_1 = 0.30$	0.302	0.057	0.055	95	0.303	0.063	0.061	95
		$\theta_2 = 0.15$	0.149	0.017	0.017	94	0.149	0.019	0.019	94
		$\theta_3 = 0.05$	0.050	0.020	0.020	95	0.049	0.023	0.023	95
		$\theta_4 = 0.08$	0.082	0.039	0.038	95	0.081	0.043	0.042	95
	400	$\theta_1 = 0.30$	0.304	0.089	0.087	94	0.306	0.101	0.097	94
		$\theta_2 = 0.15$	0.148	0.027	0.027	95	0.148	0.030	0.029	94
		$\theta_3 = 0.05$	0.050	0.032	0.031	95	0.050	0.036	0.036	95
		$\theta_4 = 0.08$	0.081	0.062	0.061	95	0.081	0.070	0.067	94
	200	$\theta_1 = 0.30$	0.303	0.121	0.123	96	0.306	0.137	0.137	95
		$\theta_2 = 0.15$	0.148	0.037	0.038	96	0.148	0.041	0.042	95
		$\theta_3 = 0.05$	0.052	0.045	0.044	94	0.052	0.052	0.050	94
		$\theta_4 = 0.08$	0.084	0.089	0.086	94	0.084	0.100	0.095	93

to the nominal level. The estimators obtained using the simple c function and optimal c function show comparable efficiency as in the individual randomization settings. These results demonstrate the validity of our inference procedure for both individual randomization and cluster randomization settings.

4.5 Application to Freemium Mobile Game Data

We apply our method to a real dataset from a freemium mobile game [Ban19], where players fight each other with robot avatars and gain the in-game reward by winning the game. Players can use these in-game rewards or real money to upgrade their robots, improve their fighting equipment, or acquire fancy game themes and backgrounds. Apart from the direct in-app purchase, there are several indirect ways of monetizing freemium games by engaging their free users in promoting the games through social

Table 4.2 Simulation results with opposite directional treatment effect and $T = 60$ under individual randomization

Propensity score	n	Parameter	Opt. c				Sim. c			
			Mean	SD	SE	C.P.	Mean	SD	SE	C.P.
constant	1000	$\theta_1 = -0.10$	-0.102	0.053	0.052	94	-0.100	0.060	0.059	94
		$\theta_2 = -0.15$	-0.149	0.016	0.016	95	-0.150	0.018	0.018	95
		$\theta_3 = 0.05$	0.051	0.018	0.017	95	0.050	0.020	0.019	94
		$\theta_4 = 0.08$	0.080	0.035	0.035	95	0.079	0.038	0.038	95
	400	$\theta_1 = -0.10$	-0.099	0.082	0.083	94	-0.099	0.094	0.094	95
		$\theta_2 = -0.15$	-0.150	0.026	0.026	95	-0.150	0.028	0.028	95
		$\theta_3 = 0.05$	0.050	0.027	0.028	95	0.049	0.031	0.030	95
		$\theta_4 = 0.08$	0.082	0.053	0.056	95	0.082	0.057	0.060	96
	200	$\theta_1 = -0.10$	-0.095	0.114	0.117	95	-0.090	0.130	0.133	95
		$\theta_2 = -0.15$	-0.152	0.034	0.036	96	-0.153	0.038	0.040	96
		$\theta_3 = 0.05$	0.051	0.038	0.039	95	0.050	0.043	0.043	94
		$\theta_4 = 0.08$	0.081	0.080	0.079	94	0.080	0.086	0.084	95
covariates dependent	1000	$\theta_1 = -0.10$	-0.102	0.051	0.054	95	-0.102	0.056	0.058	95
		$\theta_2 = -0.15$	-0.150	0.016	0.017	96	-0.150	0.017	0.018	96
		$\theta_3 = 0.05$	0.050	0.018	0.018	96	0.050	0.018	0.018	95
		$\theta_4 = 0.08$	0.082	0.035	0.036	96	0.082	0.037	0.038	96
	400	$\theta_1 = -0.10$	-0.105	0.083	0.085	95	-0.106	0.090	0.090	96
		$\theta_2 = -0.15$	-0.149	0.026	0.026	96	-0.149	0.027	0.028	95
		$\theta_3 = 0.05$	0.050	0.028	0.028	95	0.051	0.029	0.029	96
		$\theta_4 = 0.08$	0.084	0.059	0.058	94	0.084	0.061	0.060	94
	200	$\theta_1 = -0.10$	-0.107	0.120	0.121	96	-0.107	0.130	0.130	95
		$\theta_2 = -0.15$	-0.149	0.037	0.037	95	-0.148	0.039	0.039	95
		$\theta_3 = 0.05$	0.050	0.039	0.039	94	0.049	0.041	0.041	95
		$\theta_4 = 0.08$	0.085	0.082	0.082	94	0.085	0.086	0.085	95

media, for example, inviting friends on Facebook. As in Banerjee et al. [Ban19], the *daily engagement* of a player is defined as the combination of his/her in-app purchase (direct) and varied involvement with the game on social media (indirect) during the day. *Daily activity* is another metric of interest in mobile games, which measures the time a player spends playing the game during the day. Generally speaking, positive daily activity does not always imply positive daily engagement, but positive and high daily engagement often follows the persistent and incremental positive daily activity [Ban19].

There are 38,860 players in the dataset whose daily engagement and daily activity were tracked for 60 consecutive days starting from the release date of the game. Approximately 51.8% of the players have positive daily engagement (direct or indirect) during the follow-up period. All players received the same history of promotion decisions (cluster randomization), which alternatively include 40 days of promotions and 20 days of no promotions. We want to model the effect of a sequence of promotion

Table 4.3 Simulation results with same directional treatment effect and $n = 400$ under cluster randomization

T	Parameter	Opt. c				Sim. c			
		Mean	SD	SE	C.P.	Mean	SD	SE	C.P.
1000	$\theta_1 = 0.30$	0.300	0.057	0.055	93	0.301	0.053	0.052	95
	$\theta_2 = 0.15$	0.150	0.009	0.009	92	0.150	0.009	0.009	94
	$\theta_3 = 0.05$	0.050	0.011	0.011	95	0.050	0.012	0.012	95
	$\theta_4 = 0.08$	0.080	0.020	0.021	96	0.079	0.021	0.022	96
250	$\theta_1 = 0.30$	0.298	0.089	0.075	86	0.300	0.080	0.072	90
	$\theta_2 = 0.15$	0.151	0.018	0.016	88	0.150	0.017	0.015	91
	$\theta_3 = 0.05$	0.050	0.020	0.020	94	0.050	0.021	0.020	94
	$\theta_4 = 0.08$	0.080	0.035	0.035	95	0.079	0.037	0.036	94
60	$\theta_1 = 0.30$	0.293	0.136	0.098	81	0.296	0.121	0.098	87
	$\theta_2 = 0.15$	0.150	0.038	0.028	81	0.150	0.034	0.029	89
	$\theta_3 = 0.05$	0.052	0.038	0.036	92	0.052	0.038	0.036	92
	$\theta_4 = 0.08$	0.085	0.060	0.060	95	0.083	0.064	0.061	93

Table 4.4 Simulation results with opposite directional treatment effect and $n = 400$ under cluster randomization

T	Parameter	Opt. c				Sim. c			
		Mean	SD	SE	C.P.	Mean	SD	SE	C.P.
1000	$\theta_1 = -0.10$	-0.100	0.053	0.052	95	-0.099	0.055	0.054	95
	$\theta_2 = -0.15$	-0.150	0.009	0.009	95	-0.150	0.009	0.009	94
	$\theta_3 = 0.05$	0.050	0.011	0.011	95	0.050	0.012	0.012	95
	$\theta_4 = 0.08$	0.080	0.020	0.021	96	0.080	0.022	0.022	95
250	$\theta_1 = -0.10$	-0.102	0.083	0.073	90	-0.100	0.080	0.075	94
	$\theta_2 = -0.15$	-0.150	0.017	0.015	92	-0.150	0.017	0.016	93
	$\theta_3 = 0.05$	0.050	0.021	0.020	94	0.050	0.021	0.020	94
	$\theta_4 = 0.08$	0.080	0.035	0.035	95	0.079	0.037	0.037	94
60	$\theta_1 = -0.10$	-0.097	0.140	0.100	81	-0.096	0.132	0.103	86
	$\theta_2 = -0.15$	-0.151	0.040	0.029	84	-0.151	0.037	0.030	88
	$\theta_3 = 0.05$	0.052	0.037	0.036	94	0.050	0.038	0.036	93
	$\theta_4 = 0.08$	0.077	0.061	0.060	93	0.076	0.066	0.062	92

decisions on the daily engagement in the presence of other time-varying variables, for example, daily activity.

We consider three covariates at time t: the number of days t in the study, the daily activity on the previous day $l_{i,t}^{(2)}$, and the weekend indicator $l_t^{(3)}$. We fit model (4.1) with the function $f_\theta(\cdot)$ given by

$$f_{\theta_0}(\overline{v}_{i,t}) = \theta_0^{(0)} + \theta_0^{(1)} \log t + \theta_0^{(2)} \log(l_{i,t}^{(2)} + 1) + \theta_0^{(3)} l_t^{(3)}.$$

Our proposed estimator of θ_0 is $\hat{\theta} = (-0.77, 0.50, -0.26, 0.05)^T$, which is obtained based on the simple c function. Its standard error estimate is $(0.121, 0.057, 0.031, 0.166)^T$, which is obtained using the proposed sandwich formula in Theorem 4.3.2. Based on the results, we make a few observations. First, promotions have a significant positive interaction effect with time on daily engagement. In general, users may lose interest in the game over time. This implies that promotions over time are able to provide incentives and grab their attention again. Second, promotions have smaller effects in terms of increasing the engagement for active users (i.e. those with large daily activity) than for less active users, because the estimator of $\theta_0^{(2)}$ is negative with a significant p-value. This finding is coherent with the literature since active users tend to be less sensitive to the promotion compared with less active users in terms of engagement. Last, the promotion effects on weekends and weekdays have no much difference.

To visualize the treatment effects, we plot the estimated log-ratio treatment effect $f_{\hat{\theta}}(\overline{v}_t)$ as a function of t and log of daily activity plus 1, stratified by the weekend vs. weekdays (Figures 4.2a and 4.2b). Here, $f_{\hat{\theta}}(\overline{v}_t) > 0$ indicates a positive treatment effect of promotions on daily engagement, while $f_{\hat{\theta}}(\overline{v}_t) < 0$ implies a negative effect. In addition, we consider the curve with $f_{\hat{\theta}}(\overline{v}_t) = 0$ and computed its 95% point-wise confidence intervals based on the variance estimates of $\hat{\theta}$. We plot this curve and its 95% point-wise confidence intervals in Figures 4.2c and 4.2d for weekends and weekdays, respectively. The "+" indicates the area where the treatment effect is significantly positive while "-" represents the region that has significantly negative treatment effects.

4.6 Concluding Remarks

In this chapter, we propose a class of multiplicative structural nested mean models for zero-inflated nonnegative outcomes and develop a doubly robust estimation method that effectively uses the zero-inflated structure for estimating conditional mean outcomes. The asymptotic properties of the proposed estimator are established under both individual randomization and cluster randomization. In particular, we establish the bidirectional asymptotic results under individual randomization, but the limiting distribution under cluster randomization requires the follow-up time going to infinity. Furthermore, the asymptotic variance for the proposed estimator can be estimated using the typical sandwich formula without accounting for the variation of the estimation of nuisance functions and is thus easy to implement.

The multiplicative structural nested mean model describes the treatment effect on the mean shift of potential outcomes, where the shift can result from (i) the change of the probability of having positive outcomes or/and (ii) the change of the conditional mean given a positive outcome. However, the proposed model cannot distinguish and quantify these two pathways of treatment effects. In the point treatment setting, two-part models are used to model the treatment effect on (i) and (ii) separately; however, controversy exists regarding the causal interpretation of (ii) since it involves conditioning on a post-

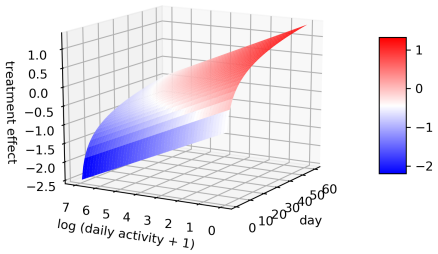

(a) $f_{\hat{\theta}}(\overline{v}_t)$ on weekends

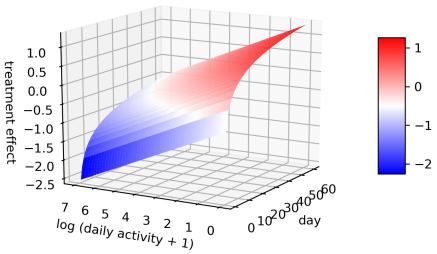

(b) $f_{\hat{\theta}}(\overline{v}_t)$ on weekdays

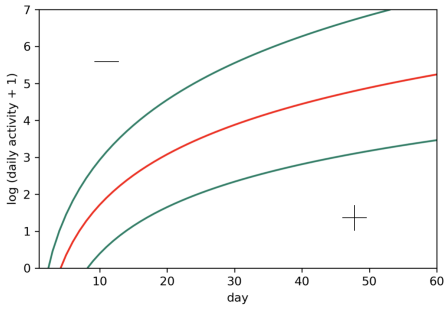

(c) The curve for $f_{\hat{\theta}}(\overline{v}_t)=0$ (red) and its 95% point-wise confidence intervals (green) for weekends

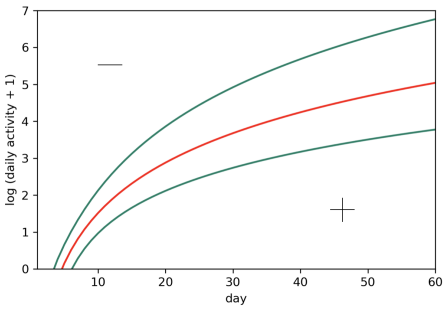

(d) The curve for $f_{\hat{\theta}}(\overline{v}_t)=0$ (red) and its 95% point-wise confidence intervals (green) for weekdays

Figure 4.2 Results of the multiplicative structural nested mean model in Freemium Mobile Game Data

treatment variable, i.e. the potential outcome being positive. Although the proposed model only provides a coarsen description of the treatment effect, it avoids the controversy associated with the two-part model.

It is important to note that the proposed structural nested mean models are designed to address the causal effects of a series of treatments. We must add a caveat that our identification strategy assumes that all time-varying confounders are measured, which however is not verifiable based on the observed data but rely on subject matter experts to assess their plausibility. If some confounders have not been captured in our data, it can distort the causal interpretation. In these cases, one can conduct sensitivity analyses [Yan18] to assess the impact of possible uncontrolled confounding.

In our current work, the ratio of conditional means of potential outcomes at time t is assumed to depend only on the current treatment a_t. We can also consider an elaborate model that accommodates not only the current but also previous treatments in the following form

$$\exp\{f_{\theta_0}^{(t)}(\overline{v}_t)\cdot a_t + f_{\theta_1}^{(t-1)}(\overline{v}_{t-1})\cdot a_{t-1} + \cdots + f_{\theta_K}^{(t-K)}(\overline{v}_{t-K})\cdot a_{t-K}\},$$

where $f_{\theta_k}^{(t-k)}(\cdot)$, $k=0,1,\ldots,K$, are known functions and $\theta=(\theta_0^T,\theta_1^T,\cdots,\theta_K^T)^T$ are the parameters of interest. Our framework requires the multiplicative structural nested mean model to be correctly specified, thus model assessment is critically important. The key insight is that we have a larger number of estimating

functions than the number of parameters by varying c in equation (4.3), leading to the over-identification of model parameters. A goodness-of-fit test can be developed based on over-identification restrictions for model diagnosis [Yan16].

BIBLIOGRAPHY

[Abh17] Abhishek, V. & Mannor, S. "A nonparametric sequential test for online randomized experiments". *Proceedings of the 26th International Conference on World Wide Web Companion*. International World Wide Web Conferences Steering Committee. 2017, pp. 610–616.

[And09] Anderson, C. *Free: The Future of a Radical Price*. Random House, 2009.

[And88] Andrews, D. W. "Laws of large numbers for dependent non-identically distributed random variables". *Econometric Theory* **4**.3 (1988), pp. 458–467.

[App20] AppBrain. *Free vs. paid Android apps*. 2020. url: `https://www.appbrain.com/stats/free-and-paid-android-applications`.

[Arm60] Armitage, P. et al. *Sequential Medical Trials*. Blackwell Scientific Publications, 1960.

[Arm69] Armitage, P. et al. "Repeated significance tests on accumulating data". *Journal of the Royal Statistical Society: Series A (General)* **132**.2 (1969), pp. 235–244.

[Bae10] Bae, J.-S. et al. "The uniform CLT for martingale difference arrays under the uniformly integrable entropy". *Bulletin of the Korean Mathematical Society* **47**.1 (2010), pp. 39–51.

[Bal15] Balsubramani, A. & Ramdas, A. "Sequential nonparametric testing with the law of the iterated logarithm". *arXiv preprint arXiv:1506.03486* (2015).

[Ban19] Banerjee, T. et al. "A large-scale constrained joint modeling approach for predicting user activity, engagement, and churn with application to freemium mobile games". *Journal of the American Statistical Association* **115**.530 (2019), pp. 538–554.

[Ben95] Benjamini, Y. & Hochberg, Y. "Controlling the false discovery rate: a practical and powerful approach to multiple testing". *Journal of the Royal Statistical Society: Series B (Methodological)* **57**.1 (1995), pp. 289–300.

[Ben01] Benjamini, Y., Yekutieli, D., et al. "The control of the false discovery rate in multiple testing under dependency". *The Annals of Statistics* **29**.4 (2001), pp. 1165–1188.

[Boo13] Boos, D. D. & Stefanski, L. "Large sample results for likelihood-based methods". *Essential Statistical Inference*. Springer, 2013, pp. 275–293.

[Bou19] Boudreau, K. et al. *Competing on free (mium): digital competition with network effects*. 2019. url: `https://ssrn.com/abstract=2984546`.

[Cha94] Chang, M. et al. "Reduction in burden of illness: a new efficacy measure for prevention trials". *Statistics in Medicine* **13**.18 (1994), pp. 1807–1814.

[Che20] Cheng, J. & Small, D. S. "Semiparametric models and inference for the effect of a treatment when the outcome is nonnegative with clumping at zero". *Biometrics* (in press) (2020).

[Cox63] Cox, D. R. "Large sample sequential tests for composite hypotheses". *Sankhyā: The Indian Journal of Statistics, Series A* **25**.1 (1963), pp. 5–12.

[Den16] Deng, A. et al. "Concise summarization of heterogeneous treatment effect using total variation regularized regression". *arXiv preprint arXiv:1610.03917* (2016).

[Dua83] Duan, N. et al. "A comparison of alternative models for the demand for medical care". *Journal of Business & Economic Statistics* **1**.2 (1983), pp. 115–126.

[Far21] Farrell, M. H. et al. "Deep neural networks for estimation and inference". *Econometrica* **89**.1 (2021), pp. 181–213.

[Gir46] Girshick, M. A. "Contributions to the theory of sequential analysis. I". *The Annals of Mathematical Statistics* **17**.2 (1946), pp. 123–143.

[Goo14] Goodson, M. *Most winning A/B test results are illusory*. 2014. url: http://www.datascienceassn.org/sites/default/files/Most%20Winning%20A-B%20Test%20Results%20are%20Illusory.pdf.

[GL83] Gordon Lan, K. & DeMets, D. L. "Discrete sequential boundaries for clinical trials". *Biometrika* **70**.3 (1983), pp. 659–663.

[Gre19] Green, A. *Mobile gaming generated 60% of the global video games revenue in 2019*. 2019. url: https://goldencasinonews.com/blog/2019/12/30/mobile-gaming-generated-60-of-the-global-video-games-revenue-in-2019.

[Gri17] Grimmer, J. et al. "Estimating heterogeneous treatment effects and the effects of heterogeneous treatments with ensemble methods". *Political Analysis* **25**.4 (2017), pp. 413–434.

[Gri01] Grimmett, G. et al. *Probability and Random Processes*. Oxford University Press, 2001.

[Hor09] Horowitz, J. L. *Semiparametric and Nonparametric Methods in Econometrics*. Vol. 12. Springer, 2009.

[Hsu96] Hsu, J. *Multiple Comparisons: Theory and Methods*. Chapman and Hall/CRC, 1996.

[Jac15] Jacobs, H. *Gaming guru explains why 'freemium' is actually the best business model for multiplayer video games*. 2015. url: https://www.businessinsider.com/sean-plott-explains-why-he-thinks-freemium-games-are-the-best-business-model-for-both-players-and-developers-2015-3.

[Joh15] Johari, R. et al. "Always valid inference: Bringing sequential analysis to A/B testing". *arXiv preprint arXiv:1512.04922* (2015).

[Joh17] Johari, R. et al. "Peeking at A/B tests: Why it matters, and what to do about it". *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM. 2017, pp. 1517–1525.

[Ju19] Ju, N. et al. “A sequential test for selecting the better variant: Online A/B testing, adaptive allocation, and continuous monitoring”. *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*. ACM. 2019, pp. 492–500.

[Kal20] Kallus, N. & Mao, X. “On the role of surrogates in the efficient estimation of treatment effects with limited outcome data”. *arXiv preprint arXiv:2003.12408* (2020).

[Kap19] Kaplan, O. *Mobile gaming is a $68.5 billion global business, and investors are buying in*. 2019. url: `https://techcrunch.com/2019/08/22/mobile-gaming-mints-money`.

[Kee19] Keele, L. & Miratrix, L. “Randomization inference for outcomes with clumping at zero”. *The American Statistician* **73**.2 (2019), pp. 141–150.

[Koh11] Kohavi, R. & Longbotham, R. “Unexpected results in online controlled experiments”. *ACM SIGKDD Explorations Newsletter* **12**.2 (2011), pp. 31–35.

[Koh09] Kohavi, R. et al. “Controlled experiments on the web: survey and practical guide”. *Data Mining and Knowledge Discovery* **18**.1 (2009), pp. 140–181.

[Koh12] Kohavi, R. et al. “Trustworthy online controlled experiments: Five puzzling outcomes explained”. *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 2012, pp. 786–794.

[Koh13] Kohavi, R. et al. “Online controlled experiments at large scale”. *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 2013, pp. 1168–1176.

[Kos07] Kosorok, M. R. *Introduction to Empirical Processes and Semiparametric Inference*. Springer Science & Business Media, 2007.

[Kul11] Kulldorff, M. et al. “A maximized sequential probability ratio test for drug and vaccine safety surveillance”. *Sequential Analysis* **30**.1 (2011), pp. 58–78.

[Lai88] Lai, T. L. “Nearly optimal sequential tests of composite hypotheses”. *The Annals of Statistics* **16**.2 (1988), pp. 856–886.

[Lai01] Lai, T. L. “Sequential analysis: some classical problems and new challenges”. *Statistica Sinica* **11**.2 (2001), pp. 303–351.

[Lai94] Lai, T. L. & Zhang, L. “A modification of Schwarz’s sequential likelihood ratio tests in multivariate sequential analysis”. *Sequential Analysis* **13**.2 (1994), pp. 79–96.

[Leh06] Lehmann, E. L. & Romano, J. P. *Testing Statistical Hypotheses*. Springer Science & Business Media, 2006.

[Lor76] Lorden, G. “2-SPRT’s and the modified Kiefer-Weiss problem of minimizing an expected sample size”. *The Annals of Statistics* **4**.2 (1976), pp. 281–291.

[Lu19] Lu, S. *Beyond A/B testing: Multi-armed bandit experiments*. 2019. url: https://towardsdatascience.com/beyond-a-b-testing-multi-armed-bandit-experiments-1493f709f804.

[Lue16] Luedtke, A. R. & Van Der Laan, M. J. "Statistical inference for the mean outcome under a possibly non-unique optimal treatment strategy". *The Annals of Statistics* **44**.2 (2016), pp. 713–742.

[Mal17] Malek, A. et al. "Sequential multiple hypothesis testing with type I error control". *Proceedings of Artificial Intelligence and Statistics*. 2017, pp. 1468–1476.

[McD17] McDonald, E. *The global games market will reach $108.9 billion in 2017 with mobile taking 42%*. 2017. url: https://newzoo.com/insights/articles/the-global-%20games-market-%20will-reach-108-%209-billion-%20in-2017-%20with-mobile-%20taking-42.

[McL74] McLeish, D. L. et al. "Dependent central limit theorems and invariance principles". *The Annals of Probability* **2**.4 (1974), pp. 620–628.

[Mil15] Miller, E. *Simple sequential A/B testing*. 2015. url: https://www.evanmiller.org/sequential-ab-testing.html.

[Mil66] Miller R.G., J. *Simultaneous Statistical Inference*. New York: McGraw-Hill Book Co., 1966.

[O'B79] O'Brien, P. C. & Fleming, T. R. "A multiple testing procedure for clinical trials". *Biometrics* **35**.3 (1979), pp. 549–556.

[Pek15] Pekelis, L. et al. *The new stats engine*. 2015. url: http://pages.optimizely.com/rs/optimizely/images/stats_engine_technical_paper.pdf.

[Poc77] Pocock, S. J. "Group sequential methods in the design and analysis of clinical trials". *Biometrika* **64**.2 (1977), pp. 191–199.

[Pol78] Pollak, M. "Optimality and almost optimality of mixture stopping rules". *The Annals of Statistics* **6**.4 (1978), pp. 910–916.

[Pow86] Powell, J. L. "Symmetrically trimmed least squares estimation for Tobit models". *Econometrica: Journal of the Econometric Society* **54**.6 (1986), pp. 1435–1460.

[Rob74] Robbins, H & Siegmund, D. "The expected sample size of some tests of power one". *The Annals of Statistics* **2**.3 (1974), pp. 415–436.

[Rob70] Robbins, H. "Statistical methods related to the law of the iterated logarithm". *The Annals of Mathematical Statistics* **41**.5 (1970), pp. 1397–1409.

[Rob94a] Robins, J. M. "Correcting for non-compliance in randomized trials using structural nested mean models". *Communications in Statistics-Theory and Methods* **23**.8 (1994), pp. 2379–2412.

[Rob00] Robins, J. M. "Marginal structural models versus structural nested models as tools for causal inference". *Statistical models in epidemiology, the environment, and clinical trials*. Springer, 2000, pp. 95–133.

[Rob09] Robins, J. M. & Hernán, M. A. "Estimation of the causal effects of time-varying exposures". *Advances in Longitudinal Data Analysis*. Ed. by Fitzmaurice, G. et al. Chapman and Hall/CRC Press, 2009. Chap. 23, pp. 533–599.

[Rob92] Robins, J. M. et al. "G-estimation of the effect of prophylaxis therapy for Pneumocystis carinii pneumonia on the survival of AIDS patients". *Epidemiology* **3**.4 (1992), pp. 319–336.

[Rob94b] Robins, J. M. et al. "Estimation of regression coefficients when some regressors are not always observed". *Journal of the American Statistical Association* **89**.427 (1994), pp. 846–866.

[Sch62] Schwarz, G. "Asymptotic shapes of Bayes sequential testing regions". *The Annals of Mathematical Statistics* **33**.1 (1962), pp. 224–236.

[Sco13] Scott, S. L. *Multi-armed bandit experiments*. 2013. url: `https://analytics.googleblog.com/2013/01/multi-armed-bandit-experiments.html`.

[Sim11] Simmons, J. P. et al. "False-positive psychology: Undisclosed flexibility in data collection and analysis allows presenting anything as significant". *Psychological Science* **22**.11 (2011), pp. 1359–1366.

[Swr14] Swrve. *Swrve finds 0.15% of mobile gamers contribute 50% of all in-game revenue*. 2014. url: `https://www.swrve.com/company/press/swrve-finds-015-of-mobile-gamers-contribute-50-of-all-in-game-revenue`.

[Tan10] Tang, D. et al. "Overlapping experiment infrastructure: More, better, faster experimentation". *Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 2010, pp. 17–26.

[Tob58] Tobin, J. "Estimation of relationships for limited dependent variables". *Econometrica: Journal of the Econometric Society* **26**.1 (1958), pp. 24–36.

[Tro11] Tropp, J. et al. "Freedman's inequality for matrix martingales". *Electronic Communications in Probability* **16** (2011), pp. 262–270.

[VDV96] Van Der Vaart, A. W. & Wellner, J. A. *Weak Convergence and Empirical Processes*. Springer, 1996.

[Van14] Vansteelandt, S., Joffe, M., et al. "Structural nested models and G-estimation: the partially realized promise". *Statistical Science* **29**.4 (2014), pp. 707–731.

[Wag18] Wager, S. & Athey, S. "Estimation and inference of heterogeneous treatment effects using random forests". *Journal of the American Statistical Association* **113**.523 (2018), pp. 1228–1242.

[Wal45] Wald, A. “Sequential tests of statistical hypotheses”. *The Annals of Mathematical Statistics* **16**.2 (1945), pp. 117–186.

[Wal48] Wald, A. & Wolfowitz, J. “Optimum character of the sequential probability ratio test”. *The Annals of Mathematical Statistics* **19**.3 (1948), pp. 326–339.

[Wan87] Wang, S. K. & Tsiatis, A. A. “Approximately optimal one-parameter boundaries for group sequential trials”. *Biometrics* **43**.1 (1987), pp. 193–199.

[Wei62] Weiss, L. “On sequential tests which minimize the maximum expected sample size”. *Journal of the American Statistical Association* **57**.299 (1962), pp. 551–566.

[Yan16] Yang, S & Lok, J. “A goodness-of-fit test for structural nested mean models”. *Biometrika* **103**.3 (2016), pp. 734–741.

[Yan18] Yang, S. & Lok, J. J. “Sensitivity analysis for unmeasured confounding in coarse structural nested mean models”. *Statistica Sinica* **28**.4 (2018), pp. 1703–1723.

[Yu20] Yu, M. et al. “A new framework for online testing of heterogeneous treatment effect”. *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 34. 06. 2020, pp. 10310–10317.

[Zha12] Zhang, B. et al. “A robust method for estimating optimal treatment regimes”. *Biometrics* **68**.4 (2012), pp. 1010–1018.

[Øy14] Øygard, A. M. *Rapid A/B testing with sequential analysis*. 2014. url: `https://www.auduno.com/2014/12/25/rapid-a-b-testing-with-sequential-analysis`.

APPENDICES

APPENDIX

A

SUPPLEMENT TO CHAPTER 2

A.1 Proof of Main Results

A.1.1 Proof of Lemma 2.3.1

From now on, we assume $a_i(\phi) = a(\phi) = 1$ for simplicity and use $\mathbf{S}_n^{(0)}(\boldsymbol{\beta})$ and $\mathbf{S}_n^{(1)}(\boldsymbol{\beta})$ in replace of $\mathbf{S}_{n,\boldsymbol{\beta}}^{(0)}(\boldsymbol{\beta})$ and $\mathbf{S}_{n,\boldsymbol{\theta}}^{(1)}(\boldsymbol{\beta},\boldsymbol{\theta}_0)$ for ease of expression:

$$\mathbf{S}_n^{(0)}(\boldsymbol{\beta}) = \sum_{i=1}^{n}\left[\frac{\partial \mu_i^{(0)}(\boldsymbol{\beta})}{\partial \boldsymbol{\beta}^T}\cdot\frac{\{Y_i^{(0)}-\mu_i^{(0)}(\boldsymbol{\beta})\}}{V_i^{(0)}(\boldsymbol{\beta})}\right]$$
$$\mathbf{S}_n^{(1)}(\boldsymbol{\beta}) = \sum_{i=1}^{n}\left[\frac{\partial \mu_i^{(1)}(\boldsymbol{\theta},\boldsymbol{\beta})}{\partial \boldsymbol{\theta}^T}\cdot\frac{\{Y_i^{(1)}-\mu_i^{(1)}(\boldsymbol{\theta},\boldsymbol{\beta})\}}{V_i^{(1)}(\boldsymbol{\theta},\boldsymbol{\beta})}\right]\Bigg|_{\boldsymbol{\theta}=\boldsymbol{\theta}_0}.$$

All other notations can find its definition in Section 2.2 and Section 2.3.

The maximum likelihood estimator $\hat{\boldsymbol{\beta}}_n$ based only on data from control group (A=0) is obtained by solving the equation $\mathbf{S}_n^{(0)}(\boldsymbol{\beta}) = 0$. Under some regularity conditions [Boo13] which are certainly satisfied by the generalized linear model, the maximum likelihood estimator $\hat{\boldsymbol{\beta}}_n$ has the following property:

$$\sqrt{n}(\hat{\boldsymbol{\beta}}_n - \boldsymbol{\beta}) \xrightarrow{d} \mathbf{MVN}_{p+1}(\mathbf{0}, \mathbf{I}^{(0)}(\boldsymbol{\beta})),$$

where $\mathbf{I}^{(0)}(\boldsymbol{\beta}) = \mathbb{E}_{(\mathbf{X},\mathbf{Y})}\left\{\bar{\mathbf{I}}_n^{(0)}(\boldsymbol{\beta})\right\} = \mathbb{E}_{(\mathbf{X},\mathbf{Y})}\left[\left\{\frac{\partial \mu_1^{(0)}(\boldsymbol{\beta})}{\partial \boldsymbol{\beta}^T}\cdot\frac{\partial \mu_1^{(0)}(\boldsymbol{\beta})}{\partial \boldsymbol{\beta}}\right\}\Big/ V_1^{(0)}(\boldsymbol{\beta})\right]$. It follows that $\hat{\boldsymbol{\beta}}_n - \boldsymbol{\beta} = \boldsymbol{O}_P(n^{-1/2})$.

Thus, $\mathbf{S}_n^{(0)}(\hat{\boldsymbol{\beta}}_n)$ can be written as

$$\mathbf{0} = \mathbf{S}_n^{(0)}(\hat{\boldsymbol{\beta}}_n) = \mathbf{S}_n^{(0)}(\boldsymbol{\beta}) - n\bar{\mathbf{I}}_n^{(0)}(\boldsymbol{\beta})(\hat{\boldsymbol{\beta}}_n - \boldsymbol{\beta}) + \boldsymbol{O}_P(1)$$

by making a Taylor expansion around $\boldsymbol{\beta}$. Applying the *weak law of large number* to $\bar{\mathbf{I}}_n^{(0)}(\boldsymbol{\beta})$, we have

$$\hat{\boldsymbol{\beta}}_n - \boldsymbol{\beta} = \left\{\mathbf{I}^{(0)}(\boldsymbol{\beta})\right\}^{-1} \frac{1}{n}\mathbf{S}_n^{(0)}(\boldsymbol{\beta}) + \boldsymbol{o}_P(n^{-\frac{1}{2}}). \tag{A.1}$$

In the same way, we can get

$$\begin{aligned}
\bar{\mathbf{S}}_n &= \frac{1}{n}\mathbf{S}_n^{(1)}(\boldsymbol{\beta}) - \bar{\mathbf{I}}_n^{(1)}(\boldsymbol{\beta})(\hat{\boldsymbol{\beta}}_n - \boldsymbol{\beta}) + \boldsymbol{O}_P(n^{-1}) && \text{(A.2)}\\
&= \frac{1}{n}\mathbf{S}_n^{(1)}(\boldsymbol{\beta}) - \mathbf{I}^{(1)}(\boldsymbol{\beta})(\hat{\boldsymbol{\beta}}_n - \boldsymbol{\beta}) + \boldsymbol{o}_P(n^{-\frac{1}{2}}) && \text{(A.3)}\\
&\overset{(A.1)}{=} \frac{1}{n}\mathbf{S}_n^{(1)}(\boldsymbol{\beta}) - \mathbf{I}^{(1)}(\boldsymbol{\beta})\left\{\mathbf{I}^{(0)}(\boldsymbol{\beta})\right\}^{-1}\frac{1}{n}\mathbf{S}_n^{(0)}(\boldsymbol{\beta}) + \boldsymbol{o}_P(n^{-\frac{1}{2}}), && \text{(A.4)}
\end{aligned}$$

where $\mathbf{I}^{(1)}(\boldsymbol{\beta}) = \mathbb{E}_{(\mathbf{X},\mathbf{Y})}\left\{\bar{\mathbf{I}}_n^{(1)}(\boldsymbol{\beta})\right\} = \mathbb{E}_{(\mathbf{X},\mathbf{Y})}\left[\left\{\frac{\partial \mu_1^{(1)}(\boldsymbol{\theta},\boldsymbol{\beta})}{\partial \boldsymbol{\beta}^T}\cdot\frac{\partial \mu_1^{(1)}(\boldsymbol{\theta},\boldsymbol{\beta})}{\partial \boldsymbol{\theta}}\right\}\Big/ V_1^{(1)}(\boldsymbol{\theta},\boldsymbol{\beta})\right]\Bigg|_{\boldsymbol{\theta}=\boldsymbol{\theta}_0}$.

Since the first two terms in (A.4) are independent, we can derive their asymptotic distributions separately. By the *central limit theorem*, the pivot quantity in the second term converges to the following distribution:

$$\frac{1}{\sqrt{n}}\mathbf{S}_n^{(0)}(\boldsymbol{\beta}) \xrightarrow{d} \mathbf{MVN}_{p+1}(\mathbf{0}, \mathbf{I}^{(0)}(\boldsymbol{\beta})). \tag{A.5}$$

Similarly, the standardized first term converges to a similar distribution under the null hypothesis $\mathrm{H}_0 : \boldsymbol{\theta} = \boldsymbol{\theta}_0$:

$$\frac{1}{\sqrt{n}}\mathbf{S}_n^{(1)}(\boldsymbol{\beta}) \xrightarrow[\mathrm{H}_0]{d} \mathbf{MVN}_{p+1}(\mathbf{0}, \mathbf{I}^{(1)}(\boldsymbol{\beta})). \tag{A.6}$$

Combining (A.5) and (A.6), it yields that

$$\sqrt{n}\bar{\mathbf{S}}_n \xrightarrow[\mathrm{H}_0]{d} \mathbf{MVN}_{p+1}(\mathbf{0}, \boldsymbol{\Sigma}(\boldsymbol{\beta})),$$

where

$$\boldsymbol{\Sigma}(\boldsymbol{\beta}) = \mathbf{I}^{(1)}(\boldsymbol{\beta}) + \mathbf{I}^{(1)}(\boldsymbol{\beta})\left\{\mathbf{I}^{(0)}(\boldsymbol{\beta})\right\}^{-1}\mathbf{I}^{(1)}(\boldsymbol{\beta}).$$

However, the asymptotic distribution of $\mathbf{S}_n^{(1)}(\boldsymbol{\beta})$ is different under the local alternative $\mathrm{H}_1 : \boldsymbol{\theta} = \boldsymbol{\theta}_0 + \boldsymbol{\delta}/\sqrt{n}$. To derive its asymptotic distribution under H_1, we reformulate it as

$$\begin{aligned}
\frac{1}{n}\mathbf{S}_n^{(1)}(\boldsymbol{\beta}) &= \frac{1}{n}\mathbf{S}_n^{(1)}(\boldsymbol{\beta},\boldsymbol{\theta}) + \left\{\frac{1}{n}\mathbf{S}_n^{(1)}(\boldsymbol{\beta}) - \frac{1}{n}\mathbf{S}_n^{(1)}(\boldsymbol{\theta},\boldsymbol{\beta})\right\}\\
&= \frac{1}{n}\mathbf{S}_n^{(1)}(\boldsymbol{\beta},\boldsymbol{\theta}) + \mathbf{I}^{(1)}(\boldsymbol{\beta})(\boldsymbol{\theta} - \boldsymbol{\theta}_0) + \boldsymbol{o}_P(n^{-\frac{1}{2}}),
\end{aligned}$$

where $\mathbf{S}_n^{(1)}(\boldsymbol{\beta},\boldsymbol{\theta}) = \sum_{i=1}^n\left[\frac{\partial \mu_i^{(1)}(\boldsymbol{\theta},\boldsymbol{\beta})}{\partial \boldsymbol{\theta}^T}\cdot\frac{\{Y_i^{(1)} - \mu_i^{(1)}(\boldsymbol{\theta},\boldsymbol{\beta})\}}{V_i^{(1)}(\boldsymbol{\theta},\boldsymbol{\beta})}\right]$. The last equality follows similarly as from (A.2) to

(A.3) since $\boldsymbol{\theta}-\boldsymbol{\theta}_0 = \boldsymbol{O}_P(n^{-1/2})$. Thus, under the local alternative $\mathrm{H}_1 : \boldsymbol{\theta} = \boldsymbol{\theta}_0 + \boldsymbol{\delta}/\sqrt{n}$,

$$\sqrt{n}\left\{\frac{1}{n}\mathbf{S}_n^{(1)}(\boldsymbol{\beta}) - \mathbf{I}^{(1)}(\boldsymbol{\beta})(\boldsymbol{\theta}-\boldsymbol{\theta}_0)\right\} \xrightarrow[\mathrm{H}_1]{d} \mathbf{MVN}_{p+1}(\mathbf{0}, \mathbf{I}^{(1)}(\boldsymbol{\beta})). \tag{A.7}$$

The asymptotic distribution of $\mathbf{S}_n^{(1)}(\boldsymbol{\beta})$ under H_1 can be obtained by combining (A.5) and (A.7):

$$\sqrt{n}\left\{\bar{\mathbf{S}}_n - \mathbf{I}^{(1)}(\boldsymbol{\beta})(\boldsymbol{\theta}-\boldsymbol{\theta}_0)\right\} \xrightarrow[\mathrm{H}_1]{d} \mathbf{MVN}_{p+1}(\mathbf{0}, \boldsymbol{\Sigma}(\boldsymbol{\beta})).$$

A.1.2 Proof of Remark 2.3.1

In this proof, we will keep the assumptions and notations in A.1.1 and add an addition assumption that $\boldsymbol{\theta}_0 = \mathbf{0}$ for simplicity. Let $\mathscr{F}_n^*$ denote another filtration that is a little bit different from $\mathscr{F}_n$ defined in Remark 2.3.1. That is,

$$\begin{aligned}
\mathscr{F}_n^* &= \{(\mathbf{X}_i^{(j)}, Y_i^{(j)}), i = 1, \cdots, n;\ j = 0, 1; \text{ and } (\mathbf{X}_{n+1}^{(0)}, Y_{n+1}^{(0)})\} \\
\mathscr{F}_n &= \{(\mathbf{X}_i^{(j)}, Y_i^{(j)}), i = 1, \cdots, n;\ j = 0, 1\}.
\end{aligned}$$

By the *law of total expectation*, we have $\mathbb{E}(\lambda_{n+1}|\mathscr{F}_n) = \mathbb{E}\left\{\mathbb{E}\left(\lambda_{n+1}|\mathscr{F}_n^*\right)|\mathscr{F}_n\right\}$.

We first compute $\mathbb{E}\left(\lambda_{n+1}|\mathscr{F}_n^*\right)$. The probability ratio λ_{n+1} defined in Section 2.3 can be simplified to

$$\lambda_{n+1} = \exp\left\{(n+1)\bar{\mathbf{S}}_{n+1}^T \boldsymbol{\Sigma}^{-1}(\boldsymbol{\beta})\mathbf{I}^{(1)}(\boldsymbol{\beta})\boldsymbol{\theta} - \frac{n+1}{2}\boldsymbol{\theta}^T \mathbf{I}^{(1)}(\boldsymbol{\beta})\boldsymbol{\Sigma}^{-1}(\boldsymbol{\beta})\mathbf{I}^{(1)}(\boldsymbol{\beta})\boldsymbol{\theta}\right\}.$$

The expectation of the random term in λ_{n+1} can calculated as

$$\begin{aligned}
&\mathbb{E}\left\{(n+1)\bar{\mathbf{S}}_{n+1}|\mathscr{F}_n^*\right\} \\
&= \mathbb{E}\left(\sum_{i=1}^{n+1}\left[\frac{\partial \mu_i^{(1)}(\boldsymbol{\theta}, \hat{\boldsymbol{\beta}}_{n+1})}{\partial \boldsymbol{\theta}^T} \cdot \frac{\{Y_i^{(1)} - \mu_i^{(1)}(\boldsymbol{\theta}, \hat{\boldsymbol{\beta}}_{n+1})\}}{V_i^{(1)}(\boldsymbol{\theta}, \hat{\boldsymbol{\beta}}_{n+1})}\right]\Bigg|_{\boldsymbol{\theta}=\mathbf{0}} \Bigg| \mathscr{F}_n^*\right) \\
&= \mathbf{S}_n^{(1)}(\hat{\boldsymbol{\beta}}_{n+1}) + \mathbb{E}\left[\left.\frac{\partial \mu_{n+1}^{(1)}(\boldsymbol{\theta}, \hat{\boldsymbol{\beta}}_{n+1})}{\partial \boldsymbol{\theta}^T}\right|_{\boldsymbol{\theta}=\mathbf{0}} \cdot \frac{\{Y_{n+1}^{(1)} - \mu_{n+1}^{(1)}(\mathbf{0}, \hat{\boldsymbol{\beta}}_{n+1})\}}{V_{n+1}^{(1)}(\mathbf{0}, \hat{\boldsymbol{\beta}}_{n+1})} \Bigg| \mathscr{F}_n^*\right] \\
&= \mathbf{S}_n^{(1)}(\hat{\boldsymbol{\beta}}_{n+1}) + \mathbb{E}\left[\left.\frac{\partial \mu_{n+1}^{(1)}(\boldsymbol{\theta}, \hat{\boldsymbol{\beta}}_{n+1})}{\partial \boldsymbol{\theta}^T}\right|_{\boldsymbol{\theta}=\mathbf{0}} \cdot \mathbb{E}\left[\frac{Y_{n+1}^{(1)} - \mu_{n+1}^{(1)}(\mathbf{0}, \hat{\boldsymbol{\beta}}_{n+1})}{V_{n+1}^{(1)}(\mathbf{0}, \hat{\boldsymbol{\beta}}_{n+1})} \Bigg| \mathbf{X}_{n+1}^{(1)}\right] \Bigg| \mathscr{F}_n^*\right] \\
&\overset{\mathrm{H}_0}{=} \mathbf{S}_n^{(1)}(\hat{\boldsymbol{\beta}}_{n+1}) + \mathbb{E}\left\{\left.\frac{\partial \mu_{n+1}^{(1)}(\boldsymbol{\theta}, \hat{\boldsymbol{\beta}}_{n+1})}{\partial \boldsymbol{\theta}^T}\right|_{\boldsymbol{\theta}=\mathbf{0}} \cdot \frac{\{\mu_{n+1}^{(1)}(\mathbf{0}, \boldsymbol{\beta}) - \mu_{n+1}^{(1)}(\mathbf{0}, \hat{\boldsymbol{\beta}}_{n+1})\}}{V_{n+1}^{(1)}(\mathbf{0}, \hat{\boldsymbol{\beta}}_{n+1})} \Bigg| \mathscr{F}_n^*\right\} \\
&= \mathbf{S}_n^{(1)}(\hat{\boldsymbol{\beta}}_{n+1}) + \mathbb{E}\left\{\left.\frac{\partial \mu_{n+1}^{(1)}(\boldsymbol{\theta}, \hat{\boldsymbol{\beta}}_{n+1})}{\partial \boldsymbol{\theta}^T}\right|_{\boldsymbol{\theta}=\mathbf{0}} \times \left.\frac{\partial \mu_{n+1}^{(1)}(\mathbf{0}, \boldsymbol{\beta})}{\partial \boldsymbol{\beta}}\right|_{\boldsymbol{\beta}=\hat{\boldsymbol{\beta}}_{n+1}} \times \frac{\boldsymbol{\beta} - \hat{\boldsymbol{\beta}}_{n+1}}{V_{n+1}^{(1)}(\mathbf{0}, \hat{\boldsymbol{\beta}}_{n+1})} + \boldsymbol{O}_P(n^{-1}) \Bigg| \mathscr{F}_n^*\right\}
\end{aligned} \tag{A.8}$$

$$\begin{aligned}
&= \mathbf{S}_n^{(1)}(\hat{\boldsymbol{\beta}}_{n+1}) + \mathbf{I}^{(1)}(\hat{\boldsymbol{\beta}}_{n+1})(\boldsymbol{\beta} - \hat{\boldsymbol{\beta}}_{n+1}) + \boldsymbol{O}_P(n^{-1}) \\
&= \mathbf{S}_n^{(1)}(\hat{\boldsymbol{\beta}}_n) - n\mathbf{I}^{(1)}(\hat{\boldsymbol{\beta}}_n)(\hat{\boldsymbol{\beta}}_{n+1} - \hat{\boldsymbol{\beta}}_n) + \boldsymbol{O}_P(1)
\end{aligned} \tag{A.9}$$

$$= n\bar{\mathbf{S}}_n - n\mathbf{I}^{(1)}(\hat{\boldsymbol{\beta}}_n)(\hat{\boldsymbol{\beta}}_{n+1} - \hat{\boldsymbol{\beta}}_n) + \boldsymbol{O}_P(1), \tag{A.10}$$

where (A.8) and (A.9) are obtained by making Taylor expansions. It follows that

$$\begin{aligned}
&\mathbb{E}\left(\lambda_{n+1}|\mathscr{F}_n^*\right) \\
&= \mathbb{E}\left[\exp\left\{(n+1)\bar{\mathbf{S}}_{n+1}^T\boldsymbol{\Sigma}^{-1}(\boldsymbol{\beta})\mathbf{I}^{(1)}(\boldsymbol{\beta})\boldsymbol{\theta} - \frac{n+1}{2}\boldsymbol{\theta}^T\mathbf{I}^{(1)}(\boldsymbol{\beta})\boldsymbol{\Sigma}^{-1}(\boldsymbol{\beta})\mathbf{I}^{(1)}(\boldsymbol{\beta})\boldsymbol{\theta}\right\}\middle|\mathscr{F}_n^*\right] \\
&\overset{\text{Delta Method}}{\approx} \exp\left[\mathbb{E}\left\{(n+1)\bar{\mathbf{S}}_{n+1}^T\boldsymbol{\Sigma}^{-1}(\boldsymbol{\beta})\mathbf{I}^{(1)}(\boldsymbol{\beta})\boldsymbol{\theta} - \frac{n+1}{2}\boldsymbol{\theta}^T\mathbf{I}^{(1)}(\boldsymbol{\beta})\boldsymbol{\Sigma}^{-1}(\boldsymbol{\beta})\mathbf{I}^{(1)}(\boldsymbol{\beta})\boldsymbol{\theta}\middle|\mathscr{F}_n^*\right\}\right] \\
&\overset{(A.10)}{=} \exp\left[\left\{n\bar{\mathbf{S}}_n^T - n(\hat{\boldsymbol{\beta}}_{n+1} - \hat{\boldsymbol{\beta}}_n)^T\mathbf{I}^{(1)}(\hat{\boldsymbol{\beta}}_n) + \boldsymbol{O}_P(1)\right\} \times \boldsymbol{\Sigma}^{-1}(\boldsymbol{\beta})\mathbf{I}^{(1)}(\boldsymbol{\beta})\boldsymbol{\theta} - \frac{n+1}{2}\boldsymbol{\theta}^T\mathbf{I}^{(1)}(\boldsymbol{\beta})\boldsymbol{\Sigma}^{-1}(\boldsymbol{\beta})\mathbf{I}^{(1)}(\boldsymbol{\beta})\boldsymbol{\theta}\right] \\
&= \lambda_n \cdot \exp\left[\left\{-n(\hat{\boldsymbol{\beta}}_{n+1} - \hat{\boldsymbol{\beta}}_n)^T\mathbf{I}^{(1)}(\hat{\boldsymbol{\beta}}_n) + \boldsymbol{O}_P(1)\right\} \times \boldsymbol{\Sigma}^{-1}(\boldsymbol{\beta})\mathbf{I}^{(1)}(\boldsymbol{\beta})\boldsymbol{\theta} - \frac{1}{2}\boldsymbol{\theta}^T\mathbf{I}^{(1)}(\boldsymbol{\beta})\boldsymbol{\Sigma}^{-1}(\boldsymbol{\beta})\mathbf{I}^{(1)}(\boldsymbol{\beta})\boldsymbol{\theta}\right] \\
&= \lambda_n \cdot \exp\left\{-n(\hat{\boldsymbol{\beta}}_{n+1} - \hat{\boldsymbol{\beta}}_n)^T\mathbf{I}^{(1)}(\hat{\boldsymbol{\beta}}_n)\boldsymbol{\Sigma}^{-1}(\boldsymbol{\beta})\mathbf{I}^{(1)}(\boldsymbol{\beta})\boldsymbol{\theta} + o_P(1)\right\}.
\end{aligned} \tag{A.11}$$

Clearly, the random term in (A.11) is $(\hat{\boldsymbol{\beta}}_{n+1} - \hat{\boldsymbol{\beta}}_n)$. Considering the equation (A.1), we have

$$\begin{aligned}
&\mathbb{E}\left(\hat{\boldsymbol{\beta}}_{n+1} - \hat{\boldsymbol{\beta}}_n|\mathscr{F}_n\right) \\
&= \mathbb{E}\left[\left\{\mathbf{I}^{(0)}(\boldsymbol{\beta})\right\}^{-1}\frac{1}{n+1}\mathbf{S}_{n+1}^{(0)}(\boldsymbol{\beta}) - \left\{\mathbf{I}^{(0)}(\boldsymbol{\beta})\right\}^{-1}\frac{1}{n}\mathbf{S}_n^{(0)}(\boldsymbol{\beta}) + \boldsymbol{o}_P(n^{-\frac{1}{2}})\middle|\mathscr{F}_n\right] \\
&= \mathbb{E}\left(\left\{\mathbf{I}^{(0)}(\boldsymbol{\beta})\right\}^{-1}\left[\frac{1}{n+1}\left\{\mathbf{S}_{n+1}^{(0)}(\boldsymbol{\beta}) - \mathbf{S}_n^{(0)}(\boldsymbol{\beta})\right\} + \underbrace{\left(\frac{1}{n+1} - \frac{1}{n}\right)\mathbf{S}_n^{(0)}(\boldsymbol{\beta})}_{\boldsymbol{O}_P(n^{-1})}\right] + \boldsymbol{o}_P(n^{-\frac{1}{2}})\middle|\mathscr{F}_n\right) \\
&= \frac{1}{n+1}\left\{\mathbf{I}^{(0)}(\boldsymbol{\beta})\right\}^{-1} \times \mathbb{E}\left[\frac{\partial\mu_{n+1}^{(0)}(\boldsymbol{\beta})}{\partial\boldsymbol{\beta}^T} \cdot \frac{\{Y_{n+1}^{(0)} - \mu_{n+1}^{(0)}(\boldsymbol{\beta})\}}{V_{n+1}^{(0)}(\boldsymbol{\beta})}\middle|\mathscr{F}_n\right] + \boldsymbol{o}_P(n^{-\frac{1}{2}}) \\
&= \boldsymbol{o}_P(n^{-\frac{1}{2}}).
\end{aligned} \tag{A.12}$$

Since $\boldsymbol{\theta} = \boldsymbol{O}_P(n^{-1/2})$ by local alternative, it follows that

$$\begin{aligned}
&\mathbb{E}(\lambda_{n+1}|\mathscr{F}_n) \\
&= \mathbb{E}\left\{\mathbb{E}\left(\lambda_{n+1}|\mathscr{F}_n^*\right)|\mathscr{F}_n\right\} \\
&\overset{(A.11)}{=} \lambda_n \cdot \mathbb{E}\left[\exp\left\{-n(\hat{\boldsymbol{\beta}}_{n+1} - \hat{\boldsymbol{\beta}}_n)^T\mathbf{I}^{(1)}(\hat{\boldsymbol{\beta}}_n)\boldsymbol{\Sigma}^{-1}(\boldsymbol{\beta})\mathbf{I}^{(1)}(\boldsymbol{\beta})\boldsymbol{\theta} + o_P(1)\right\}\middle|\mathscr{F}_n\right] \\
&\overset{\text{Delta Method}}{\approx} \lambda_n \cdot \exp\left[\mathbb{E}\left\{-n(\hat{\boldsymbol{\beta}}_{n+1} - \hat{\boldsymbol{\beta}}_n)^T\mathbf{I}^{(1)}(\hat{\boldsymbol{\beta}}_n)\boldsymbol{\Sigma}^{-1}(\boldsymbol{\beta})\mathbf{I}^{(1)}(\boldsymbol{\beta})\boldsymbol{\theta} + o_P(1)\middle|\mathscr{F}_n\right\}\right] \\
&\overset{(A.12)}{=} \lambda_n \cdot \exp\{o_P(1)\},
\end{aligned}$$

which finalizes our proof.

A.2 Additional Simulation Results

We also compare SST with mSPRT on data generated from generalized linear models with homogeneous treatment effects ((2.3)-(2.5)). We consider the same three generalized linear models and five covariates as in Section 2.5.1. In this setting, the covariates are only used in SST and will not impact the response generation. The true values of the testing parameter θ and nuisance parameter β are shown in the tables. The significance level $\alpha = 0.5$ and the null value of the testing parameter $\theta_0 = 0$ are fixed for all experiments. Each experiment is repeated 1000 times.

In A/B testing, data are generated in batches with batch size 200 and equally assigned to the control group and treatment group. After each batch, we compute the corresponding test statistic and reject the null hypothesis the first time it exceeds its corresponding predetermined threshold. We choose a failure time $N = 10000$. We can see that both mSPRT and SST are able to control the type I error (Table A.1), but SST has a slightly lower power than SST (Table A.2). This is the price that SST pays to account for the covariates in the absence of heterogeneous treatment effects.

Table A.1 Estimated type I error of SST and mSPRT in A/B testing on data with homogeneous treatment effects

GLM	θ	β	Covariates	Type I error (SST)	Type I error (mSPRT)
Logistic Regression	0	0.5	N(0,1)	0.013	0.004
			U[-1,1]	0.011	
			Ber(0.5)	0.014	
			MVN	0.012	
			N(0,1)+U[-1,1]	0.017	
Linear Regression	0	0	N(0,1)	0.001	0.005
			U[-1,1]	0.008	
			Ber(0.5)	0.004	
			MVN	< 0.001	
			N(0,1)+U[-1,1]	< 0.001	
Log Regression	0	1	N(0,1)	0.007	NA
			U[-1,1]	0.009	
			Ber(0.5)	0.009	
			MVN	0.003	
			N(0,1)+U[-1,1]	0.009	

In multiple testing, we consider $m = 64$ hypotheses, among which 3/4 are true null hypotheses $\theta = 0$, and the remaining 1/4 true alternatives are equally placed at $\theta = 0.01, 0.03, 0.05, 0.07$ for logistic regression and $\beta = 0.05, 0.07, 0.09, 0.11$ for linear regression and log regression. For each comparison, we wait until the data are accumulated to $N = 10000$ (for each group) and compute the sequential p-value

Table A.2 Estimated power of SST and mSPRT in A/B testing on data with homogeneous treatment effects

GLM	θ	β	Covariates	Power (SST)	Power (mSPRT)
Logistic Regression	0.03	0.5	N(0,1)	0.548	0.686
			U[-1,1]	0.597	
			Ber(0.5)	0.618	
			MVN	0.301	
			N(0,1)+U[-1,1]	0.332	
	0.04	0.5	N(0,1)	0.921	0.969
			U[-1,1]	0.942	
			Ber(0.5)	0.950	
			MVN	0.787	
			N(0,1)+U[-1,1]	0.810	
Linear Regression	0.05	0	N(0,1)	0.548	0.686
			U[-1,1]	0.597	
			Ber(0.5)	0.618	
			MVN	0.301	
			N(0,1)+U[-1,1]	0.332	
	0.07	0	N(0,1)	0.921	0.969
			U[-1,1]	0.942	
			Ber(0.5)	0.950	
			MVN	0.787	
			N(0,1)+U[-1,1]	0.810	
Log Regression	0.07	1	N(0,1)	0.702	NA
			U[-1,1]	0.739	
			Ber(0.5)	0.754	
			MVN	0.417	
			N(0,1)+U[-1,1]	0.420	
	0.09	1	N(0,1)	0.972	NA
			U[-1,1]	0.972	
			Ber(0.5)	0.972	
			MVN	0.873	
			N(0,1)+U[-1,1]	0.875	

p_N. Table A.3 shows that both mSPRT and SST are able to control the false discovery rate, but mSPRT achieves a higher true positive rate than SST.

Table A.3 Estimated false discovery rate and true positive rate of SST and mSPRT in multiple testing on data with homogeneous treatment effects

GLM	Covariates	FDR (SST)	FDR (mSPRT)	TPR (SST)	TPR (mSPRT)
Logistic Regression	N(0,1)	0.0039	0.0005	0.5565	0.5877
	U[-1,1]	0.0033		0.5726	
	Ber(0.5)	0.0034		0.5759	
	MVN	0.0055		0.5118	
	N(0,1)+U[-1,1]	0.0051		0.5191	
Linear Regression	N(0,1)	0.0005	0.0016	0.5815	0.7208
	U[-1,1]	0.0009		0.6054	
	Ber(0.5)	0.0010		0.6113	
	MVN	0.0001		0.4587	
	N(0,1)+U[-1,1]	0.0007		0.4672	
Log Regression	N(0,1)	0.0029	NA	0.5884	NA
	U[-1,1]	0.0036		0.6063	
	Ber(0.5)	0.0038		0.6114	
	MVN	0.0027		0.4668	
	N(0,1)+U[-1,1]	0.0017		0.4785	

APPENDIX

B

SUPPLEMENT TO CHAPTER 3

B.1 Proof of Main Results

B.1.1 Proof of Theorem 3.3.1

Let $\mathscr{F}_j$, $0 \le j \le k$, denote a filtration generated by observations in first $(j+1)$ batches $\overline{\mathscr{C}}_j = \cup_{r=0}^{j} \mathscr{C}_r$, and $\bar{D}_j(\mathscr{C}_j; \hat{d}_{j-1}^{opt}, \mu, \theta, p)$ denote an AIPW estimator for Δ with only optimal decision rule estimated by previous batches:

$$\begin{aligned}\bar{D}_j(\mathscr{C}_j; \hat{d}_{j-1}^{opt}, \mu, \theta, p) \coloneqq \frac{1}{m} \sum_{\mathbf{O}_i \in \mathscr{C}_j} \Bigg[& \frac{Y_i \cdot 1\left(A_i = 1\{\hat{\theta}_{j-1}(\mathbf{X}_i) > 0\}\right)}{p_{A_i}(\mathbf{X}_i)} - \left\{ \frac{1\left(A_i = 1\{\hat{\theta}_{j-1}(\mathbf{X}_i) > 0\}\right)}{p_{A_i}(\mathbf{X}_i)} - 1 \right\} \\ & \times g^{-1}\left(\mu(\mathbf{X}_i) + \theta(\mathbf{X}_i) 1\{\hat{\theta}_{j-1}(\mathbf{X}_i) > 0\}\right) \Bigg] \\ & - \Bigg[\frac{Y_i \cdot 1(A_i = 0)}{1 - p(\mathbf{X}_i)} - \left\{ \frac{1(A_i = 0)}{1 - p(\mathbf{X}_i)} - 1 \right\} \times g^{-1}\left(\mu(\mathbf{X}_i)\right) \Bigg].\end{aligned}$$

Then

$$\begin{aligned}& \frac{1}{k}\left(\sum_{j=1}^{k} \hat{\sigma}_j^{-1}\right)(\hat{\Delta}_k - \Delta) \\ = & \frac{1}{k} \sum_{j=1}^{k} \hat{\sigma}_j^{-1} \left\{ \bar{D}_j(\mathscr{C}_j; \overline{\mathscr{C}}_{j-1}) - \Delta \right\} \\ = & \frac{1}{k} \sum_{j=1}^{k} \hat{\sigma}_j^{-1} \left(\left[\bar{D}_j(\mathscr{C}_j; \overline{\mathscr{C}}_{j-1}) - \mathbb{E}\left\{ \bar{D}_j(\mathscr{C}_j; \hat{d}_{j-1}^{opt}, \mu, \theta, p) | \overline{\mathscr{C}}_{j-1} \right\} \right] + \left[\mathbb{E}\left\{ \bar{D}_j(\mathscr{C}_j; \hat{d}_{j-1}^{opt}, \mu, \theta, p) | \overline{\mathscr{C}}_{j-1} \right\} - \Delta \right] \right)\end{aligned}$$

$$= \frac{1}{k}\sum_{j=1}^{k}\hat{\sigma}_j^{-1}\Big[\bar{D}_j(\mathscr{C}_j;\overline{\mathscr{C}}_{j-1}) - \mathbb{E}\left\{\bar{D}_j(\mathscr{C}_j;\hat{d}_{j-1}^{opt},\mu,\theta,p)|\overline{\mathscr{C}}_{j-1}\right\}\Big] + o_P(k^{-1/2}) \tag{B.1}$$

$$= \frac{1}{k}\sum_{j=1}^{k}\hat{\sigma}_j^{-1}\Big[\bar{D}_j(\mathscr{C}_j;\overline{\mathscr{C}}_{j-1}) - \mathbb{E}\left\{\bar{D}_j(\mathscr{C}_j;\overline{\mathscr{C}}_{j-1})|\mathscr{F}_{j-1}\right\}\Big] + o_P(k^{-1/2}). \tag{B.2}$$

Above (B.1) follows by condition (C5) and (B.2) follows by condition (C4). For $j = 1, 2, \cdots, k$, let

$$M_{k,j} = \frac{1}{\sqrt{k}} \cdot \frac{\bar{D}_j(\mathscr{C}_j;\overline{\mathscr{C}}_{j-1}) - \mathbb{E}\{\bar{D}_j(\mathscr{C}_j;\overline{\mathscr{C}}_{j-1})|\mathscr{F}_{j-1}\}}{\hat{\sigma}_j}.$$

It is obvious that for each k, $M_{k,j}$, $1 \le j \le k$, is a martingale with respect to the filtration $\mathscr{F}_j$. In particular, for all $j \ge 1$, $\mathbb{E}(M_{k,j}|\mathscr{F}_{j-1}) = 0$ and $\sum_{j=1}^{k}\mathbb{E}(M_{k,j}^2|\mathscr{F}_{j-1}) = \frac{1}{k}\sum_{j=1}^{k}\sigma_j^2/\hat{\sigma}_j^2 \xrightarrow{p} 1$ as $k \longrightarrow \infty$ by (C3). The conditional Lindeberg condition holds in (C2), so the *martingale central limit theory* for triangular arrays gives

$$\sum_{j=1}^{k} M_{k,j} \xrightarrow{d} N(0,1).$$

Plugging it back into (B.2), we can get

$$\frac{1}{\sqrt{k}}\left(\sum_{j=1}^{k}\hat{\sigma}_j^{-1}\right)\left(\hat{\Delta}_k - \Delta\right) \xrightarrow{d} N(0,1).$$

B.1.2 Proof of Proposition 3.3.1

We first simplify the formula of λ_k to:

$$\begin{aligned}
\lambda_k &= \frac{\psi_{\left(\frac{1}{\sqrt{k}}\left(\sum_{j=1}^{k}\hat{\sigma}_j^{-1}\right)\Delta,1\right)}(R_k)}{\psi_{(0,1)}(R_k)} \\
&= \exp\left\{\frac{1}{\sqrt{k}}\sum_{j=1}^{k}\hat{\sigma}_j^{-1}\cdot\Delta\cdot R_k - \frac{1}{2k}(\sum_{j=1}^{k}\hat{\sigma}_j^{-1})^2\cdot\Delta^2\right\} \\
&= \exp\left\{\frac{1}{k}\sum_{j=1}^{k}\hat{\sigma}_j^{-1}\cdot\Delta\cdot\sum_{j=1}^{k}(\hat{\sigma}_j^{-1}\bar{D}_j) - \frac{1}{2k}(\sum_{j=1}^{k}\hat{\sigma}_j^{-1})^2\cdot\Delta^2\right\},
\end{aligned}$$

where we denote $\bar{D}_j(\mathscr{C}_j;\overline{\mathscr{C}}_{j-1})$ with $\bar{D}_j$ for simplicity. Let

$$\hat{\Delta}_k \coloneqq \frac{\sum_{j=1}^{k}\hat{\sigma}_j^{-1}\bar{D}_j}{\sum_{j=1}^{k}\hat{\sigma}_j^{-1}}. \tag{B.3}$$

and remember that Theorem 3.1 gives

$$\frac{1}{\sqrt{k}}\left(\sum_{j=1}^{k}\hat{\sigma}_j^{-1}\right)\left(\hat{\Delta}_k - \Delta\right) \xrightarrow{d} N(0,1), \tag{B.4}$$

where $\hat{\sigma}_j$ is estimated from the first j batches $\overline{\mathscr{C}}_{j-1},\ j=1,2,\cdots,k$. Since the true value difference Δ is not very large in practice, we assume local alternative $\Delta = O_P(k^{-1/2})$ here as in Theorem 3.1.

Then,

$$
\begin{aligned}
&\mathbb{E}_{H_0}(\lambda_{k+1}|\mathscr{F}_k)\\
&= \mathbb{E}_{H_0}\left[\exp\left\{\frac{1}{k+1}\sum_{j=1}^{k+1}\hat{\sigma}_j^{-1}\cdot\Delta\cdot\sum_{j=1}^{k+1}(\hat{\sigma}_j^{-1}\bar{D}_j)-\frac{1}{2(k+1)}(\sum_{j=1}^{k+1}\hat{\sigma}_j^{-1})^2\cdot\Delta^2\right\}\middle|\mathscr{F}_k\right]\\
&\overset{\text{Delta Method}}{\approx} \exp\left[\mathbb{E}_{H_0}\left\{\frac{1}{k+1}\sum_{j=1}^{k+1}\hat{\sigma}_j^{-1}\cdot\Delta\cdot\sum_{j=1}^{k+1}(\hat{\sigma}_j^{-1}\bar{D}_j)-\frac{1}{2(k+1)}(\sum_{j=1}^{k+1}\hat{\sigma}_j^{-1})^2\cdot\Delta^2\middle|\mathscr{F}_k\right\}\right]\\
&= \exp\left[\frac{1}{k+1}\sum_{j=1}^{k+1}\hat{\sigma}_j^{-1}\cdot\Delta\cdot\left\{\sum_{j=1}^{k}(\hat{\sigma}_j^{-1}\bar{D}_j)+\hat{\sigma}_{k+1}^{-1}\cdot\underbrace{\mathbb{E}_{H_0}(\bar{D}_{k+1}|\mathscr{F}_k)}_{0}\right\}-\frac{1}{2(k+1)}(\sum_{j=1}^{k+1}\hat{\sigma}_j^{-1})^2\cdot\Delta^2\right]\\
&= \exp\left\{\frac{1}{k+1}\sum_{j=1}^{k+1}\hat{\sigma}_j^{-1}\cdot\Delta\cdot\sum_{j=1}^{k}(\hat{\sigma}_j^{-1}\bar{D}_j)-\frac{1}{2(k+1)}(\sum_{j=1}^{k+1}\hat{\sigma}_j^{-1})^2\cdot\Delta^2\right\}\\
&= \exp\left[\frac{1}{k}\sum_{j=1}^{k}\hat{\sigma}_j^{-1}\cdot\Delta\cdot\sum_{j=1}^{k}(\hat{\sigma}_j^{-1}\bar{D}_j)+\left(\frac{1}{k+1}\sum_{j=1}^{k+1}\hat{\sigma}_j^{-1}-\frac{1}{k}\sum_{j=1}^{k}\hat{\sigma}_j^{-1}\right)\cdot\Delta\cdot\sum_{j=1}^{k}(\hat{\sigma}_j^{-1}\bar{D}_j)\right.\\
&\qquad\left.-\frac{1}{2k}(\sum_{j=1}^{k}\hat{\sigma}_j^{-1})^2\cdot\Delta^2-\left\{\frac{1}{2(k+1)}(\sum_{j=1}^{k+1}\hat{\sigma}_j^{-1})^2-\frac{1}{2k}(\sum_{j=1}^{k}\hat{\sigma}_j^{-1})^2\right\}\cdot\Delta^2\right]\\
&= \lambda_k\cdot\exp\left[\left(\frac{1}{k+1}\sum_{j=1}^{k+1}\hat{\sigma}_j^{-1}-\frac{1}{k}\sum_{j=1}^{k}\hat{\sigma}_j^{-1}\right)\cdot\Delta\cdot\sum_{j=1}^{k}(\hat{\sigma}_j^{-1}\bar{D}_j)-\left\{\frac{1}{2(k+1)}(\sum_{j=1}^{k+1}\hat{\sigma}_j^{-1})^2-\frac{1}{2k}(\sum_{j=1}^{k}\hat{\sigma}_j^{-1})^2\right\}\cdot\Delta^2\right]\\
&\overset{(B.3)}{=} \lambda_k\cdot\exp\left\{\left(\frac{1}{k+1}\sum_{j=1}^{k+1}\hat{\sigma}_j^{-1}-\frac{1}{k}\sum_{j=1}^{k}\hat{\sigma}_j^{-1}\right)\cdot\Delta\cdot\sum_{j=1}^{k}\hat{\sigma}_j^{-1}\cdot\hat{\Delta}_k-\left\{\frac{1}{2(k+1)}(\sum_{j=1}^{k+1}\hat{\sigma}_j^{-1})^2-\frac{1}{2k}(\sum_{j=1}^{k}\hat{\sigma}_j^{-1})^2\right\}\cdot\Delta^2\right\}\\
&= \lambda_k\cdot\exp\left\{\underbrace{\left(\frac{1}{k+1}\sum_{j=1}^{k+1}\hat{\sigma}_j^{-1}-\frac{1}{k}\sum_{j=1}^{k}\hat{\sigma}_j^{-1}\right)}_{O_P(k^{-1})}\times\underbrace{\sum_{j=1}^{k}\hat{\sigma}_j^{-1}\hat{\Delta}_k}_{O_P(k^{1/2})\text{by (B.4)}}\times\Delta-\underbrace{\left\{\frac{1}{2(k+1)}(\sum_{j=1}^{k+1}\hat{\sigma}_j^{-1})^2-\frac{1}{2k}(\sum_{j=1}^{k}\hat{\sigma}_j^{-1})^2\right\}}_{O_P(1)}\times\Delta^2\right\}\\
&= \lambda_k\cdot\exp\{o_P(1)\}
\end{aligned}
$$

B.2 Additional Simulation Results

To evaluate the robustness of SUBTLE to the mixture density variance τ^2, we conduct additional simulations with $\tau^2\in\{0.0001, 0.001, 0.01, 0.1, 1, 10\}$. The data are generated from Model I with $c=0,1$, and batch size $m=200$ for computation efficiency. When $c=0$ we estimate the type I error, while when $c=1$ we estimate the power. The results in Table B.1 show that the type I error is always controlled

below significance level 0.05 and the power has considerable robustness to the choice of τ^2.

Table B.1 Estimated type I error and power of SUBTLE with varying mixture density variance

τ^2	0.0001	0.001	0.01	0.1	1	10
Type I error	0.003	0.024	0.021	0.016	0.005	0.002
Power	0.887	0.976	0.956	0.932	0.892	0.825

APPENDIX

C

SUPPLEMENT TO CHAPTER 4

C.1 Proof of Main Results

C.1.1 Proof of Theorem 4.3.1

Before proving Theorem 4.3.1, we introduce the following two lemmas. The proofs of them are given at the end of the Appendix.

Lemma C.1.1 *Assume Conditions (C2)-(C3) hold. Then, as $nT \to \infty$, we have*

$$\sup_{\theta\in\Theta,\eta\in\mathscr{G}_{\eta_0}} \left\|\mathbb{P}_n\psi(\theta,\eta)-\mathbb{P}\psi(\theta,\eta)\right\|_2 \xrightarrow{P} 0$$

and

$$\sup_{\theta\in\Theta,\eta\in\mathscr{G}_{\eta_0}} \left\|\mathbb{P}_n\dot{\psi}_\theta(\theta,\eta)-\mathbb{P}\dot{\psi}_\theta(\theta,\eta)\right\|_2 \xrightarrow{P} 0,$$

where $\dot{\psi}_\theta(\theta,\eta)=\frac{\partial\psi(\theta,\eta)}{\partial\theta^T}$.

Lemma C.1.2 *Assume Conditions (C2)-(C3) and (C5)-(C6) hold. Then, as $nT \to \infty$, we have*

$$\sqrt{nT}(\mathbb{P}_n-\mathbb{P})\psi(\theta_0,\eta) \rightsquigarrow \boldsymbol{Z} \text{ as random elements in } l^\infty(\mathscr{G}_{\eta_0}),$$

where "$\rightsquigarrow$" represents the weak convergence *of a stochastic process, and $l^\infty(\mathscr{G}_{\eta_0})$ is the collection of all bounded functions $f:\mathscr{G}_{\eta_0}\to\mathbb{R}^p$. The limiting process $\boldsymbol{Z}=\{\boldsymbol{Z}(\eta):\eta\in\mathscr{G}_{\eta_0}\}$ is a mean zero multivariate Gaussian process and the sample paths of $\boldsymbol{Z}$ belong to the set $UC(\mathscr{G}_{\eta_0},\|\cdot\|_{2,P})=\{z\in l^\infty(\mathscr{G}_{\eta_0}) : z$ is uniformly continuous with respect to $\|\cdot\|_{2,P}\}$.*

First, we prove the consistency of $\hat{\theta}$ by showing $\left\|\mathbb{P}\psi(\hat{\theta},\eta_0)\right\|_2 \xrightarrow{P} 0$. Since we have

$$\begin{aligned}\left\|\mathbb{P}\psi(\hat{\theta},\eta_0)\right\|_2 &\le \left\|\mathbb{P}\psi(\hat{\theta},\eta_0)-\mathbb{P}\psi(\hat{\theta},\hat{\eta})\right\|_2 + \left\|\mathbb{P}\psi(\hat{\theta},\hat{\eta})\right\|_2 \\ &= \left\|\mathbb{P}\psi(\hat{\theta},\eta_0)-\mathbb{P}\psi(\hat{\theta},\hat{\eta})\right\|_2 + \left\|\mathbb{P}\psi(\hat{\theta},\hat{\eta})-\mathbb{P}_n\psi(\hat{\theta},\hat{\eta})\right\|_2 \\ &\le \left\|\mathbb{P}\psi(\hat{\theta},\eta_0)-\mathbb{P}\psi(\hat{\theta},\hat{\eta})\right\|_2 + \sup_{\theta\in\Theta,\eta\in\mathscr{G}_{\eta_0}} \left\|\mathbb{P}_n\psi(\theta,\eta)-\mathbb{P}\psi(\theta,\eta)\right\|_2, \end{aligned} \tag{C.1}$$

it only needs to show that both terms in (C.1) are negligible. By Taylor's expansion,

$$\begin{aligned}\left\|\psi_{0,t}(\hat{\theta},\hat{\eta})-\psi_{0,t}(\hat{\theta},\eta_0)\right\|_2 &= \left\|\frac{\partial\psi_{0,t}(\theta,\eta)}{\partial\eta^T}\Big|_{\eta=\tilde{\eta},\theta=\hat{\theta}}(\hat{\eta}-\eta_0)\right\|_2 \\ &\le \left\|\frac{\partial\psi_{0,t}(\theta,\eta)}{\partial\eta^T}\Big|_{\eta=\tilde{\eta},\theta=\hat{\theta}}\right\|_2 \left\|\hat{\eta}-\eta_0\right\|_2,\end{aligned}$$

where $\|\tilde{\eta}-\eta_0\|_2 < \|\hat{\eta}-\eta_0\|_2$. It follows that by the Cauchy-Schwartz inequality

$$\begin{aligned}&\left\|\mathbb{P}\psi(\hat{\theta},\hat{\eta})-\mathbb{P}\psi(\hat{\theta},\eta_0)\right\|_2 \\ &\le \mathbb{P}\left\|\psi(\hat{\theta},\hat{\eta})-\psi(\hat{\theta},\eta_0)\right\|_2 \le \mathbb{P}\left\{\left\|\frac{\partial\psi(\theta,\eta)}{\partial\eta^T}\Big|_{\eta=\tilde{\eta},\theta=\hat{\theta}}\right\|_2 \cdot \left\|\hat{\eta}-\eta_0\right\|_2\right\} \\ &\le \frac{1}{T}\sum_{t=1}^{T}\left(\left[\mathbb{E}\left\{\left\|\frac{\partial\psi_{0,t}(\theta,\eta)}{\partial\eta^T}\Big|_{\eta=\tilde{\eta},\theta=\hat{\theta}}\right\|_2^2\right\}\right]^{1/2} \cdot \left[\mathbb{E}\left\{\|\hat{\eta}(\overline{V}_{0,t})-\eta_0(\overline{V}_{0,t})\|_2^2\right\}\right]^{1/2}\right) \\ &\le \left[\frac{1}{T}\sum_{t=1}^{T}\mathbb{E}\left\{\left\|\frac{\partial\psi_{0,t}(\theta,\eta)}{\partial\eta^T}\Big|_{\eta=\tilde{\eta},\theta=\hat{\theta}}\right\|_2^2\right\}\right]^{1/2} \cdot \left[\frac{1}{T}\sum_{t=1}^{T}\mathbb{E}\left\{\|\hat{\eta}(\overline{V}_{0,t})-\eta_0(\overline{V}_{0,t})\|_2^2\right\}\right]^{1/2} \\ &\le b^*\|\hat{\eta}-\eta_0\|_{2,P} = o_P(1).\end{aligned} \tag{C.2}$$

By Lemma C.1.1, we have

$$\sup_{\theta\in\Theta,\eta\in\mathscr{G}_{\eta_0}} \left\|\mathbb{P}_n\psi(\theta,\eta)-\mathbb{P}\psi(\theta,\eta)\right\|_2 \xrightarrow{P} 0 \text{ as } nT\to\infty.$$

Combining with (C.1) and (C.2), it yields that $\left\|\mathbb{P}\psi(\hat{\theta},\eta_0)\right\|_2 = o_P(1)$. Then, by the identification condition (C1), we have $\hat{\theta}\xrightarrow{P}\theta_0$ as $nT\to\infty$.

Next, we prove the asymptotic distribution of $\hat{\theta}$. By Taylor expansion, we have

$$0 = \sqrt{nT}\mathbb{P}_n\psi(\hat{\theta},\hat{\eta}) = \sqrt{nT}\mathbb{P}_n\psi(\theta_0,\hat{\eta}) + \mathbb{P}_n\dot{\psi}_\theta(\tilde{\theta},\hat{\eta})\sqrt{nT}(\hat{\theta}-\theta_0),$$

where $\|\tilde{\theta}-\theta_0\|_2 < \|\hat{\theta}-\theta_0\|_2$. By Lemma C.1.1, we have

$$\sup_{\theta\in\Theta,\eta\in\mathscr{G}_{\eta_0}} \left\|\mathbb{P}_n\dot{\psi}_\theta(\theta,\eta)-\mathbb{P}\dot{\psi}_\theta(\theta,\eta)\right\|_2 \xrightarrow{P} 0 \text{ as } nT\to\infty.$$

It follows that

$$\mathbb{P}_n\dot{\psi}_\theta(\tilde{\theta},\hat{\eta})\xrightarrow{P}\mathbb{P}\dot{\psi}_\theta(\theta_0,\eta_0) \text{ as } nT\to\infty,$$

because $\tilde{\theta} \xrightarrow{P} \theta_0$ and $\hat{\eta}$ is also consistent by Condition (C4). Therefore,

$$\begin{aligned}\sqrt{nT}(\hat{\theta}-\theta_0) &= -\left\{\mathbb{P}\dot{\psi}_\theta(\theta_0,\eta_0)\right\}^{-1}\sqrt{nT}\mathbb{P}_n\psi(\theta_0,\hat{\eta})+o_P(1)\\ &= -\left\{\mathbb{P}\dot{\psi}_\theta(\theta_0,\eta_0)\right\}^{-1}\sqrt{nT}\left\{(\mathbb{P}_n-\mathbb{P})\psi(\theta_0,\hat{\eta})+\mathbb{P}\psi(\theta_0,\hat{\eta})\right\}+o_P(1). \qquad \text{(C.3)}\end{aligned}$$

Moreover, we have

$$\begin{aligned}\mathbb{P}\psi(\theta_0,\hat{\eta}) &= \frac{1}{T}\sum_{t=1}^{T}\mathbb{E}\Big(c(\overline{V}_{0,t})\cdot\big[Y_{0,t}\exp\left\{-f_{\theta_0}(\overline{V}_{0,t})A_{0,t}\right\}-\hat{\mu}_1(\overline{V}_{0,t})\exp\{-f_{\theta_0}(\overline{V}_{0,t})\}\hat{\pi}(\overline{V}_{0,t})\\ &\quad -\hat{\mu}_0(\overline{V}_{0,t})\{1-\hat{\pi}(\overline{V}_{0,t})\}\big]\cdot\left\{I(A_{0,t}=1)-\hat{\pi}(\overline{V}_{0,t})\right\}\Big)\\ &= \frac{1}{T}\sum_{t=1}^{T}\mathbb{E}\Big(c(\overline{V}_{0,t})\cdot\big[\mu_{1,0}(\overline{V}_{0,t})\exp\{-f_{\theta_0}(\overline{V}_{0,t})\}\pi_0(\overline{V}_{0,t})+\mu_{0,0}(\overline{V}_{0,t})\{1-\pi_0(\overline{V}_{0,t})\}\\ &\quad -\hat{\mu}_1(\overline{V}_{0,t})\exp\{-f_{\theta_0}(\overline{V}_{0,t})\}\hat{\pi}(\overline{V}_{0,t})-\hat{\mu}_0(\overline{V}_{0,t})\{1-\hat{\pi}(\overline{V}_{0,t})\}\big]\cdot\left\{\pi_0(\overline{V}_{0,t})-\hat{\pi}(\overline{V}_{0,t})\right\}\Big)\\ &= \frac{1}{T}\sum_{t=1}^{T}\mathbb{E}\Big(c(\overline{V}_{0,t})\cdot\big[\mu_{1,0}(\overline{V}_{0,t})\exp\{-f_{\theta_0}(\overline{V}_{0,t})\}\pi_0(\overline{V}_{0,t})-\mu_{1,0}(\overline{V}_{0,t})\exp\{-f_{\theta_0}(\overline{V}_{0,t})\}\hat{\pi}(\overline{V}_{0,t})\\ &\quad +\mu_{1,0}(\overline{V}_{0,t})\exp\{-f_{\theta_0}(\overline{V}_{0,t})\}\hat{\pi}(\overline{V}_{0,t})-\hat{\mu}_1(\overline{V}_{0,t})\exp\{-f_{\theta_0}(\overline{V}_{0,t})\}\hat{\pi}(\overline{V}_{0,t})\\ &\quad +\mu_{0,0}(\overline{V}_{0,t})\{1-\pi_0(\overline{V}_{0,t})\}-\mu_{0,0}(\overline{V}_{0,t})\{1-\hat{\pi}(\overline{V}_{0,t})\}\\ &\quad +\mu_{0,0}(\overline{V}_{0,t})\{1-\hat{\pi}(\overline{V}_{0,t})\}-\hat{\mu}_0(\overline{V}_{0,t})\{1-\hat{\pi}(\overline{V}_{0,t})\}\big]\cdot\left\{\pi_0(\overline{V}_{0,t})-\hat{\pi}(\overline{V}_{0,t})\right\}\Big).\end{aligned}$$

Applying the Cauchy-Schwartz inequality, after some calculations we have

$$\begin{aligned}&\left\|\mathbb{P}\psi(\theta_0,\hat{\eta})\right\|_2\\ &\le \frac{1}{T}\sum_{t=1}^{T}\mathbb{E}\Big(\Big\|c(\overline{V}_{0,t})\cdot\big[\mu_{1,0}(\overline{V}_{0,t})\exp\{-f_{\theta_0}(\overline{V}_{0,t})\}\pi_0(\overline{V}_{0,t})-\mu_{1,0}(\overline{V}_{0,t})\exp\{-f_{\theta_0}(\overline{V}_{0,t})\}\hat{\pi}(\overline{V}_{0,t})\\ &\quad +\mu_{1,0}(\overline{V}_{0,t})\exp\{-f_{\theta_0}(\overline{V}_{0,t})\}\hat{\pi}(\overline{V}_{0,t})-\hat{\mu}_1(\overline{V}_{0,t})\exp\{-f_{\theta_0}(\overline{V}_{0,t})\}\hat{\pi}(\overline{V}_{0,t})\\ &\quad +\mu_{0,0}(\overline{V}_{0,t})\{1-\pi_0(\overline{V}_{0,t})\}-\mu_{0,0}(\overline{V}_{0,t})\{1-\hat{\pi}(\overline{V}_{0,t})\}\\ &\quad +\mu_{0,0}(\overline{V}_{0,t})\{1-\hat{\pi}(\overline{V}_{0,t})\}-\hat{\mu}_0(\overline{V}_{0,t})\{1-\hat{\pi}(\overline{V}_{0,t})\}\big]\cdot\left\{\pi_0(\overline{V}_{0,t})-\hat{\pi}(\overline{V}_{0,t})\right\}\Big\|_2\Big)\\ &\le \frac{1}{T}\sum_{t=1}^{T}\mathbb{E}\Big[\left\|c(\overline{V}_{0,t})\right\|_2\cdot\left|\mu_{1,0}(\overline{V}_{0,t})\exp\{-f_{\theta_0}(\overline{V}_{0,t})\}\pi_0(\overline{V}_{0,t})-\mu_{1,0}(\overline{V}_{0,t})\exp\{-f_{\theta_0}(\overline{V}_{0,t})\}\hat{\pi}(\overline{V}_{0,t})\right.\\ &\quad +\mu_{1,0}(\overline{V}_{0,t})\exp\{-f_{\theta_0}(\overline{V}_{0,t})\}\hat{\pi}(\overline{V}_{0,t})-\hat{\mu}_1(\overline{V}_{0,t})\exp\{-f_{\theta_0}(\overline{V}_{0,t})\}\hat{\pi}(\overline{V}_{0,t})\\ &\quad +\mu_{0,0}(\overline{V}_{0,t})\{1-\pi_0(\overline{V}_{0,t})\}-\mu_{0,0}(\overline{V}_{0,t})\{1-\hat{\pi}(\overline{V}_{0,t})\}\\ &\quad \left.+\mu_{0,0}(\overline{V}_{0,t})\{1-\hat{\pi}(\overline{V}_{0,t})\}-\hat{\mu}_0(\overline{V}_{0,t})\{1-\hat{\pi}(\overline{V}_{0,t})\}\right|\cdot\left|\pi_0(\overline{V}_{0,t})-\hat{\pi}(\overline{V}_{0,t})\right|\Big]\\ &\preceq \frac{1}{T}\sum_{t=1}^{T}\mathbb{E}\Big[\left\{\left|\pi_0(\overline{V}_{0,t})-\hat{\pi}(\overline{V}_{0,t})\right|+\left|\mu_{1,0}(\overline{V}_{0,t})-\hat{\mu}_1(\overline{V}_{0,t})\right|+\left|\mu_{0,0}(\overline{V}_{0,t})-\hat{\mu}_0(\overline{V}_{0,t})\right|\right\}\\ &\quad \cdot\left|\pi_0(\overline{V}_{0,t})-\hat{\pi}(\overline{V}_{0,t})\right|\Big]\end{aligned}$$

$$
\begin{aligned}
&\preceq \frac{1}{T}\sum_{t=1}^{T}\left\{\left(\mathbb{E}\Big[\Big\{\Big|\pi_0(\overline{V}_{0,t})-\hat{\pi}(\overline{V}_{0,t})\Big|+\Big|\mu_{1,0}(\overline{V}_{0,t})-\hat{\mu}_1(\overline{V}_{0,t})\Big|+\Big|\mu_{0,0}(\overline{V}_{0,t})-\hat{\mu}_0(\overline{V}_{0,t})\Big|\Big\}^2\Big]\right)^{1/2}\right.\\
&\quad\left.\cdot\left(\mathbb{E}\Big[\big\{\pi_0(\overline{V}_{0,t})-\hat{\pi}(\overline{V}_{0,t})\big\}^2\Big]\right)^{1/2}\right\}\\
&\preceq\left(\frac{1}{T}\sum_{t=1}^{T}\mathbb{E}\Big[\Big\{\Big|\pi_0(\overline{V}_{0,t})-\hat{\pi}(\overline{V}_{0,t})\Big|+\Big|\mu_{1,0}(\overline{V}_{0,t})-\hat{\mu}_1(\overline{V}_{0,t})\Big|+\Big|\mu_{0,0}(\overline{V}_{0,t})-\hat{\mu}_0(\overline{V}_{0,t})\Big|\Big\}^2\Big]\right)^{1/2}\\
&\quad\cdot\left(\frac{1}{T}\sum_{t=1}^{T}\mathbb{E}\Big[\big\{\pi_0(\overline{V}_{0,t})-\hat{\pi}(\overline{V}_{0,t})\big\}^2\Big]\right)^{1/2}\\
&=\Big\||\hat{\pi}-\pi_0|+|\hat{\mu}_1-\mu_{1,0}|+|\hat{\mu}_0-\mu_{0,0}|\Big\|_{2,P}\cdot\Big\|\hat{\pi}-\pi_0\Big\|_{2,P}\\
&\leq\big(\|\hat{\pi}-\pi_0\|_{2,P}+\|\hat{\mu}_1-\mu_{1,0}\|_{2,P}+\|\hat{\mu}_0-\mu_{0,0}\|_{2,P}\big)\cdot\|\hat{\pi}-\pi_0\|_{2,P}=o_P(1/\sqrt{nT}),
\end{aligned}
$$

where we use the notation $\preceq$ to represent the left-hand side is bounded by a constant times the right-hand side. The last equality is due to Condition (C4). This shows

$$\sqrt{nT}\mathbb{P}\psi(\theta_0,\hat{\eta})=o_p(1). \tag{C.4}$$

In addition, by Lemma C.1.2, we have

$$\sqrt{nT}(\mathbb{P}_n-\mathbb{P})\psi(\theta_0,\eta)\rightsquigarrow \boldsymbol{Z} \text{ in } l^\infty(\mathscr{G}_{\eta_0}) \text{ as } nT\to\infty.$$

Since $\hat{\eta}\xrightarrow{P}\eta_0$ in the semimetric space $\mathscr{G}_{\eta_0}$ relative to the metric $\|\cdot\|_{2,P}$, it follows that

$$(\sqrt{nT}(\mathbb{P}_n-\mathbb{P})\psi(\theta_0,\eta),\hat{\eta})\rightsquigarrow(\boldsymbol{Z},\eta_0) \text{ in the space } l^\infty(\mathscr{G}_{\eta_0})\times\mathscr{G}_{\eta_0} \text{ as } nT\to\infty.$$

Define a function $s: l^\infty(\mathscr{G}_{\eta_0})\times\mathscr{G}_{\eta_0}\mapsto\mathbb{R}^p$ by $s(z,\eta)=z(\eta)-z(\eta_0)$. Notice that the function s is continuous at every point (z,η) such that $\eta\mapsto z(\eta)$ is continuous. By Lemma C.1.2, almost all sample paths of $\boldsymbol{Z}$ are continuous on $\mathscr{G}_{\eta_0}$. Thus the function s is continuous at almost every point $(\boldsymbol{Z},\eta_0)$. By the *continuous-mapping theorem*,

$$\sqrt{nT}(\mathbb{P}_n-\mathbb{P})\big\{\psi(\theta_0,\hat{\eta})-\psi(\theta_0,\eta_0)\big\}=s(\sqrt{nT}(\mathbb{P}_n-\mathbb{P})\psi(\theta_0,\eta),\hat{\eta})\rightsquigarrow s(\boldsymbol{Z},\eta_0)=0.$$

Thus, we have

$$\sqrt{nT}(\mathbb{P}_n-\mathbb{P})\psi(\theta_0,\hat{\eta})=\sqrt{nT}(\mathbb{P}_n-\mathbb{P})\psi(\theta_0,\eta_0)+o_P(1).$$

Combining it with (C.3) and (C.4), we have

$$\sqrt{nT}(\hat{\theta}-\theta_0)=-\big\{\mathbb{P}\dot{\psi}_\theta(\theta_0,\eta_0)\big\}^{-1}\sqrt{nT}(\mathbb{P}_n-\mathbb{P})\psi(\theta_0,\eta_0)+o_P(1).$$

Lastly, we only need to show

$$\sqrt{nT}(\mathbb{P}_n - \mathbb{P})\psi(\theta_0, \eta_0) \xrightarrow{d} MVN(\mathbf{0}, \mathbf{\Sigma}) \text{ as } nT \to \infty. \tag{C.5}$$

For any integer $1 \le g \le nT$, let $i(g)$ be the quotient of $g + T$ divided by T, and $t(g)$ be the integer that satisfies

$$g = (i(g)-1)T + t(g) \text{ and } 1 \le t(g) \le T.$$

Then proving (C.5) is equivalent to proving

$$\sum_{g=1}^{nT} \frac{1}{\sqrt{nT}} \mathbf{\Sigma}^{-1/2} \psi_{i(g),t(g)}(\theta_0, \eta_0) \xrightarrow{d} MVN(\mathbf{0}, \mathbf{I}) \text{ as } nT \to \infty. \tag{C.6}$$

Let $\mathscr{F}_0 = \{L_{1,1}\}$ and iteratively define $\{\mathscr{F}_g\}_{1 \le g \le nT}$ as follows:

$$\begin{aligned}
&\mathscr{F}_g = \mathscr{F}_{g-1} \cup \left\{A_{i(g),t(g)}, Y_{i(g),t(g)}, L_{i(g),t(g)+1}\right\}, \text{ if } \quad t(g) < T \\
&\mathscr{F}_g = \mathscr{F}_{g-1} \cup \left\{A_{i(g),T}, Y_{i(g),T}, L_{i(g)+1,1}\right\}, \text{ otherwise.}
\end{aligned}$$

By Proposition 2, we have $\mathbb{E}\left\{\psi_{i(g),t(g)}(\theta_0, \eta_0) | \mathscr{F}_{g-1}\right\} = 0$. Hence, the left-hand side of (C.6) forms a martingale difference sequence with respect to the filtration $\{\sigma(\mathscr{F}_g)\}_{g \ge 1}$, where $\sigma(\mathscr{F}_g)$ stands for the σ-algebra generated by $\mathscr{F}_g$. To show the asymptotic normality, we apply the martingale central limit theorem for triangular arrays (Theorem 2.3 of McLeish et al. [McL74]), which requires to verify the following two conditions:

(a) $\max_{1 \le g \le nT} \left\| \frac{1}{\sqrt{nT}} \mathbf{\Sigma}^{-1/2} \psi_{i(g),t(g)}(\theta_0, \eta_0) \right\|_2 \xrightarrow{P} 0$ as $nT \to \infty$.

(b) $\frac{1}{nT} \sum_{g=1}^{nT} \mathbf{\Sigma}^{-1/2} \psi_{i(g),t(g)}(\theta_0, \eta_0) \psi_{i(g),t(g)}^T(\theta_0, \eta_0) (\mathbf{\Sigma}^{-1/2})^T \xrightarrow{P} \mathbf{I}$ as $nT \to \infty$.

First, we have

$$\begin{aligned}
&\max_{1 \le g \le nT} \left\| \frac{1}{\sqrt{nT}} \mathbf{\Sigma}^{-1/2} \psi_{i(g),t(g)}(\theta_0, \eta_0) \right\|_2 \\
\le\ &\frac{1}{\sqrt{nT}} \left\| \mathbf{\Sigma}^{-1/2} \right\|_2 \cdot \max_{1 \le g \le nT} \left\| \psi_{i(g),t(g)}(\theta_0, \eta_0) \right\|_2 \\
\le\ &\frac{1}{\sqrt{nT}} \left\| \mathbf{\Sigma}^{-1/2} \right\|_2 \cdot \max_{1 \le g \le nT} \Big\{ \left\| c(\overline{V}_{i(g),t(g)}) \right\|_2 \cdot \Big| I(A_{i(g),t(g)} = 1) - \pi_0(\overline{V}_{i(g),t(g)}) \Big| \\
&\cdot \Big| Y_{i(g),t(g)} \exp\{-f_{\theta_0}(\overline{V}_{i(g),t(g)}) A_{i(g),t(g)}\} - \mu_{1,0}(\overline{V}_{i(g),t(g)}) \exp\{-f_{\theta_0}(\overline{V}_{i(g),t(g)})\} \pi_0(\overline{V}_{i(g),t(g)}) \\
&- \mu_{0,0}(\overline{V}_{i(g),t(g)}) \{1 - \pi_0(\overline{V}_{i(g),t(g)})\} \Big| \Big\} = O_p(1/\sqrt{nT}),
\end{aligned}$$

which completes the proof of (a).

To verify (b), we have

$$\left\| \frac{1}{nT} \sum_{g=1}^{nT} \mathbf{\Sigma}^{-1/2} \psi_{i(g),t(g)}(\theta_0, \eta_0) \psi_{i(g),t(g)}^T(\theta_0, \eta_0) (\mathbf{\Sigma}^{-1/2})^T - \mathbf{I} \right\|_2$$

$$= \left\| \boldsymbol{\Sigma}^{-1/2} \left\{ \frac{1}{nT} \sum_{g=1}^{nT} \boldsymbol{\psi}_{i(g),t(g)}(\theta_0, \eta_0) \boldsymbol{\psi}_{i(g),t(g)}^T(\theta_0, \eta_0) - \boldsymbol{\Sigma} \right\} (\boldsymbol{\Sigma}^{-1/2})^T \right\|_2$$
$$\leq \left\| \boldsymbol{\Sigma}^{-1/2} \right\|_2^2 \left\| \frac{1}{nT} \sum_{g=1}^{nT} \boldsymbol{\psi}_{i(g),t(g)}(\theta_0, \eta_0) \boldsymbol{\psi}_{i(g),t(g)}^T(\theta_0, \eta_0) - \boldsymbol{\Sigma} \right\|_2 .$$

It suffices to show

$$\left\| \frac{1}{nT} \sum_{g=1}^{nT} \boldsymbol{\psi}_{i(g),t(g)}(\theta_0, \eta_0) \boldsymbol{\psi}_{i(g),t(g)}^T(\theta_0, \eta_0) - \boldsymbol{\Sigma} \right\|_2 = o_P(1).$$

Define $\mathbf{M}_g = \boldsymbol{\psi}_{i(g),t(g)}(\theta_0, \eta_0) \boldsymbol{\psi}_{i(g),t(g)}^T(\theta_0, \eta_0) - \mathbb{E}\left\{ \boldsymbol{\psi}_{i(g),t(g)}(\theta_0, \eta_0) \boldsymbol{\psi}_{i(g),t(g)}^T(\theta_0, \eta_0) | \mathscr{F}_{g-1} \right\}$. Then $\{\mathbf{M}_g\}_{g \geq 1}$ forms a martingale difference sequence with respect to the filtration $\{\sigma(\mathscr{F}_g)\}_{g \geq 1}$. Since $\mathbb{E}\left(\mathbf{M}_g \mathbf{M}_{g'}^T\right) = 0$ for $g \neq g'$ and $\mathbb{E}\left(\mathbf{M}_g \mathbf{M}_g^T\right)$ is bounded for all g, we have $\left\| \frac{1}{nT} \sum_{g=1}^{nT} \mathbf{M}_g \right\|_2 \xrightarrow{P} 0$ as $nT \to \infty$. That is,

$$\left\| \frac{1}{nT} \sum_{g=1}^{nT} \left[\boldsymbol{\psi}_{i(g),t(g)}(\theta_0, \eta_0) \boldsymbol{\psi}_{i(g),t(g)}^T(\theta_0, \eta_0) - \mathbb{E}\left\{ \boldsymbol{\psi}_{i(g),t(g)}(\theta_0, \eta_0) \boldsymbol{\psi}_{i(g),t(g)}^T(\theta_0, \eta_0) | \mathscr{F}_{g-1} \right\} \right] \right\|_2 \xrightarrow{P} 0.$$

This proves (b). The proof of Theorem 4.3.1 is hence completed.

C.1.2 Proof of Theorem 4.3.2

In the case of cluster randomization, we can establish similar results as for Lemmas C.1.1 and C.1.2 under modified conditions. The remaining proof is similar to that of Theorem 4.3.1, which is omitted here.

C.2 Proof of Auxiliary Results

C.2.1 Proof of Lemma C.1.1

For any integer $1 \leq g \leq nT$, let $i(g)$ be the quotient of $g + T$ divided by T, and $t(g)$ be the integer that satisfies

$$g = (i(g) - 1)T + t(g) \text{ and } 1 \leq t(g) \leq T.$$

Let $\mathscr{F}_0 = \{L_{1,1}\}$, and $\{\mathscr{F}_g\}_{1 \leq g \leq nT}$ iteratively defined as follows:

$$\mathscr{F}_g = \mathscr{F}_{g-1} \cup \left\{ A_{i(g),t(g)}, Y_{i(g),t(g)}, L_{i(g),t(g)+1} \right\}, \text{ if } \quad t(g) < T$$
$$\mathscr{F}_g = \mathscr{F}_{g-1} \cup \left\{ A_{i(g),T}, Y_{i(g),T}, L_{i(g)+1,1} \right\}, \text{ otherwise.}$$

For any $(\theta, \eta) \in \mathscr{U}$, define

$$\mathbf{m}_g(\theta, \eta) = \boldsymbol{\psi}_{i(g),t(g)}(\theta, \eta) - \mathbb{E}\{ \boldsymbol{\psi}_{i(g),t(g)}(\theta, \eta) | \mathscr{F}_{g-1} \}.$$

Then for any fixed (θ,η), $\mathbf{m}_g(\theta,\eta)$ is a martingale difference sequence adapted to $\{\sigma(\mathscr{F}_g)\}_{g\geq 1}$. By the continuousness of $\mathbf{m}_g(\cdot)$ in θ and η and condition (C2), we have that for any $\delta>0$, there exists a finite ϵ-net $\mathscr{U}_\epsilon$ such that

$$\sup_{(\theta,\eta)\in\mathscr{U}} \frac{1}{nT}\left\|\sum_{g=1}^{nT}\mathbf{m}_g(\theta,\eta)\right\|_2 \leq \sup_{(\theta,\eta)\in\mathscr{U}_\epsilon} \frac{1}{nT}\left\|\sum_{g=1}^{nT}\mathbf{m}_g(\theta,\eta)\right\|_2 + \delta. \tag{C.7}$$

For any $(\theta,\eta)\in\mathscr{U}_\epsilon$, define

$$\mathbf{W}_{col,nT} \coloneqq \sum_{g=1}^{nT}\mathbb{E}\left\{\mathbf{m}_g(\theta,\eta)\mathbf{m}_g^T(\theta,\eta)|\mathscr{F}_{g-1}\right\} \text{ and}$$
$$\mathbf{W}_{row,nT} \coloneqq \sum_{g=1}^{nT}\mathbb{E}\left\{\mathbf{m}_g^T(\theta,\eta)\mathbf{m}_g(\theta,\eta)|\mathscr{F}_{g-1}\right\}.$$

By condition (C3), we can show that for $\forall(\theta,\eta)\in\mathscr{U}$

$$\max\left\{\|\mathbf{W}_{col,nT}\|_2, \|\mathbf{W}_{row,nT}\|_2\right\} \leq nT\sigma^2 \quad a.s.$$
$$\|\mathbf{m}_g(\theta,\eta)\|_2 \leq r \quad a.s.$$

for some constant σ^2 and r. Then by the *Martingale concentration inequality* [Tro11], we have

$$P\left\{\left\|\sum_{g=1}^{nT}\mathbf{m}_g(\theta,\eta)\right\|_2 \geq \tau\right\} \leq (1+p)\exp\left(\frac{-\tau^2}{\sigma^2 nT + r\tau/3}\right).$$

Setting $\tau=\sqrt{nT\log nT}$, we can show that the following event occurs with a probability of at least $1-O(1/nT)$ for any $(\theta,\eta)\in\mathscr{U}$,

$$\left\|\sum_{g=1}^{nT}\mathbf{m}_g(\theta,\eta)\right\|_2 \preceq \sqrt{nT\log nT}.$$

Since $\mathscr{U}_\epsilon$ is finite, we have

$$\sup_{(\theta,\eta)\in\mathscr{U}_\epsilon}\left\|\sum_{g=1}^{nT}\mathbf{m}_g(\theta,\eta)\right\|_2 \preceq \sqrt{nT\log nT},$$

and hence

$$\sup_{(\theta,\eta)\in\mathscr{U}_\epsilon}\left\|\frac{1}{nT}\sum_{g=1}^{nT}\mathbf{m}_g(\theta,\eta)\right\|_2 = O_P(\sqrt{\frac{\log nT}{nT}}) = o_P(1).$$

By (C.7), we obtain

$$\sup_{\theta\in\Theta,\eta\in\mathscr{G}_{\eta_0}}\left\|\mathbb{P}_n\psi(\theta,\eta)-\mathbb{P}\psi(\theta,\eta)\right\|_2 = o_P(1).$$

In a similar way, we can show

$$\sup_{\theta\in\Theta,\eta\in\mathscr{G}_{\eta_0}} \left\|\mathbb{P}_n\dot{\psi}_\theta(\theta,\eta)-\mathbb{P}\dot{\psi}_\theta(\theta,\eta)\right\|_2 = o_P(1).$$

C.2.2 Proof of Lemma C.1.2

We prove this lemma by applying the uniform central limit theorem for function-indexed martingale difference arrays (Theorem 2 of Bae et al. [Bae10]). Specifically, proving Lemma C.1.2 is equivalent to proving

$$\frac{1}{\sqrt{nT}}\sum_{g=1}^{nT}\mathbf{m}_g(\eta) \rightsquigarrow \boldsymbol{Z} \text{ in } l^\infty(\mathscr{G}_{\eta_0}) \text{ as } nT\to\infty.$$

Note that for any fixed η, $\mathbf{m}_g(\eta)$ is a martingale difference sequence adapted to $\{\sigma(\mathscr{F}_g)\}_{g\geq 1}$. Let $\mathbf{m}_g^{(j)}(\eta)$ denote the j-th component of $\mathbf{m}_g(\eta)$. To show the weak convergence, we need to verify the following conditions:

(a') $\mathscr{G}_{\eta_0}$ has uniformly integrable entropy.

(b') For $j=1,\cdots,p$, there exists a constant M such that as $nT\to\infty$

$$P\left(\sup_{\eta,\eta'\in\mathscr{G}_{\eta_0}} \frac{\frac{1}{nT}\sum_{g=1}^{nT}\mathbb{E}\left[\left\{\mathbf{m}_g^{(j)}(\eta)-\mathbf{m}_g^{(j)}(\eta')\right\}^2|\overline{V}_g\right]}{\|\eta-\eta'\|_{2,P}^2}\geq M\right)\to 0.$$

(c') For $j=1,\cdots,p$, $\frac{1}{nT}\sum_{g=1}^{nT}\mathbb{E}\left[\{\mathbf{m}_g^{(j)}(F)\}^2\cdot 1\left\{\frac{1}{\sqrt{nT}}\mathbf{m}_g^{(j)}(F)>\epsilon\right\}\right]\xrightarrow{P}0$ as $nT\to\infty$, where F is an envelope function for the class of functions $\{\eta:\eta\in\mathscr{G}_{\eta_0}\}$.

(d') $\frac{1}{nT}\sum_{g=1}^{nT}\mathbb{E}\left\{\mathbf{m}_g(\eta)\mathbf{m}_g^T(\eta)|\mathscr{F}_{g-1}\right\}\xrightarrow{P}\boldsymbol{\Sigma}(\eta)$ as $nT\to\infty$ for each $\eta\in\mathscr{G}_{\eta_0}$, where $\boldsymbol{\Sigma}(\eta)$ is a positive definite matrix.

Here, (a'), (b'), and (d') are Conditions (C2), (C6), and (C5), respectively, and their validity is discussed after introducing these conditions. So we only show (c'). An envelop function F can be chosen as a constant vector $(b_0,b_0,b_0)^T$ due to the boundedness condition (C3)(iii). Moreover, we can show that $\|\mathbf{m}_g(F)\|_2$ is bounded almost surely. Therefore, for $j=1,\cdots,p$,

$$\mathbb{E}\left[\{\mathbf{m}_g^{(j)}(F)\}^2\cdot 1\left\{\frac{1}{\sqrt{nT}}\mathbf{m}_g^{(j)}(F)>\epsilon\right\}\right]\preceq P\left\{\frac{1}{\sqrt{nT}}\mathbf{m}_g^{(j)}(F)>\epsilon\right\}\xrightarrow{P}0,$$

which completes the proof of (c').

C.3 Additional Simulation Results

We also conduct simulations for individual randomization with a fixed T and increasing n. The generation of covariates, treatments, and outcomes is the same as in Section 4.4. Tables C.1 and C.2 respectively

show the results for the settings with the same and opposite direction of treatment effects.

Table C.1 Simulation results with same directional treatment effect and $n = 100$ under individual randomization

Propensity score	T	Parameter	Opt. c				Sim. c			
			Mean	SD	SE	C.P.	Mean	SD	SE	C.P.
constant	1000	$\theta_1 = 0.30$	0.294	0.091	0.089	94	0.296	0.099	0.096	93
		$\theta_2 = 0.15$	0.151	0.016	0.015	94	0.151	0.017	0.016	93
		$\theta_3 = 0.05$	0.052	0.019	0.019	95	0.051	0.021	0.021	95
		$\theta_4 = 0.08$	0.080	0.039	0.040	95	0.080	0.041	0.042	95
	250	$\theta_1 = 0.30$	0.302	0.122	0.121	94	0.302	0.135	0.131	95
		$\theta_2 = 0.15$	0.149	0.027	0.026	94	0.149	0.029	0.028	95
		$\theta_3 = 0.05$	0.051	0.032	0.032	95	0.050	0.036	0.035	95
		$\theta_4 = 0.08$	0.082	0.066	0.065	95	0.082	0.072	0.069	94
	60	$\theta_1 = 0.30$	0.307	0.170	0.166	95	0.307	0.182	0.179	94
		$\theta_2 = 0.15$	0.149	0.051	0.051	95	0.150	0.054	0.055	96
		$\theta_3 = 0.05$	0.049	0.056	0.056	95	0.047	0.060	0.061	96
		$\theta_4 = 0.08$	0.084	0.114	0.111	95	0.082	0.121	0.118	95
covariates dependent	1000	$\theta_1 = 0.30$	0.301	0.095	0.095	96	0.303	0.105	0.108	96
		$\theta_2 = 0.15$	0.150	0.016	0.016	95	0.149	0.018	0.018	96
		$\theta_3 = 0.05$	0.052	0.022	0.022	95	0.052	0.026	0.026	94
		$\theta_4 = 0.08$	0.079	0.044	0.044	94	0.079	0.049	0.049	95
	250	$\theta_1 = 0.30$	0.305	0.128	0.129	96	0.306	0.147	0.145	94
		$\theta_2 = 0.15$	0.149	0.028	0.028	95	0.149	0.032	0.031	95
		$\theta_3 = 0.05$	0.050	0.036	0.037	95	0.050	0.042	0.042	94
		$\theta_4 = 0.08$	0.080	0.071	0.072	95	0.081	0.079	0.081	95
	60	$\theta_1 = 0.30$	0.309	0.173	0.175	95	0.312	0.192	0.195	96
		$\theta_2 = 0.15$	0.146	0.053	0.054	96	0.146	0.058	0.059	96
		$\theta_3 = 0.05$	0.051	0.064	0.063	95	0.051	0.074	0.071	93
		$\theta_4 = 0.08$	0.089	0.120	0.121	95	0.088	0.136	0.134	95

Table C.2 Simulation results with opposite directional treatment effect and $n = 100$ under individual randomization

Propensity score	T	Parameter	Opt. c				Sim. c			
			Mean	SD	SE	C.P.	Mean	SD	SE	C.P.
constant	1000	$\theta_1 = -0.10$	-0.099	0.090	0.090	94	-0.101	0.105	0.103	95
		$\theta_2 = -0.15$	-0.150	0.015	0.015	95	-0.150	0.018	0.017	95
		$\theta_3 = 0.05$	0.050	0.019	0.019	96	0.050	0.022	0.021	94
		$\theta_4 = 0.08$	0.080	0.039	0.040	95	0.080	0.043	0.043	95
	250	$\theta_1 = -0.10$	-0.104	0.119	0.121	95	-0.101	0.136	0.140	96
		$\theta_2 = -0.15$	-0.150	0.026	0.027	95	-0.150	0.029	0.030	95
		$\theta_3 = 0.05$	0.052	0.031	0.032	96	0.049	0.035	0.035	96
		$\theta_4 = 0.08$	0.084	0.066	0.066	95	0.082	0.071	0.070	95
	60	$\theta_1 = -0.10$	-0.104	0.164	0.166	96	-0.102	0.190	0.188	95
		$\theta_2 = -0.15$	-0.150	0.050	0.051	95	-0.150	0.056	0.057	95
		$\theta_3 = 0.05$	0.050	0.057	0.056	95	0.048	0.063	0.061	94
		$\theta_4 = 0.08$	0.079	0.112	0.111	94	0.077	0.121	0.120	95
covariates dependent	1000	$\theta_1 = -0.10$	-0.093	0.091	0.092	95	-0.090	0.098	0.099	96
		$\theta_2 = -0.15$	-0.151	0.016	0.016	95	-0.152	0.016	0.016	96
		$\theta_3 = 0.05$	0.050	0.019	0.019	95	0.051	0.019	0.019	95
		$\theta_4 = 0.08$	0.080	0.043	0.041	93	0.080	0.043	0.042	94
	250	$\theta_1 = -0.10$	-0.104	0.120	0.125	95	-0.098	0.133	0.135	95
		$\theta_2 = -0.15$	-0.150	0.026	0.027	96	-0.151	0.028	0.029	96
		$\theta_3 = 0.05$	0.051	0.031	0.032	95	0.050	0.033	0.033	95
		$\theta_4 = 0.08$	0.084	0.067	0.068	95	0.083	0.069	0.070	96
	60	$\theta_1 = -0.10$	-0.109	0.165	0.172	95	-0.110	0.181	0.185	95
		$\theta_2 = -0.15$	-0.148	0.051	0.053	96	-0.147	0.054	0.056	96
		$\theta_3 = 0.05$	0.050	0.057	0.056	95	0.051	0.061	0.058	95
		$\theta_4 = 0.08$	0.081	0.115	0.116	96	0.082	0.119	0.120	96