

Estimation Methods with Recurrent Causal Events

by

Wenling Zhang

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Statistics

Waterloo, Ontario, Canada, 2024

© Wenling Zhang 2024

Examining Committee Membership

The following served on the Examining Committee for this thesis. The decision of the Examining Committee is by majority vote.

External Examiner: **Robert W. Platt**
Professor,
Department of Epidemiology, Biostatistics, and
Occupational Health and of Pediatrics,
McGill University

Supervisors: **Cecilia A. Cotton**
Associate Professor,
Department of Statistics and Actuarial Science,
University of Waterloo

Lan Wen
Assistant Professor,
Department of Statistics and Actuarial Science,
University of Waterloo

Internal Members: **Yeying Zhu**
Associate Professor,
Department of Statistics and Actuarial Science,
University of Waterloo

Joel A. Dubin
Professor,
Department of Statistics and Actuarial Science,
University of Waterloo

Internal-External Member: **JM Gamble**
Clinical Associate Professor,
School of Pharmacy,
University of Waterloo

Author's Declaration

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

This dissertation presents a comprehensive exploration on causal effects of treatment strategies on recurrent events within complex longitudinal settings. Utilizing a series of advanced statistical methodologies, this work focuses on addressing challenges in causal inference when faced with the complexities related to various treatment strategies, recurrent outcomes and time-varying covariates that are confounded or censored.

The first chapter lays the groundwork by introducing two real-life datasets that provide a practical context for investigating recurrent causal events. In this chapter, we establish the foundation of essential concepts and terminologies. An overview of conventional causal estimands and various estimation methods in non-recurrent event settings is described, providing the necessary tools and knowledge base for effective causal analysis in more intricate longitudinal studies with recurrent event outcomes discussed in subsequent chapters.

Chapter 2 extends the traditional time-fixed measure of marginal odds ratios (MORs) to a more complex, causal longitudinal setting. The novel Aggregated Marginal Odds Ratio (AMOR) is introduced to manage scenarios where treatment varies in time and outcome also recurs. Through Monte Carlo simulations, we demonstrate that AMOR can be estimated with low bias and stable variance, when employing appropriate stabilized weight models, for both absorbing and non-absorbing treatment settings. With the 1997 National Longitudinal Study of Youth dataset, we investigate the causal effect of youth smoking on their recurrent enrollment and dropout from school, with the proposed AMOR estimator.

In the third chapter, the focus shifts to the causal effect of static treatment on recurrent event outcomes with time-varying covariates. We derive the identifying assumptions and employ a variety of estimators for the average causal effect estimation, addressing the issues of time-varying confounding and censoring. We conduct simulations to verify the robustness of these methods against potential model misspecifications. Among the proposed estimators, we conclude that the targeted maximum likelihood (TML) estimator is the appropriate one for complex longitudinal settings. Therefore, we implement targeted maximum likelihood estimation to the Systolic Blood Pressure Intervention Trial (SPRINT) dataset. Adopting an intention-to-treat analysis, we estimate the average causal effect of intensive versus standard blood pressure lowering therapy on acute kidney injury recurrences for participants surviving the first four years of SPRINT.

Chapter 4 further investigates the average causal effect of time-varying treatments on the recurrence outcome of interest with censoring. Building on the methodologies in Chapter 3, this chapter explores the singly and doubly robust estimators, especially the TML estimator, in the time-varying treatment context. Then simulation studies are conducted to support the theoretical derivations and validate the robustness of the estimators. The application of the proposed methods on the SPRINT yields some insightful findings. By incorporating participants' medication adherence levels over time as part of the treatment, we are able to investigate various adherence-related questions, and shifting from intention-to-treat to per-protocol analysis for causal effects estimation comparing the intensive versus standard blood pressure therapies.

The dissertation concludes with a summary of the main findings and a discussion of significant and promising areas for future research in Chapter 5. The studies conducted demonstrate the potential of advanced causal inference methods in handling the complexities of longitudinal data in medical and social research, offering valuable insights into how treatment strategies affect the recurrent causal outcomes over time. This work not only contributes to the theoretical advancements in statistical methodologies but also provides practical implications for the analysis of clinical trials and observational studies involving recurrent events.

Acknowledgments

First and foremost, I would like to express my deepest gratitude to my supervisors, Dr. Cecilia Cotton and Dr. Lan Wen. This dissertation would not have been possible without their knowledgeable guidance and caring support, both in research and in life. I feel immensely fortunate to have had Cecilia and Lan as my supervisors, and they will always be my role models in life.

I am sincerely grateful to Dr. Yeying Zhu, Dr. Joel Dubin, Dr. JM Gamble, and Dr. Robert Platt for their kind service as committee members for my dissertation. Their insights have greatly enriched this work.

I would also like to extend my heartfelt thanks to the professors, staff members, and fellow graduate students in our department for their indispensable help. I will always cherish the memories of being part of this supportive community.

Finally, I am deeply grateful to my friends and family for their unwavering encouragement and support during challenging times. Their belief in me has been a constant source of strength.

Table of Contents

Examining Committee Membership	ii
Author’s Declaration	iv
Abstract	v
Acknowledgments	vii
List of Figures	xiii
List of Tables	xvi
1 Introduction	1
1.1 Longitudinal studies with Recurrent Events	1
1.1.1 1997 National Longitudinal Study of Youth (NLSY97)	2
1.1.2 Systolic Blood Pressure Intervention Trial (SPRINT)	3
1.2 The Counterfactual Causal Framework	5
1.2.1 History & Setup	5
1.2.2 Assumptions	7
1.2.3 Graphical Representations	9
1.2.4 Estimands	10
1.2.5 Estimation Methods	13
1.3 Outline of the Thesis	15

2	Causal Inference for Recurrent Events via Aggregated Marginal Odds Ratio	17
2.1	Background	18
2.2	Methodology	20
2.2.1	Notation	21
2.2.2	An Intermediate Approach for Partially Time-varying Setting	24
2.2.3	Proposed Approach for the Time-varying Setting	27
2.2.4	Variations of IPW Models	31
2.3	Simulations	33
2.3.1	Results for Time-Fixed Setting	34
2.3.2	Results for Partially Time-varying Setting	37
2.3.3	Results for Time-varying Treatment Setting	40
2.4	Application	48
2.5	Discussion	50
3	Targeted Maximum Likelihood Estimation with Static Point Treatment and Longitudinal Follow-ups	52
3.1	Background	52
3.2	Methodology	53
3.2.1	Notation	53
3.2.2	Definition and Identification of Causal Estimands	55
3.2.3	Inverse probability weighted (IPW) estimator: $\hat{\psi}_{\text{IPW}}$	58
3.2.4	Hájek inverse probability weighted (HIPW) estimator: $\hat{\psi}_{\text{HIPW}}$	59
3.2.5	Iterative conditional expectation (ICE) estimator: $\hat{\psi}_{\text{ICE}}$	59
3.2.6	Augmented inverse probability weighted (AIPW) estimator: $\hat{\psi}_{\text{AIPW}}$	61
3.2.7	Weighted iterative conditional expectation (WICE) estimator: $\hat{\psi}_{\text{WICE}}$	62
3.2.8	Targeted maximum likelihood (TML) estimator: $\hat{\psi}_{\text{TML}}$	63

3.3	Simulation	64
3.3.1	Model Specifications	66
3.3.2	Results	67
3.4	Application	74
3.4.1	SPRINT Data	74
3.4.2	Effect Estimand	75
3.4.3	Data Coarsening	75
3.4.4	Results	78
3.5	Conclusion	78
4	Targeted Maximum Likelihood Estimation with Time-varying Treatments	79
4.1	Background	79
4.2	Methodology	80
4.2.1	Notation	80
4.2.2	Inverse probability weighted (IPW) estimator: $\hat{\psi}_{\text{IPW}}$	82
4.2.3	Hájek inverse probability weighted (HIPW) estimator: $\hat{\psi}_{\text{HIPW}}$	83
4.2.4	Iterative conditional expectation (ICE) estimator: $\hat{\psi}_{\text{ICE}}$	83
4.2.5	Augmented inverse probability weighted (AIPW) estimator: $\hat{\psi}_{\text{AIPW}}$	85
4.2.6	Weighted iterative conditional expectation (WICE) estimator: $\hat{\psi}_{\text{WICE}}$	85
4.2.7	Targeted maximum likelihood (TML) estimator: $\hat{\psi}_{\text{TML}}$	86
4.3	Simulation	87
4.3.1	Model Specifications	89
4.3.2	Results	90
4.4	Application	98
4.4.1	Data Description	98
4.4.2	Data Transformation	99

4.4.3	Total Per-Protocol Effect of the Intensive vs. Standard Blood Pressure Lowering Therapy	100
4.5	Conclusion	101
5	Discuss	103
5.1	Conclusion	103
5.2	Future Work	105
5.2.1	Extensions of AMOR	105
5.2.2	Improvements on Data Pre-processing	106
5.2.3	Dynamic Treatment Strategies	107
5.2.4	Methods for Unmeasured Confounding	109
5.2.5	Competing events and Separable Effects	110
5.2.6	Enhancements for the current TML estimator	111
	References	114
	APPENDICES	126
A	Weight Estimation and Detailed Simulation Process for Chapter 2	127
A.1	Simulation Details for Section 2.3.1	127
A.2	Weight Estimation for Section 2.3.1	128
A.3	Simulation Details for Section 2.3.2	129
A.4	Weight Estimation for Section 2.3.2	129
A.5	Simulation Details for Section 2.3.3	131
A.6	Weight Estimation for Section 2.3.3	134
B	Supplementary Information for Chapter 2	138
B.1	Definition and Identification of Causal Estimands	138
B.2	Weight Distributions for Simulation Studies	139
B.3	Additional Results for Simulation Studies	141
B.4	Variable Information for the Real Data Example in Section 2.4	144

C	Supplementary Information for Chapter 3	146
C.1	Identification	146
C.1.1	Without Censoring	146
C.1.2	With Censoring	148
C.2	EIF Derivation for $\psi_{I,k}$	150
C.2.1	Without Censoring	151
C.2.2	With Censoring	151
C.3	Generalized form for some of the estimators	158
C.3.1	Iterative conditional expectation (ICE) estimator: $\hat{\psi}_{ICE}$	158
C.3.2	Augmented inverse probability weighted (AIPW) estimator: $\hat{\psi}_{AIPW}$	159
C.3.3	Weighted iterative conditional expectation (WICE) estimator: $\hat{\psi}_{WICE}$	160
C.3.4	Targeted maximum likelihood (TML) estimator: $\hat{\psi}_{TML}$	161
D	Supplementary Information for Chapter 4	162
D.1	Identification	162
D.2	EIF for $\psi_{II,k}$	165
D.3	Generalized procedures for some estimators	168
D.3.1	Iterative conditional expectation (ICE) estimator: $\hat{\psi}_{ICE}$	168
D.3.2	Augmented inverse probability weighted (AIPW) estimator: $\hat{\psi}_{AIPW}$	169
D.3.3	Weighted iterative conditional expectation (WICE) estimator: $\hat{\psi}_{WICE}$	170
D.3.4	Targeted maximum likelihood (TML) estimator: $\hat{\psi}_{TML}$	170
D.4	Supplementary Application Results with the SPRINT dataset	171
D.4.1	Total Intention-to-treat Effect of Intensive Blood Pressure Ther- apy	171
D.4.2	Population Intervention Effect of Adherence	173
D.4.3	Summery Table for Results	175

List of Figures

1.1	Example of a causal acyclic diagram illustrating the relationships between time-fixed treatment A and outcome Y , with a potential confounder X .	9
1.2	Example of a single world intervention graph (SWIG) that identifies a (static) causal treatment effect from the realized treatment or intervention a to the potential outcome Y^a with a potential confounder X .	10
2.1	Causal graph illustrating the relationships between baseline covariates, time-fixed treatment and outcome for the setting specified in Section 2.2.1	23
2.2	Causal graph that illustrates the relationship between the baseline covariates, the time-varying treatment, and the time-fixed outcome for the setting specified in Section 2.2.2	26
2.3	Causal graph illustrating the relationship between the baseline covariates, time-varying treatment, and the recurrent outcome for the setting specified in Section 2.2.3	29
2.4	Simulation-Based Estimation for Different Weight Models (WM): Time-fixed Treatment and Time-fixed Outcome	35
2.5	Simulation-Based Estimation for Different Weight Models (WM): Partially Time-varying Setting	38
2.6	Simulation-Based Estimation for Different Weight Models (WM): the original Nonabsorbing Treatment Time-varying Setting Specified in Section 2.3.3	41
2.7	Simulation-Based Estimation for Different Weight Models (WM): the Absorbing Treatment Time-varying Setting Specified in Section 2.3.3	42

2.8	Simulation-Based Estimation for Different Weight Models (WM): the Nonabsorbing Treatment Time-varying Scenario with Comparative Settings Specified in Section 2.3.3	44
2.9	Causal graph illustrating the relationship among exposures ($A(t)$), events ($Y(t)$), and baseline covariates classified by X_{trt} , X_{out} and X_{cfd} for $t = 0, 1, \dots, 5$	49
3.1	A single world intervention graph (SWIG) that identifies the average causal effect and illustrates the relationship between the measured baseline variables L_0 , the static treatment A , the censoring indicator $\bar{C}_{k+1}$, the time-varying variable $\bar{L}_k$ and the (intermediate) recurrent outcome $\bar{Y}_{k+1}$, assuming no unmeasured confounding. Note that for any variable X_l in the counterfactual setting, the superscript $\bar{c}_k = \bar{0}$ in $X_l^{\bar{c}_k = \bar{0}}$ is suppressed as $\bar{c} = 0$, implying X_l under history up to and include the most recent c_k before X_l ($k \leq l$). e.g., $Y_k^{\bar{c}=0} = Y_k^{\bar{c}_k = \bar{0}}$	57
3.2	A single world intervention graph (SWIG) that identifies the average causal effect and illustrates the relationship between the measured baseline variables L_0 , the initial treatment assignment (A), the censoring status $\bar{C}_2$, the confounders $\bar{L}_1$, and the (intermediate) recurrent outcome $\bar{Y}_2$, assuming there is no unmeasured confounding.	65
4.1	A single world intervention graph (SWIG) that illustrates the connections and dependencies among the key variables: L_0 (the baseline covariates), A_{k-1} (the time-varying treatment and the censoring indicator in interval k), L_{k-1} (the time-varying covariate and the intermediate recurrent outcome in interval k), and Y_{K+1} (the last intermediate outcome), where $k \in \{1, \dots, K+1\}$	82
4.2	A single world intervention graph (SWIG) that illustrates the relationship between the key variables and identifies the causal estimand of $Y_2^{\bar{a}_1}$ with $k = 2$ time intervals	88
4.3	A single world intervention graph (SWIG) that identifies the average causal effect and illustrates the relationships between the measured baseline variables L_0 , the joint treatment variable of initial treatment assignment & time-varying intervention of adherence ($Z, \bar{A}_1$), the semi-competing event of death $\bar{D}_2$, the time-varying confounder L_1 , and the (intermediate) recurrent outcome $\bar{Y}_2$, assuming there is no unmeasured confounding.	100

5.1	A single world intervention graph (SWIG) that identifies the causal effect and illustrates the relationship between the measured baseline variables L_0 , the joint treatment variable of initial treatment assignment & time-varying intervention of adherence $(Z, \bar{A}_1)$, the semi-competing event of death $\bar{D}_2$, the confounders $\bar{L}_1$, and the (intermediate) recurrent outcome $\bar{Y}_2$, assuming there is no unmeasured confounding.	108
5.2	A single world intervention graph (SWIG) identifies the separable effects of some time-varying treatment intervention A_k on the semi-competing event D_{k+1} and the (intermediate) recurrent outcome Y_{k+1} , with the measured time-varying confounder L_k for $k \in \{0, \dots, K\}$, assuming no unmeasured confounding. Note that we use red arrows to indicate the hypothetical separation of the treatment effect.	112
B.1	Estimated Weight Distributions when $T = 4$ Specified in Section 2.3.3	140
C.1	A single world intervention graph (SWIG) that identifies the average causal effect and illustrates the relationship between the measured baseline variables L_0 , the static treatment $A = a$, the time-varying variable $\bar{L}_k$ and the (intermediate) recurrent outcome $\bar{Y}_{k+1}$, assuming no unmeasured confounding.	147
D.1	A single world intervention graph (SWIG) that identifies the causal treatment effect and illustrates the relationship between the measured baseline variables L_0 , the initial treatment assignment (Z) , the competing event of death $\bar{D}_2$, the confounders $\bar{L}_1$, and the (intermediate) recurrent outcome $\bar{Y}_2$, assuming there is no unmeasured confounding.	172
D.2	A single world intervention graph (SWIG) that identifies the (conditional) population intervention effect and illustrates the relationship between the measured baseline variables L_0 , time-varying intervention of participants' adherence $\bar{A}_1$, the competing event of death $\bar{D}_2$, the confounders $\bar{L}_1$, and the (intermediate) recurrent outcome $\bar{Y}_2$, assuming there is no unmeasured confounding.	174

List of Tables

2.1	Simulation-Based Estimation of MOR with Different Weight Models: the Time-fixed Setting	36
2.2	Simulation-Based Estimation of Time-varying Treatment MOR _{2trt} Using Different Weight Models: the Partially Time-varying Setting . . .	39
2.3	Simulation-Based Estimation of AMOR with Different Weight Models: the original Nonabsorbing Treatment Time-varying Setting Specified in Section 2.3.3	45
2.4	Simulation-Based Estimation of AMOR with Different Weight Models: the Absorbing Treatment Time-varying Setting Specified in Section 2.3.3	46
2.5	Simulation-Based Estimation of the AMOR with Different Weight Models: the Nonabsorbing Treatment Time-varying Scenario with Comparative Setting specified in Section 2.3.3	47
3.1	Average of key time-varying variables (Y_1 , Y_2 , L_1 and R_2) under null effect	68
3.2	Average of key time-varying variables (Y_1 , Y_2 , L_1 and R_2) under moderate effect	68
3.3	Average of key time-varying variables (Y_1 , Y_2 , L_1 and R_2) under strong effect	69
3.4	Table of estimation results under null effect based on 1,000 sets of 5,000 simulated individuals with six estimators: IPW, HIPW, ICE, AIPW, WICE and TML respectively	71
3.5	Table of estimation results under moderate effect based on 1,000 sets of 5,000 simulated individuals with six estimators: IPW, HIPW, ICE, AIPW, WICE and TML respectively	72

3.6	Table of estimation results under strong effect based on 1,000 sets of 5,000 simulated individuals with six estimators: IPW, HIPW, ICE, AIPW, WICE and TML respectively	73
3.7	Baseline Characteristics for the SPRINT Study Population of 9,322 Subjects Included in the Analysis	76
3.8	Event Count of Acute Kidney Injury Episode within Time Interval 1 and 2 for the Intensive and Standard Therapy Arms ($A = 1$ and $A = 0$ Respectively)	76
3.9	Average of Key Time-varying Variables in the first interval ($t = 1$) or in the second interval ($t = 2$) for the SPRINT Study Population Included in the Analysis	77
4.1	Average of key time-varying variables (Y_1 , Y_2 , L_1 and R_2) under null effect	91
4.2	Table of estimation results under null effect based on 2,000 sets of 2,000 simulated individuals and 1,000 Bootstrap samples with six estimators: IPW, HIPW, ICE, AIPW, WICE and TML respectively	93
4.3	Average of key time-varying variables (Y_1 , Y_2 , L_1 and R_2) under moderate effect	94
4.4	Table of estimation results under moderate effect based on 2,000 sets of 2,000 simulated individuals with six estimators: IPW, HIPW, ICE, AIPW, WICE and TML respectively	95
4.5	Average of key time-varying variables (Y_1 , Y_2 , L_1 and R_2) under strong effect	96
4.6	Table of estimation results under strong effect based on 2,000 sets of 2,000 simulated individuals with six estimators: IPW, HIPW, ICE, AIPW, WICE and TML respectively	97
4.7	The Average or Count of Key Time-varying Variables in the First Interval ($t = 1$) or in the Second Interval ($t = 2$) for the SPRINT Study Population Included in the Analysis	101
A.1	Table of $\alpha_{0,1}$ (for the original setting), $\alpha_{0,2}$ (for the comparative setting) and β_{trt} (in terms of AMOR) Values for the Time-varying Nonabsorbing Treatment settings	136

A.2	Table of α_0 and β_{trt} (in terms of AMOR) values for the Time-varying Absorbing treatment settings	136
A.3	Table comparing the five settings specified in Sections 2.3.1, 2.3.2, and 2.3.3.	137
B.1	Simulation-Based Estimation of AMOR with Different Weight Models: Non-absorbing Time-varying Treatment, $T = 10$	141
B.2	Simulation-Based Estimation of AMOR with Different Weight Models: Absorbing Treatment Case with $T = 10$	142
B.3	Simulation-Based Estimation of AMOR with Different Weight Models: Non-absorbing Treatment with Comparative Setting, $T = 10$	143
D.1	Table of estimated causal effects with the SPRINT	176

Chapter 1

Introduction

1.1 Longitudinal studies with Recurrent Events

Recurrent events, occurring repeatedly over time, hold considerable importance for policymakers, health professionals, and statistical methodologists. For instance, the study of hospitalization frequencies in children can guide pediatric care policies [49], while examining the incidence of acute kidney injuries among elderly patients can enhance geriatric care protocols [48].

In exploring the causal effects on recurrent events, it necessitates a clear understanding of the nature of treatment strategies, which can be categorized as either deterministic or stochastic and further as static or dynamic.

This dissertation primarily focuses on static, deterministic treatments, which involve consistent application of interventions based on established guidelines without adjustments for individual variations or random elements. However, the potential of dynamic treatments, which adapt interventions based on ongoing assessments and introduce random components to account for uncertainty in treatment outcomes, is also acknowledged [107, 112].

Deterministic strategies are predictable and uniform, applying the same treatment protocol across all patients [41]. Static, deterministic treatments are defined by a uniform application of the same treatment protocol across all subjects over the study period, without adaptation to changing conditions or individual responses. Examples include uniform sentencing in criminal justice, standardized educational curricula, or fixed vaccination schedules [22, 59, 90, 113].

Dynamic treatment strategies, on the other hand, are tailored to individual differences and adapt to changing circumstances over time. In other words, these intervention strategies introduce past patient characteristics into the treatment decision process, allowing for adaptations based on real-time feedback and individual patient circumstances. Dynamic treatment strategies, despite their higher implementation costs, can offer advantages such as reducing resource waste and providing more precise treatment tailored to individual patient needs, particularly in managing chronic diseases and other complex clinical scenarios [50, 113, 61]. A notable example is combined antiretroviral therapy (cART) used to manage human immunodeficiency virus (HIV) infection, which adapts based on past CD4 cell counts and diagnoses of related illnesses, demonstrating a dynamic treatment regimen [112].

Observational studies and randomized controlled trials (RCTs) are two primary approaches used in research to explore causal relationships and evaluate intervention effects. Observational studies, such as cohort studies or panel studies [95], follow individuals or groups over an extended period, collecting data at multiple time points to analyze relationships between exposures or interventions and outcomes. While these studies offer valuable insights into real-world scenarios and the influence of various factors on outcomes, they are susceptible to biases like confounding and selection bias, potentially undermining the validity of findings. In contrast, RCTs are experimental studies where participants are assigned to different treatment groups with fixed or known probabilities, enabling researchers to assess intervention effects while minimizing bias [16]. Note that even though RCTs are distinct from observational studies, sometimes they both can be included within a longitudinal research.

This dissertation will utilize two datasets of longitudinal studies, to apply and compare the effects of static treatment strategies on recurrent events, thereby demonstrating their implications for policy and care protocols.

1.1.1 1997 National Longitudinal Study of Youth (NLSY97)

The National Longitudinal Study of Youth (NLSY97) [13] is an ongoing cohort study that aims to follow a sample of American youth born between 1980 and 1984. Initially comprised of 8,984 respondents who were between the ages of 12-17 at the first interview in 1997, the NLSY97 offers a rich, detailed account of the transition from adolescence into adulthood across various facets of life.

The NLSY97 utilized a series of surveys (the “survey instrument” [58]) for data collection, employing distinct questionnaires tailored to different respondents: the

Screeners, Household Roster, and Nonresident Roster Questionnaire for household informants; the Youth Questionnaire for the NLSY97 respondents themselves; and the Parent Questionnaire for the respondents' parents. This method ensured a broad and detailed collection of data from multiple perspectives within the respondent's environment. In subsequent rounds (rounds 2-5), the Youth Questionnaire continued to be the primary tool for gathering information directly from the participants, supplemented by a Household Income Update from 1998 to 2001 to capture household income data from a parent.

The questionnaires covered a wide array of topics organized into general categories, such as Education, Training & Achievement Scores; Employment; Household, Geography & Contextual Variables; Parents, Family Process & Childhood; Dating, Marriage & Cohabitation; Sexual Activity, Pregnancy & Fertility; Children; Income, Assets & Program Participation; Health; Attitudes, Expectations, Non-Cognitive Tests, Activities; and Crime & Substance Use. The diversity of subject areas and variables covered in the survey questionnaires ensures a comprehensive overview of the participants' lives and experiences.

In the dataset of NLSY97, we are interested in understanding the potential (causal) relationships between American youth's smoking status and school enrollment. Smoking status is a time-varying treatment variable defined by self-reported smoking within the last 30 days before each annual interview. School enrollment status, serving as the recurrent event outcome, is defined by an indicator of current enrollment in educational institutions (high school or college) or the achievement of a college degree, as youths may exit and re-enter the educational system. This focus aligns with the broader objectives of the NLSY97 to capture the evolving nature of various life domains, including health and education, as participants navigate the critical transition from adolescence into adulthood.

1.1.2 Systolic Blood Pressure Intervention Trial (SPRINT)

The Systolic Blood Pressure Intervention Trial (SPRINT) [111] is an open-label clinical trial funded by the U.S. National Institute of Health. It was designed to assess the effects of intensive versus standard blood pressure management in adults aged 50 or older. The cohort consisted of 9,361 subjects who were selected based on their high systolic blood pressure (≥ 130 mm Hg), elevated risk of cardiovascular diseases, and absence of diabetes. At the outset of the study, participants were randomly allocated to either a standard treatment group, with a goal of reducing

systolic blood pressure to less than 140 mm Hg, or an intensive treatment group, targeting a systolic blood pressure of less than 120 mm Hg.

Running from November 2010 to July 2016, SPRINT provided critical insights into the management of hypertension in older adults at risk for cardiovascular events. With a median follow-up period of 3.26 years, it was observed that the intensive treatment group experienced higher rates of acute kidney injury episodes compared to their standard treatment counterparts. However, this same group also saw a reduction in all-cause mortality rates [111].

Our interest lies in examining the causal effect of medication treatments related to the blood pressure therapies on the recurrence of acute kidney injury episodes. The occurrence of acute kidney injury as a recurrent outcome in the context of hypertension management presents a complex challenge of balancing the benefits of blood pressure control via medications against the risks associated with increased acute kidney injury episodes. Understanding this balance is crucial for guiding clinical decisions, especially in populations at elevated risk for both cardiovascular disease and kidney injury. We seek to delve deeper into the dynamics of the variables and potentially inform more effective approaches to hypertension management in older adults.

In the context of discussing the SPRINT dataset, it is important to introduce two critical methodological concepts in clinical trials: Intent-to-Treat analysis and Per-Protocol analysis.

Intent-to-Treat (ITT) is a principle in RCTs where all participants are included in the analysis according to the group they were originally assigned to, regardless of whether they completed the study, adhered to the treatment protocol, or switched to a different treatment [38]. The ITT analysis preserves the benefits of randomization, ensuring that the comparison between treatment levels remains unbiased by post-randomization events [29, 62].

A Per-Protocol (PP) analysis, on the other hand, investigates the hypothetical causal outcomes that would have occurred if all subjects had strictly complied with their assigned treatment as specified in the trial protocol. [38, 69]. This approach focuses on estimating the treatment effect under ideal conditions where participants fully comply with the intervention as planned.

Compared to ITT analysis, the results of a PP analysis may be of greater interest to patients who intend to adhere closely to their prescribed treatments [41]. Nevertheless, both ITT and PP analyses can provide valuable insights depending on the specific research question of interest [69].

It is worth noting that conventional naive PP analysis differs from As-Treated (AT) analysis. The AT analysis evaluates the effect based on the intervention participants actually received. In contrast, the naive PP analysis focuses solely on participants who strictly adhere to their assigned protocol, excluding those who deviate from the study protocol [93]. However, in this thesis, appropriate measures are taken to adjust for confounding and selection bias for our PP analysis. This approach aligns with the counterfactual model framework, which will be discussed in the following section.

1.2 The Counterfactual Causal Framework

While real life datasets of NLSY97 and SPRINT exemplify the challenge of inferring causality in longitudinal studies, our goal is to discern the potential cause(s) behind the patterns observed in these studies. By adopting the counterfactual theories of causation [41, 85], we aim to frame these causal questions within a broader context, seeking not only to understand the specific phenomena at hand but also to uncover broader implications and potential solutions grounded in this framework.

1.2.1 History & Setup

The basic concept of counterfactuals was initially proposed by the 18-th century philosopher David Hume [3]. Then, in the 1970s, the framework was further expanded and attracted more attention from researchers in other fields [9]. With its recent rapid development, the field of causal inference saw formal integration into a statistical context by Donald Rubin in the 1970s [85]. This foundational work set the stage for significant expansions and refinements by subsequent scholars. Notably, Paul Holland contributed to the clarification and formalization of causal inference principles, especially regarding the concept of causality in statistics [44]. Judea Pearl introduced a graphical approach to causal inference in the 2000s, further enriching the methodological toolkit available for causal analysis [64]. Building on these advancements, Miguel Hernán and James Robins have crucially advanced the operationalization of counterfactual concepts in epidemiology and public health, enhancing the framework’s applicability to real-world contexts [41], especially in medical studies. For example, causal models have been used to study the effects of different interventions on preventing depression symptoms [104], aspirin on cardiovascular mortality [17], and iron supplement use during pregnancy on anemia [11].

To better understand causality, it is essential to understand how outcomes are influenced by treatment [43]. The counterfactual framework provides a rigorous method for conceptualizing what might have happened to individuals in a population under different interventions. Within this framework, the outcome under a hypothetical treatment intervention being studied is called the “potential outcome” [87] or “counterfactual outcome”. For introductory illustration purposes, consider a time-fixed setting. This will later be extended to more complex settings in the thesis. Consider binary treatment and outcome variables, denoted by A and Y respectively. Let a represent the realization of A , where $a = 1$ or 0 signify whether or not an individual is exposed to a specific intervention. Similarly, $Y = 1$ indicates the presence of a specified outcome, and $Y = 0$ indicates its absence. The potential outcomes under intervention and control conditions are denoted as $Y^{a=1}$ and $Y^{a=0}$, respectively.

A fundamental challenge within the counterfactual framework is the inherent limitation that for any given individual, indexed by i , we can observe only one of the potential outcomes [41]. For individual i , denote the observed outcome as Y_i and the realized values of treatment A_i as a_i ($a_i = 0, 1$), respectively. If the individual is actually treated, then we have the observed outcome $Y_i = Y_i^{a_i=1}$; otherwise, we have $Y_i = Y_i^{a_i=0}$, but not both. Such limitation prevents us from drawing causal conclusions directly, as the two potential outcomes $Y_i^{a_i=1}$ and $Y_i^{a_i=0}$ do not co-exist in the same reality for comparison. Nevertheless, we are often interested in estimating the counterfactual mean outcomes rather than individual counterfactual outcomes in real world examples. Based on these counterfactual outcome means, we would be able to estimate the average causal effects (ACE) by comparing the expected potential outcome of a treated population with that of the same population but under control: $E(Y^{a=1}) - E(Y^{a=0})$, where Y^a denotes the potential outcome of the target population under treatment a . By making appropriate causal assumptions [14], ACEs in the potential outcome framework can then be identified as estimable statistical quantities. We will discuss these assumptions in Section 1.2.2.

It is crucial to understand the impact of confounders in causal inference. Confounders are variables that are associated with both the treatment/exposure and the outcome of interest. In epidemiology settings, for example, children’s hospitalisation history could potentially be associated with both their current vaccination statuses (the exposure) and future hospitalisations due to infection (the recurrent outcome of interest) [49]. Ignoring the impact of confounding will lead to bias in causal effect estimation [36]. To address the challenges posed by confounding in causal effect identification and estimation, we first assume that all confounders in the causal relationship have been observed, as part of the causal assumptions in Section 1.2.2.

Various methods have been developed to control for the measured confounding effect, which we will discuss in Section 1.2.5.

1.2.2 Assumptions

Positivity, consistency, exchangeability, and the stable unit treatment value assumption (SUTVA) are often referred to as key causal assumptions in the causal framework to identify causal effects [41]. The details of these key assumptions of causal inference are discussed below.

Positivity

The assumption of positivity refers to the condition that there is a positive probability of receiving every level of treatment in the population, given the observed covariates. This assumption is crucial for estimating causal effects across the entire range of treatment levels since it guarantees that there is information available on the outcome under each treatment condition. In other words, causal effects, such as the ACE, generally cannot be estimated in a subset of the study population if everyone is exposed (or unexposed) in that subset. Mathematically, it can be expressed as, for every combination of values of treatment and covariate levels that occur in the population i.e., $0 < P(A = a|X = x) < 1$, for a and x as any possible level of A and X . The positivity assumption can be easily controlled in randomized experiments, by randomly assigning positive treatment probability to each stratum of subjects. In observational studies, we can assess potential near violations to positivity assumption (in theory) by checking the covariate distributions in the treatment groups [118].

Consistency

In statistical theory, the “consistency” of an estimator means that the estimator converges in probability to the true value of its corresponding parameter [25]. However, the definition is different in the context of causal inference. The consistency assumption is defined as the condition that the potential outcome for an individual under the specific treatment level they actually receive is the same as the observed outcome [14]. In other words, the consistency assumption requires that if any individual i in the population is observed to have treatment level $A_i = a_i$, its corresponding observed outcome Y_i is precisely the potential outcome under a_i , i.e., $Y_i = Y_i^{a_i}$. For a study with binary treatment, the consistency assumption states

that $Y_i = a_i Y_i^{a_i=1} + (1 - a_i) Y_i^{a_i=0}$. Additionally, the consistency assumption implies a well-defined treatment such that each treatment level has only one version for all subjects receiving it [41]. This foundational assumption links the theoretical potential outcome to the observed data in reality. Without the consistency assumption, one cannot identify the actual cause of the outcome, which means that the estimation results would be meaningless. Testing for consistency, however, is inherently challenging since potential outcomes under unobserved treatment conditions are, by definition, not observable.

Exchangeability

The term exchangeability (or ignorability) refers to the condition where the assignment of treatment is independent of the potential outcomes, given the observed covariates [41]. Exchangeability is crucial to infer causation from association, as it enables the comparison of the exposed and unexposed groups as if they were randomized experiments [40]. That is, we assume that given all observed confounders X , the potential outcome Y^a is not affected by the treatment assignment A , or $Y^a \perp A | X$ for any a . Even though this conditional independence assumption is weaker than the marginal independence condition, (i.e., $Y^a \perp A$ for any a), only the former is needed to obtain an unbiased estimate of the causal effect [80]. However, exchangeability can be difficult to test directly because it involves counterfactual outcomes. Researchers often assess it by examining the plausibility of no unmeasured confounding variables, which can suggest the extent to which this assumption may hold. Meanwhile, if we consider censoring or lost-to-followup in our study, a similar exchangeability assumption is used for censoring - we assume that censoring is conditionally independent of the potential outcome given past measured covariates.

Stable Unit Treatment Value Assumption (SUTVA)

The Stable Unit Treatment Value Assumption (SUTVA) is the “*no interference between units*” condition, which states that the treatment assignment of a subject does not affect the potential outcome of any other subject in the population. In formal terms, for any subject i in the population, its potential outcome under treatment realization a_i is independent of the treatment assignments of any other units, i.e., $Y_i^{a_i} \perp A_j$, for any subject $j \neq i$ in the population. This concept does not apply to scenarios where an intervention on one individual influences the risk of outcomes for others, as seen in infectious disease models, which would be a violation of the SUTVA. [31, 63].

These four causal assumptions are fundamental in causal inference that facilitates the interpretation and estimation of causal effects. The first three assumptions, in particular, link the causal effects described in potential outcome framework to the data observed in reality, thereby enabling the identification of causal relationships from observational data. Violation of these assumptions may lead to biased causal effect estimation or even meaningless causal analysis [68, 92]. We will make the four assumptions discussed above throughout this thesis.

1.2.3 Graphical Representations

In this section, we introduce graphical representations that illustrate the causal relationships among a set of variables. These variables, depicted as nodes within the graph, are connected through directed edges (arrows) that signify causal influence from one variable to another [60]. Figure 1.1 shows an example of a simple directed acyclic diagram (DAG) [41]. This graphical framework provides a visual tool to understand causal directions and the role of confounding variables that may effect the relationships between the treatment and outcome variables of interest.




Figure 1.1: Example of a causal acyclic diagram illustrating the relationships between time-fixed treatment A and outcome Y , with a potential confounder X .

Expanding on the concept of DAGs, Single-World Intervention Graphs (SWIGs) offer a refined perspective by explicitly representing interventions and their effects on outcomes [72]. Figure 1.2 shows an example of a SWIG. The SWIGs graphically “implement” interventions by “cutting” the arrows leading into the intervened variable (A), and setting it to a specific intervention level (a). This extension of causal graphs to include interventions provides a critical link between theoretical causal models and the empirical estimation of causal effects, offering a clear visualization of how interventions could lead to different outcomes.

The SWIGs are instrumental in the conceptualization of counterfactual outcomes—what the outcome for an individual or a group would be under different intervention scenarios. They offer a graphical approach for assessing potential violations of the exchangeability assumption [72], which is central to the validity of causal inferences.




Figure 1.2: Example of a single world intervention graph (SWIG) that identifies a (static) causal treatment effect from the realized treatment or intervention a to the potential outcome Y^a with a potential confounder X

1.2.4 Estimands

Building on the foundation of the causal assumptions, we can now explore the various types of causal estimands, which are the specific effects we seek to measure in order to understand the impact of an intervention or treatment. We will begin our discussion of the causal contrast by first examining differences, followed by an exploration of ratios.

Causal Differences

The average causal effect (ACE), also referred to as the average treatment effect, is a fundamental concept in causal inference literature to quantify the causal impact of a treatment (or intervention) on an outcome of interest, compared to a control condition in the “parallel world” [41]. As we briefly mentioned above, the ACE quantifies the expected difference in potential outcomes of the entire population under two scenarios: treated and untreated, expressed as

$$E(Y^{a=1}) - E(Y^{a=0}),$$

where $E(Y^{a=1} = 1)$ is the expected potential outcome had the population been exposed to treatment, and $E(Y^{a=0} = 1)$ is the expected potential outcome had the population not been exposed to treatment.

The conditional average treatment effect (CATE), on the other hand, refines the ACE by estimating the effect of treatment for specific subgroups within the population, based on certain characteristics or conditions. Denote the condition by X , it can be expressed by

$$E(Y^{a=1}|X) - E(Y^{a=0}|X),$$

where $E(Y^a|X)$, $a = 0, 1$ represents the conditional expected outcome under treatment level a given covariate(s) X . The CATE is particularly useful when the effectiveness of an intervention varies across different segments of the population, allowing for more targeted and efficient policy or treatment implementations [47].

A special case of the CATE is the average treatment effect on the treated (ATT), which focuses on the average effect of the treatment specifically in those who were actually treated, with the expression

$$E(Y^{a=1}|A = 1) - E(Y^{a=0}|A = 1).$$

These extensions of the ACE allow for a more nuanced understanding of causal effects, especially when the treatment effect is not uniform across all individuals. The ATT is particularly useful for evaluating the effectiveness of a treatment among those who receive it, but it may not be generalizable to the entire population.

A slight variant to traditional ACE, the population intervention effect (PIE) is a concept used in epidemiology and public health to quantify the overall impact of applying an intervention across an entire population compared to the observed state of that population [109]. This concept is articulated as the difference between the expected outcome if the entire population were exposed to the intervention and the observed average outcome, i.e.,

$$E(Y^{a=1}) - E(Y).$$

It is particularly valuable in scenarios where it is impractical or unethical to withhold treatment from any group of individuals in the study population, making traditional ACE less relevant. For example, McQuade *et al.* (2022) [82] estimated the PIE to show a significant positive impact of enhanced sanitation practices on children's growth. In this case, an ACE estimate may overestimate the potential positive effect of the exposure, as the comparative scenario of an unexposed population is generally not of interest. Therefore, the PIE estimate is more favorable than the ACE estimate because it compares the outcome under full exposure with the outcome under exposure level observed in reality.

Causal Ratios

The Risk Ratio (RR), also known as the Relative Risk (RR) quantifies the likelihood of an outcome occurring in one group compared to another and can be interpreted as the factor by which the risk of the outcome changes with exposure to the treatment.

In the causal inference context, it is calculated as the ratio of the probability of the potential outcome in the treatment group to the probability of the potential outcome in the control group. Mathematically, it can be represented as

$$\frac{P(Y^{a=1} = 1)}{P(Y^{a=0} = 1)},$$

where $P(Y^{a=1} = 1)$ is the probability of the outcome had the population been exposed to treatment, and $P(Y^{a=0} = 1)$ is the probability of potential outcome had the population not been exposed to treatment. The RR is especially valuable in situations where the baseline risk of the outcome is well-defined and consistent across groups, and when the study design involves prospective cohort studies or randomized controlled trials (RCTs). In these cases, the RR provides a straightforward interpretation of the treatment effect as the relative change in risk due to the intervention.

An odds ratio (OR) is defined to compare the odds of outcome under different levels of treatment through a ratio. Although its concept may not be straightforward to grasp for the general public like the RR, the OR is commonly used to quantify effects because it can provide a more stable estimate, especially in case-control studies where the outcome prevalence is low [66, 4, 117, 110, 24].

In the context of causal inference, the OR is defined by the ratio between the odds of the potential outcome if all were treated and the odds of the potential outcome if all were untreated, where odds of an event is represented by the probability of having the event divided by one minus that probability. Specifically, the OR can be expressed as

$$\frac{P(Y^{a=1} = 1)/(1 - P(Y^{a=1} = 1))}{P(Y^{a=0} = 1)/(1 - P(Y^{a=0} = 1))}.$$

Note that the numerator in the expression is the odds of the outcome given a particular exposure; and the corresponding denominator is the odds of the outcome in the absence of the exposure for time-fixed binary counterfactual treatment and outcome.

The Hazard Ratio (HR) from survival analysis compares the event rates between treatment and control groups [35]. Without the causal context, it represents the ratio of the hazard (event rate) in the treatment group to the hazard in the control group at any given point in time. A HR greater than 1 indicates a higher risk of the event in the treatment group, while a HR less than 1 suggests a lower risk. However, since the HR implicitly conditions on survival up to a certain time point, the groups of survivors being compared (the treatment group versus the control group) may evolve over the study period. Such changes can be influenced by the treatment itself

or other external factors, leading to a skewed representation of the true treatment effect. This phenomenon, known as selection bias, prevents the HR from representing a consistent treatment effect over time [114, 37].

1.2.5 Estimation Methods

After exploring various estimands that quantify causal effects, it is essential to examine methods for estimating these effects. Given the focus on recurrent events in this thesis, advanced techniques for time-varying scenarios become particularly relevant. Robins’ generalization methods, also known as g-methods, aim to identify and consistently estimate causal effects in complex data structures [74]. Both marginal structural models and g-formulas are integral components of the g-methods family.

Marginal structural models (MSMs), as proposed by Robins *et al.* [80], provide a framework to model a counterfactual outcome by considering its marginal means across time, or over different treatment strategies, while accounting for time-varying confounders [39, 103]. We will delve into these type of models in Chapter 2.

The g-formula, also known as the generalized computation algorithm formula, or the g method that models the joint distribution of the treatments, outcomes and confounders across time [73, 75]. Under certain key causal assumptions, the g-formula can be used to identify counterfactual outcomes and estimate the causal effect with the presence of confounders in some longitudinal settings [112]. It is crucial to note that the g-formula can be represented as different estimable forms that are algebraically equivalent [107]. Specifically, two representations of the g-formula will be of our interest in this thesis: one is a weighted expectation and the other is an iterative conditional expectation, with corresponding estimators. We will discuss both estimation methods in this section, followed by possibilities to improve them, with the introduction of the influence function.

Inverse Probability of treatment Weighting (IPW)

Balancing covariates, including potential confounders, across different treatment groups is essential to obtain unbiased causal effect estimates. Central to covariate balancing is the concept of the propensity score, which is essentially the probability of a participant receiving a particular treatment, given a set of observed covariates [84]. Within strata of the same propensity score, it has been shown that treated and untreated units with the same propensity score will have similar distributions of

baseline covariates [5]. Popular propensity score methods include matching, stratification, covariate adjustment, and **inverse probability of treatment weighting (IPW)** [4, 23, 27]. Previous work has suggested that the methods of matching, stratifying, and covariate adjustment all yield positively biased estimators and poor confidence interval coverage in OR estimation in simulation settings [4, 23], while one can obtain essentially unbiased estimator with the IPW method [54]. Indeed, the IPW, originated from the Horvitz-Thompson estimator [45], is frequently employed to balance the covariates and adjust for potential confounders [5]. For longitudinal data, we estimate the propensity scores at each time point (person-time) [39, 15]. Then each subject will be assigned the weight of a cumulative product of its inverse propensity score computed at each time point, which can also be referred to as the “cumulative weight.” The use of IPW essentially creates a pseudo-population where treatment assignment is statistically independent of baseline and time-varying confounders [14]. By adjusting for the probability of receiving treatments, IPW helps mimic a randomized trial scenario, thus addressing confounding effectively.

Iterative Conditional Expectation (ICE)

The **iterative conditional expectation (ICE)** estimator is a parametric g-formula used for estimating causal treatment effect in complex longitudinal settings involving time-varying treatments in epidemiology [67, 99, 107, 106]. As the name suggests, the ICE method estimates the conditional expectation by iteratively standardizing over past measured covariates. For longitudinal data, ICE starts with the outcome at the end of the study period and iteratively estimates conditional expectations of the outcome given covariate history backwards in time. The advantage of this estimator is that it does not require parametric specifications for the joint distribution of the entire observed data.

Based on the g-formula, an IPW estimator and an ICE estimator would require different parametric models to be correctly specified. They are referred to as singly robust estimators, meaning that they only offer one chance of attaining consistency. More specifically, an IPW estimator is consistent if the propensity score model is correctly specified; while an ICE estimator requires the true parameterization for the series of conditional outcome regression models.

Efficient influence function (EIF)-based Estimation

With efficient influence functions, we aim to improve estimation with more efficiency and robustness of the aforementioned IPW and ICE estimators. In robust statistics, an influence function measures how sensitive an estimator is responding to the changes to each individual data point. In particular, an efficient influence function (EIF) is defined to be the influence function that achieves the semiparametric efficiency bound [32, 100]. The derivation of EIFs can be approached in mainly three ways: (i) taking the pathwise derivative along some parametric submodel; (ii) taking the Gateaux derivative in the direction of some contamination model [42] and (iii) projecting on the tangent space of some influence function [101]. For this thesis, we will primarily employ approach (i) for EIF derivations.

After deriving on the EIF, we are able to construct estimators with desirable asymptotic properties including doubly robustness and efficiency guarantees. Specifically, it has been demonstrated that, in some longitudinal settings, we can extend from basic IPW and ICE estimators to more advanced forms such as the **augmented IPW (AIPW)** estimator and the **weighted ICE (WICE)** estimator, based on different EIF representations [106]. The AIPW, for instance, combines the strengths of IPW and regression adjustment through the EIF, allowing it to remain consistent if either the model for treatment or the outcome model is correctly specified, thus enhancing its robustness. Similarly, the WICE estimator applies additional weights to the ICE procedure, augmenting its confounding adjustment capabilities [8].

Similar to other EIF-based doubly robust estimators, the **targeted maximum likelihood estimation (TMLE)** algorithm [102, 101] is a semi-parametric estimation approach that adjusts the initial model estimates with a target step based on the influence function to reduce bias and enhance estimation efficiency. Furthermore, TMLE is designed to be fully adaptive to the data, utilizing machine learning algorithms to model complex relationships within the data effectively. This adaptability makes TMLE especially suitable for epidemiological studies and other areas of health research where the relationships between variables are often complex and influenced by numerous confounders.

1.3 Outline of the Thesis

In this thesis, we tackle estimation challenges related to recurrent causal events. In Chapter 2, we introduce a novel statistical measure for quantifying the causal effect

of time-varying treatments on recurrent event outcomes, which has been published in *Statistics in Medicine* [116] and is presented here in its entirety. Moving forward, Chapter 3 delves into exploring the average causal effect of a static point treatment on recurrent events within longitudinal settings with the presence of censoring, highlighting the efficacy of the targeted maximum likelihood estimation method over alternative estimators. Chapter 4 extends our methodology to encompass recurrent causal events with time-varying treatment strategies, where the targeted maximum likelihood estimator is developed tailored to this context. We aim to submit Chapters 3 and 4 for publication in academic journals in the near future. Finally, in Chapter 5, we draw conclusions from our findings, discussing limitations, potential extensions beyond proposed methods, and charting out avenues for future research.

Chapter 2

Causal Inference for Recurrent Events via Aggregated Marginal Odds Ratio

Researchers often work with treatments and outcomes that vary over time. For example, psychologists are interested in the curative effect of cognitive behavior therapies on patients' recurrent depression symptoms. While there are various causal effect measures designed for one-time treatment, the causal effect measures for time-varying treatment and recurrent events are relatively under-developed. In this chapter, a new causal measure is proposed to quantify the causal effect of time-varying treatments on recurrent events. We suggest estimators with robust standard errors that are based on various weight models for both conventional causal measures and the proposed measure in different time settings. We outline the approaches and describe how using some stabilized inverse probability weight models are more advantageous than others. We demonstrate that the proposed causal estimand can be estimated with low bias and stable variance for study periods of moderate length, and the estimation results are compared under different treatment settings with various weight models. We also find that the proposed method is suitable for both absorbing and nonabsorbing treatments. The methods are applied to the 1997 National Longitudinal Study of Youth as an illustrative example.

2.1 Background

In scientific or medical studies where subjects experience the same outcome multiple times, researchers often seek to understand the relationship between certain factors and the so-called recurrent event. It is very common to experience recurrent events in the fields of public health and social studies. For example, unvaccinated children tend to have more frequent hospitalizations than the vaccinated ones; youth may leave and re-enter school multiple times before graduation [49, 89]. From the standard statistical point of view, this problem can often be simplified to estimations of association between certain covariates and the outcome (event) of interest that vary over time. Therefore, statisticians have developed systematic methods for analysis with recurrent event data [18]. Meanwhile, it is known that association cannot guarantee causation in statistics, which makes causal inference an appealing topic. The counterfactual framework of causation marked an important milestone in the development of causal inference [85, 44, 64]. The framework aims to capture the causal relationship by comparing the theoretical outcomes of the population given different levels of treatment, holding everything else the same. However, it is impossible to simultaneously observe the outcomes of the same population corresponding to different treatment arms. The framework defines treatment arm that is opposed to reality as the counterfactual treatment, and the subsequent outcome corresponding to the theoretical treatment to be the potential outcome. Certain assumptions must be verified or accepted to infer from association to causation and achieve a valid causal inference [76].

Positivity, consistency, exchangeability, and stable unit treatment value assumption (SUTVA) are often referred to as key causal assumptions in the causal framework. The assumption of positivity requires a positive probability of receiving each level of treatment for each combination of treatment values and confounder levels that occur among individuals in the population. The consistency assumption states that the observed outcome of any subject in the population must always coincide with the potential outcome under the observed treatment level. The assumption of exchangeability or ignorability refers to the condition that there are no unmeasured confounders. In other words, the potential outcome is independent of the treatment assignment, given all observed confounders considered. The SUTVA states that each subject's treatment has no effect on the potential outcome of any other subject. Violation of these assumptions will lead to biased causal effect estimation or even meaningless causal analysis [68, 92]. In observational studies, we can assess the positivity assumption by checking the covariate distributions in the treatment groups

[118]. It should also be verified that there is no known violation of the untestable assumptions of consistency, exchangeability and SUTVA before proceeding with causal analysis.

In epidemiological studies, researchers usually need to deal with confounding associated with the causal effect of interest. For example, researchers have investigated the causal effects of different interventions on preventing depression symptoms, aspirin on cardiovascular mortality, and iron supplement use during pregnancy on anemia [11, 17, 104]. When there exists risk factor(s) that are associated with both treatments and outcomes, ignoring its effect will lead to bias in causal effect estimation [36]. Therefore advanced methods must be used to control for confounding in a proper causal analysis [43]. The marginal structural models (MSM) proposed by Robins *et al.* [80] yields less biased estimation in the presence of time-varying confounders in observational studies [39, 103]. The method of MSM first balance all other covariates (including confounders) among different treatment groups via the so-called inverse probability weighting (IPW) method. Then, the treatment effect is estimated with an appropriate model using the reweighted population. Note that the standard use of MSM requires both the treatment and outcome models to be correctly specified [76].

The simplest form of IPW assigns each person with a weight that is either the inverse of their estimated probability of treatment if they receive the treatment, or the probability of no treatment if they are not treated. For longitudinal data, we calculate the weight accordingly for each person at each time point (person-time). The weight is also referred to as the “nonstabilized weight” or “unstabilized weight” [15, 39]. It is possible for the unstabilized weights to have extreme values if the denominators are around zero, which may lead to large standard errors for the estimated treatment effect [15]. Unlike unstabilized weights, stabilized weights usually include the estimated conditional probability of receiving the observed treatment (or no treatment) in the numerator. In longitudinal studies, the increment in time yields more extreme weights in an unstabilized weight model. Thus, the stabilized weight, or stabilized IPW, is often preferred due to its protection against extreme weights [39, 15, 49]. Meanwhile in longitudinal studies with time-varying data, the weight assigned to each subject will be a cumulative product of its weights computed at each time point, which is later referred to as the “cumulative weight”.

It is important to select an appropriate measure of the causal effect on a particular outcome of interest. An odds ratio (OR) is defined to compare the odds of outcome under different levels of treatment through a ratio. Although its concept may not be straightforward to grasp for the general public, the OR is commonly used to quantify

such an effect in the fields of medical and health studies due to its unique design. With logistic regression being the standard model for binary outcomes, the OR can be conveniently derived from the estimates of a logistic regression and interpreted as the strength of the (causal) effect of interest [4, 117, 110]. Nevertheless, the definition of the OR limits its usage in complex studies of time-varying treatment and recurrent events. We aim to extend the concept of the OR to create a causal measure that preserves its advantages but can be used in longitudinal settings.

In the case of recurrent events, Robins *et al.* [79] have demonstrated how to use the g-computation algorithm and generalized estimation equations (GEE) to estimate the causal effect, marginal mean, of time-dependent treatment on binary outcomes under derived conditions. They also briefly discuss the estimation of the marginal mean using the MSMs. To adjust for time-varying confounding, MSMs with IPW have been implemented to estimate causal effects based on various causal measures for recurrent outcomes in observational studies [49, 34]. This method is shown to be able to consistently estimate the marginal hazard ratio for recurrent events when censoring is absent [53]. Our goal is to develop a simple yet effective causal estimand for time-varying treatment and recurrent outcomes using MSMs.

The remainder of this article is organized as follows. In Section 2.2, we outline methods to estimate causal treatment effects in basic time settings with time-fixed outcome, and propose a causal measure for time-varying treatments and recurrent outcomes. We present in Section 2.3 the results from a series of Monte Carlo simulations, which examine the performance of the estimation methods under different weight models and various time settings. Next (Section 2.4), we illustrate an application of the concepts and estimators using a real-life data example. Finally, Section 2.5 discusses the possible modifications that can be made to broaden the definition of the proposed causal measure and optimize its usage.

2.2 Methodology

In this section, we discuss the methods to estimate conventional and new causal effect measures for three different time settings. We first outline the standard approach to estimate a causal effect measure for time-fixed treatment and outcome. In Section 2.2.2, we describe an intermediate setting in which we estimate the marginal causal effect with a time-varying treatment and a time-fixed outcome. In Section 2.2.3, we introduce a new causal measure for a time-varying treatment and recurrent events. We also propose five IPW models that differ in their numerators for the re-weighting

process in estimation. The notation in this chapter is based on Robins[76] and Austin [6]. Capital letters represent random variables and their corresponding lower-case letters represent realizations of their possible values.

2.2.1 Notation

As discussed in Section 2.1, the OR is a common measure of the association between a binary treatment and a binary outcome. It can be further characterized as a marginal odds ratio (MOR) or a conditional odds ratio (COR) defined as follows. Denote A and Y as an arbitrary treatment and the outcome of interest in a study. Suppose both A and Y are binary, the MOR is defined as

$$\text{MOR} = \frac{P(Y = 1|A = 1)}{1 - P(Y = 1|A = 1)} \bigg/ \frac{P(Y = 1|A = 0)}{1 - P(Y = 1|A = 0)}. \quad (2.1)$$

Let $\mathbf{X}$ be a vector of all other covariates that are associated with A and/or Y , then the COR is defined as

$$\text{COR} = \frac{P(Y = 1|A = 1, \mathbf{X} = \mathbf{x})}{1 - P(Y = 1|A = 1, \mathbf{X} = \mathbf{x})} \bigg/ \frac{P(Y = 1|A = 0, \mathbf{X} = \mathbf{x})}{1 - P(Y = 1|A = 0, \mathbf{X} = \mathbf{x})}.$$

In general, the COR can be empirically estimated from the observed samples, but it does not necessarily equal the MOR when the causal relationship is confounded. Since the causal effect of treatment A is our focus, the MOR is more desirable than the COR in observational studies where confounders are usually present. If the causal assumptions given in Section 2.1 hold for a particular set of observed data, the MOR is identifiable and will have a causal interpretation.

Under the causal framework, we use the superscript to represent a potential condition that is possibly counterfactual. For example, $Y^{a=1}$ represents the potential outcome when subjects are treated regardless of whether subjects are actually treated or not in reality. We define a Causal OR as the ratio of the odds of the potential outcome if all were treated and the odds of the potential outcome if all were untreated, where the odds of an event are represented by the probability of having the event divided by one minus that probability [117]. That is,

$$\text{Causal OR} = \frac{P(Y^{a=1} = 1)}{1 - P(Y^{a=1} = 1)} \bigg/ \frac{P(Y^{a=0} = 1)}{1 - P(Y^{a=0} = 1)}, \quad (2.2)$$

In observational studies, one can only attempt to estimate causal associations after examining the key causal assumptions of the counterfactual framework (discussed in Section 2.1). The simulation process should be designed so that the key causal assumptions are not violated. The Causal OR of interest is then equivalent to the conventional statistical measure, the MOR. Based on Austin[4] and Forbes and Shortreed[23], we outline the method to derive the MOR empirically from a large dataset and estimate the MOR of a time-fixed treatment on the outcome of interest. Assume that we have a study with M samples from a population (M is extremely large). Let A_i and Y_i be the binary treatment and outcome of the individual i ($i = 1, 2, \dots, M$). For simplicity, assume that treatment and outcome have Bernoulli distributions with subject-specific probabilities that are characterized by logistic regression models. Specifically, suppose that their parametric forms are as follows.

$$A_i \sim \text{Bernoulli}(p_{i,\text{treat}}),$$

where

$$\text{logit}(p_{i,\text{treat}}) = \alpha_0 + \mathbf{X}_{i,\text{treat}}^T \alpha_{\text{treat}} + \mathbf{X}_{i,\text{cfd}}^T \alpha_{\text{cfd}}, \quad (2.3)$$

and

$$Y_i \sim \text{Bernoulli}(p_{i,\text{out}}),$$

where

$$\text{logit}(p_{i,\text{out}}) = \beta_0 + \beta_A A_i + \mathbf{X}_{i,\text{out}}^T \beta_{\text{out}} + \mathbf{X}_{i,\text{cfd}}^T \beta_{\text{cfd}}. \quad (2.4)$$

The vector $\mathbf{X}_{i,\text{treat}}$ denotes the covariates that are only associated with the treatment (exposure) A_i . The vector $\mathbf{X}_{i,\text{cfd}}$ represents another vector of covariates that are associated with both A_i and Y_i ; and $\mathbf{X}_{i,\text{out}}$ is a vector of covariates that are only associated with the outcome (events) Y_i . Let the realizations of $\mathbf{X}_{i,\text{treat}}$, $\mathbf{X}_{i,\text{cfd}}$, $\mathbf{X}_{i,\text{out}}$, A_i , and Y_i be $\mathbf{x}_{i,\text{treat}}$, $\mathbf{x}_{i,\text{cfd}}$, $\mathbf{x}_{i,\text{out}}$, a_i , and y_i , respectively. Assume that each covariate belongs to only one of $\mathbf{X}_{i,\text{treat}}$, $\mathbf{X}_{i,\text{cfd}}$, and $\mathbf{X}_{i,\text{out}}$. Any nuisance covariates are ignored. The parameters α_0 , α_{treat} , α_{cfd} , β_0 , β_A , β_{out} , and β_{cfd} are the corresponding covariate coefficients or coefficient vectors. The parameters $p_{i,\text{treat}}$ and $p_{i,\text{out}}$ are the subject-specific probabilities for treatment and outcome, respectively. Figure 2.1 shows the relationships implied by Equations (2.3) and (2.4).

Although the value of the COR is simply the exponential of β_{treat} , the value of the MOR remains unknown. Austin[4] demonstrated an empirical method to obtain the MOR value using the Markov chain Monte Carlo (MCMC) simulations. For each




Figure 2.1: Causal graph illustrating the relationships between baseline covariates, time-fixed treatment and outcome for the setting specified in Section 2.2.1

subject i , the subject-specific probability of experiencing an outcome if the subject are treated, $p_{i,\text{out}1}$, is represented by

$$\text{logit}(p_{i,\text{out}1}) = \beta_0 + \beta_A + \mathbf{X}_{i,\text{out}}^T \beta_{\text{out}} + \mathbf{X}_{i,\text{cfd}}^T \beta_{\text{cfd}}.$$

And if the subject is not treated, the probability of outcome, $p_{i,\text{out}0}$, is represented by

$$\text{logit}(p_{i,\text{out}0}) = \beta_0 + \mathbf{X}_{i,\text{out}}^T \beta_{\text{out}} + \mathbf{X}_{i,\text{cfd}}^T \beta_{\text{cfd}}.$$

Then, the MOR defined by Equation (2.1) can be empirically represented by

$$\text{MOR} = \frac{\dot{p}_1 / (1 - \dot{p}_1)}{\dot{p}_0 / (1 - \dot{p}_0)}, \quad (2.5)$$

where $\dot{p}_1$ is the average probability of the outcome had all subjects treated ($\dot{p}_1 = \frac{1}{M} \sum_{i=1}^M p_{i,\text{out}1}$), and $\dot{p}_0$ is the average probability of the outcome had all subjects not treated ($\dot{p}_0 = \frac{1}{M} \sum_{i=1}^M p_{i,\text{out}0}$). For a sufficiently large sample size M , this empirical MOR converges to the true value of the MOR [4]. Note that later we use the data size of N (in contrast with M) to differentiate the process of estimating the MOR using MSM with IPW from the process of empirically deriving the MOR via MCMC described above.

Recall from Section 2.1, the IPW method adjusts for confounding by assigning subject-specific weights to each subject and producing a new pseudopopulation [7]. For the i -th subject, $i = 1 \dots N$, its (unstabilized) weight w_i is the reciprocal of the

probability that the i -th subject received their observed treatment given all covariates that are associated with the treatment. That is,

$$w_i = \frac{1}{P(A_i = a_i | \mathbf{X}_{i,\text{treat}} = \mathbf{x}_{i,\text{treat}}, \mathbf{X}_{i,\text{cfd}} = \mathbf{x}_{i,\text{cfd}})}.$$

One method to estimate the weight denominator is regressing the realization of treatment a_i on the baseline covariates $\mathbf{x}_{i,\text{treat}}$ and $\mathbf{x}_{i,\text{cfd}}$ using model (2.3) to obtain the predicted probability of treatment conditional on the related covariates $\hat{p}_{i,\text{treat}}$:

$$\text{logit}(\hat{p}_{i,\text{treat}}) = \hat{\beta}_0 + \hat{\mathbf{X}}_{i,\text{treat}}^T \hat{\beta}_{\text{treat}} + \mathbf{X}_{i,\text{cfd}}^T \hat{\beta}_{\text{cfd}}.$$

Therefore, the estimated (unstabilized) IPW is the inverse of $\hat{p}_{i,\text{treat}}$, if the i -th subject is observed to be treated; or the inverse of $(1 - \hat{p}_{i,\text{treat}})$, if not treated. That is,

$$\hat{w}_i = \frac{1}{\hat{p}_{i,\text{treat}} a_i + (1 - \hat{p}_{i,\text{treat}})(1 - a_i)} = \frac{a_i}{\hat{p}_{i,\text{treat}}} + \frac{1 - a_i}{1 - \hat{p}_{i,\text{treat}}}.$$

Note that other variations of the weight model will be discussed in Section 2.2.4.

To estimate the marginal treatment effect, regress the outcome Y on the treatment A via a weighted logistic regression with the estimated weights of the form

$$\text{logit}(\hat{p}_{i,\text{out}}) = \hat{\beta}_{0,\text{marg}} + \hat{\beta}_{\text{treat, marg}} a_i. \quad (2.6)$$

2.2.2 An Intermediate Approach for Partially Time-varying Setting

As an intermediate approach, suppose that a binary treatment is performed at two time points, we define the causal two-course treatment OR, Causal OR_{int.}, by comparing the distribution of potential outcomes under treatment at two time points with the distribution of potential outcomes without treatment. Specifically,

$$\text{Causal OR}_{\text{int.}} = \frac{P(Y^{(A(1)=1, A(2)=1)} = 1)}{1 - P(Y^{(A(1)=1, A(2)=1)} = 1)} \bigg/ \frac{P(Y^{(A(1)=0, A(2)=0)} = 1)}{1 - P(Y^{(A(1)=0, A(2)=0)} = 1)}, \quad (2.7)$$

and

$$\text{MOR}_{\text{int.}} = \frac{P(Y = 1 | A(1) = 1, A(2) = 1)}{1 - P(Y = 1 | A(1) = 1, A(2) = 1)} \bigg/ \frac{P(Y = 1 | A(1) = 0, A(2) = 0)}{1 - P(Y = 1 | A(1) = 0, A(2) = 0)}.$$

Meanwhile, suppose that X is a vector of all other covariates that are associated with A and/or Y at time t . Then, the conditional counterpart of $\text{MOR}_{\text{int.}}$ can be represented by

$$\text{COR}_{\text{int.}} = \frac{P(Y = 1|A(1) = 1, A(2) = 1, X = x)}{1 - P(Y = 1|A(1) = 1, A(2) = 1, X = x)} \bigg/ \frac{P(Y = 1|A(1) = 0, A(2) = 0, X(t) = x)}{1 - P(Y = 1|A(1) = 0, A(2) = 0, X = x)}.$$

Next, we demonstrate how to empirically derive the value of $\text{MOR}_{\text{int.}}$ via MCMC. The setting is an extension of Austin's work [6] where the treatments and outcomes for M subjects are assumed to follow Bernoulli distributions and the vectors $\mathbf{X}_{i,\text{treat}}$, $\mathbf{X}_{i,\text{cfd}}$ and $\mathbf{X}_{i,\text{out}}$ also follow the same definitions as in Section 2.2.1. Suppose that $A(1)$ does not directly affect Y , but affects Y through $A(2)$. For subject i , $i = 1, \dots, N$, assume

$$A_i(1) \sim \text{Bernoulli}(\pi(1)_{i,\text{treat}})$$

where

$$\text{logit}(\pi(1)_{i,\text{treat}}) = \alpha_{0,t1} + \mathbf{X}_{i,\text{treat}}^T \alpha_{\text{treat},t1} + \mathbf{X}_{i,\text{cfd}}^T \alpha_{\text{cfd},t1},$$

and

$$A_i(2) \sim \text{Bernoulli}(\pi(2)_{i,\text{treat}})$$

where

$$\text{logit}(\pi(2)_{i,\text{treat}}) = \alpha_{0,t2} + \alpha_{1,t2}A_i(1) + \mathbf{X}_{i,\text{treat}}^T \alpha_{\text{treat}} + \mathbf{X}_{i,\text{cfd}}^T \alpha_{\text{cfd},t2},$$

and

$$Y_i \sim \text{Bernoulli}(\pi_{i,\text{out}})$$

where

$$\text{logit}(\pi_{i,\text{out}}) = \beta_0 + \beta_{A(2)}A(2)_i + \mathbf{X}_{i,\text{out}}^T \beta_{\text{out}} + \mathbf{X}_{i,\text{cfd}}^T \beta_{\text{cfd}}. \quad (2.8)$$

The vector $\mathbf{X}_{i,\text{treat}}$, $\mathbf{X}_{i,\text{cfd}}$, and $\mathbf{X}_{i,\text{out}}$ are defined the same way as in Section 2.2.1. The parameters $\alpha_{0,t1}$, $\alpha_{\text{treat},t1}$, $\alpha_{\text{cfd},t1}$, $\alpha_{0,t2}$, $\alpha_{\text{treat},t2}$, $\alpha_{\text{cfd},t2}$, β_0 , $\beta_{A(1)}$, $\beta_{A(2)}$, β_{out} , and β_{cfd} are the corresponding covariate coefficients or vectors of covariate coefficients. The parameters $\pi(1)_{i,\text{treat}}$, $\pi(2)_{i,\text{treat}}$, and $\pi_{i,\text{out}}$ are the subject-specific probabilities for treatment at times $t = 1$ and 2 and the outcome, respectively. The process is visualized in Figure 2.2.

Assuming that all β 's are known, the value of $\text{MOR}_{\text{int.}}$ remains unknown, and $\text{COR}_{\text{int.}} = \exp(\beta_{A(2)})$. Parallel to Section 2.2.1, MCMC can be used to empirically obtain $\text{MOR}_{\text{int.}}$. The subject-specific probability of the outcome, $\pi_{i,\text{out}}$, is modeled by




Figure 2.2: Causal graph that illustrates the relationship between the baseline covariates, the time-varying treatment, and the time-fixed outcome for the setting specified in Section 2.2.2

Equation (2.8). Then, the marginal time-varying treatment Causal OR is empirically represented by

$$\text{MOR}_{\text{int.}} = \frac{\dot{\pi}_1 / (1 - \dot{\pi}_1)}{\dot{\pi}_0 / (1 - \dot{\pi}_0)},$$

where $\dot{\pi}_1$ is the average probability of the outcomes had all subjects treated at both time points ($\dot{\pi}_1 = \frac{1}{M} \sum_{i=1}^M \pi_{i,\text{out}}$ with $A(1) = 1$ and $A(2) = 1$), and $\dot{\pi}_0$ is the average probability of the outcomes had all subjects not treated at both time points ($\dot{\pi}_0 = \frac{1}{M} \sum_{i=1}^M \pi_{i,\text{out}}$ with $A(1) = 0$ and $A(2) = 0$). For a sufficiently large M , this empirically derived value can be used to represent the true $\text{MOR}_{\text{int.}}$.

To employ the IPW approach for confounding in causal models, subject-specific weights are defined for all individuals at all time points. Specifically, subject i 's (unstabilized) weights at time $t = 1$ and time $t = 2$ are

$$w_i(1) = \frac{1}{P(A_i(1) = a_i(1) | \mathbf{X}_{i,\text{treat}} = \mathbf{x}_{i,\text{treat}}, \mathbf{X}_{i,\text{cfd}} = \mathbf{x}_{i,\text{cfd}})},$$

and

$$w_i(2) = \frac{1}{P(A_i(2) = a_i(2) | A_i(1) = a_i(1), \mathbf{X}_{i,\text{treat}} = \mathbf{x}_{i,\text{treat}}, \mathbf{X}_{i,\text{cfd}} = \mathbf{x}_{i,\text{cfd}})}.$$

After the weights are defined and estimated, the outcome is regressed in the two-course treatment with a weighted logistic model:

$$\text{logit}(\hat{\pi}_{\text{out}}) = \hat{\beta}_{0, \text{ marg, 2trt}} + \hat{\beta}_{\text{treat, marg, 2trt}} A(2). \quad (2.9)$$

Finally, in the counterfactual framework, the estimated causal treatment effect can be represented by $\exp(\hat{\beta}_{\text{treat, marg, 2trt}})$.

2.2.3 Proposed Approach for the Time-varying Setting

Introduction of the AMOR

In the previous sections, we have discussed the definition and estimation of causal effect measures in the fixed-time setting and the partially time-varying setting. In this section, we propose a causal measure that quantifies the marginal causal effect of time-varying treatments on recurrent outcomes over discrete time points. Under the counterfactual framework, the aggregated marginal odds ratio (AMOR) is designed to represent the odds of having the outcome at every time point when all were exposed to the treatment at all time points, versus the odds of having the outcome when all were never treated at any time points. The odds of an outcome refer to the ratio between the probability of experiencing the outcome had a given treatment history, versus one minus this probability. Let $A(t)$ be the binary treatment (or exposure) status at time t , and $Y(t)$ be the binary outcome (or event) at time $t = 1, \dots, T$. The outcome is assumed to be fully observed (no loss to follow-up), and $Y(t)$ is temporally immediately after $A(t)$ at each time point. Let $\bar{A}(t)$ denote the complete treatment history up to and including time t , that is, $\bar{A}(t) = (A(1), A(2), \dots, A(t))$. Similarly, let $\bar{Y}(t)$ represent the complete outcome history until time t , that is, $\bar{Y}(t) = (Y(1), Y(2), \dots, Y(t))$. The definition of AMOR is as follows.

$$\text{AMOR} = \frac{P(\bar{Y}(T)^{(\bar{A}(T)=\bar{1})} = \bar{1})}{1 - P(\bar{Y}(T)^{(\bar{A}(T)=\bar{1})} = \bar{1})} \bigg/ \frac{P(\bar{Y}(T)^{(\bar{A}(T)=\bar{0})} = \bar{1})}{1 - P(\bar{Y}(T)^{(\bar{A}(T)=\bar{0})} = \bar{1})},$$

where $P(\bar{Y}(T)^{(\bar{A}(T)=\bar{1})} = \bar{1})$ is the probability that the outcome occurs at all time points under treatment at all time points; and $P(\bar{Y}(T)^{(\bar{A}(T)=\bar{0})} = \bar{1})$ is the probability that the outcome occurs at all time points had there been no treatment at all time points. Suppose that the key causal assumptions hold, we can focus on the simpler

version of the AMOR that only involves association:

$$\text{AMOR}_{\text{association}} = \frac{P(\bar{Y}(T) = \bar{1} | (\bar{A}(T) = \bar{1}))}{1 - P(\bar{Y}(T) = \bar{1} | (\bar{A}(T) = \bar{1}))} \bigg/ \frac{P(\bar{Y}(T) = \bar{1} | (\bar{A}(T) = \bar{0}))}{1 - P(\bar{Y}(T) = \bar{1} | (\bar{A}(T) = \bar{0}))},$$

The definition of AMOR is a natural extension of the static causal measure MOR. The connection between the MOR and the AMOR is that they both aim to measure the marginal causal effect in the sense that the result is not conditioned on any other measured covariates. Moreover, both describe the causal effect with a comprehensive statistic. While the MOR only concerns the treatment and the outcome without considering the time component, the defined AMOR is more suitable for longitudinal settings where treatment and outcome both vary over a moderate length of time. For example, we can use the MOR to measure the causal effect of an antiviral drug on the recovery of patients with COVID-19, while the AMOR can be used to determine the causal effect of daily salbutamol treatment daily on asthma attacks within a week.

In the above setting, the status of treatment at each time point can be 0 or 1 with a positive probability. However, for certain medical treatments such as kidney transplantation or wisdom tooth removal, once treatment occurs, the subject will remain treated for the remaining time of the study period. We define this type of treatment as an “absorbing treatment”, in contrast to the previously described “nonabsorbing treatment” scenario.

For any population under an absorbing treatment, AMOR can be interpreted as the ratio of odds of having the outcome recurrently over time when the population is treated from the start of the study period, versus the odds when the population is never treated. Regardless of the treatment being absorbing or not, the AMOR still measures the causal effect on the recurrent outcome. The difference in the treatment property will also be addressed in the estimation process discussed next.

Estimation of the AMOR

We describe the following scenario to demonstrate AMOR estimation. Let the realizations of $A(t)$ and $Y(t)$ be $a(t)$ and $y(t)$, respectively. The vectors $\mathbf{X}_{i,\text{treat}}$, $\mathbf{X}_{i,\text{cfd}}$ and $\mathbf{X}_{i,\text{out}}$ inherit the definitions in Section 2.2.2. Suppose that $A(0)$ and $Y(0)$ follow Bernoulli(p_0) and are considered as baseline covariates for $A(t)$ and $Y(t)$. Then, assume that M subjects have the following parametric forms for treatment and outcome models at time t ($t = 1, \dots, T$). For subject i , $i \in \{1, \dots, M\}$:

$$\pi(t)_{i,\text{treat}} = P(A_i(t) = 1 | \mathbf{X}_{i,\text{treat}} = \mathbf{x}_{i,\text{treat}}, \mathbf{X}_{i,\text{cfd}} = \mathbf{x}_{i,\text{cfd}}, \bar{A}_i(t-1), \bar{Y}_i(t-1));$$




Figure 2.3: Causal graph illustrating the relationship between the baseline covariates, time-varying treatment, and the recurrent outcome for the setting specified in Section 2.2.3

and

$$\pi(t)_{i,\text{out}} = P(Y_i(t) = 1 | \mathbf{X}_{i,\text{out}} = \mathbf{x}_{i,\text{out}}, \mathbf{X}_{i,\text{cfd}} = \mathbf{x}_{i,\text{cfd}}, \bar{A}_i(t), \bar{Y}_i(t-1)).$$

For simplicity, assume that only on the one-step history of previous treatment and outcome will be included in the models. The processes of $A(t)$ and $Y(t)$ are visualized in Figure 2.3.

Starting at $t = 1$, we assume that the following models hold:

$$A_i(t) \sim \text{Bernoulli}(\pi(t)_{i,\text{treat}}),$$

where

$$\text{logit}(\pi(t)_{i,\text{treat}}) = \alpha_0 + \mathbf{X}_{i,\text{treat}}^T \alpha_1 + \mathbf{X}_{i,\text{cfd}}^T \alpha_2 + \alpha_3 A_i(t-1) + \alpha_4 Y_i(t-1); \quad (2.10)$$

and

$$Y_i(t) \sim \text{Bernoulli}(\pi(t)_{i,\text{out}}),$$

where

$$\text{logit}(\pi(t)_{i,\text{out}}) = \beta_0 + \mathbf{X}_{i,\text{out}}^T \beta_1 + \mathbf{X}_{i,\text{cfd}}^T \beta_2 + \beta_3 A_i(t) + \beta_4 Y_i(t-1). \quad (2.11)$$

We use an MSM with IPW to estimate the AMOR for the setting that involves time-varying confounders. First, assume that the key causal assumptions hold and the models are correctly specified. For nonabsorbing treatment scenario, Figure 2.3 shows that the recent outcome history $Y(t-1)$ affects both $A(t)$ and $Y(t)$, while its value depends on $A(t-1)$. Therefore, $Y(t-1)$ is considered a time-dependent confounder affected by prior exposure. The bias induced by the confounding effect can be adjusted by weighting each subject-time unit by its inverse probability weight that accumulates from the beginning to the current time point. For each subject at time s , $s = 1, \dots, T$, the corresponding (unstabilized) weight is written as:

$$\text{UW}(s) = \prod_{t=1}^s \frac{1}{P(A(t) = a(t) | \mathbf{x}_{\text{treat}}, \mathbf{x}_{\text{cfd}}, a(t-1), y(t-1))}. \quad (2.12)$$

Note that from this point on, the condition in the model is written as its realization for conciseness. For example, the condition of $(\mathbf{X}_{\text{treat}} = \mathbf{x}_{\text{treat}}, \mathbf{X}_{\text{cfd}} = \mathbf{x}_{\text{cfd}}, A(t-1) = a(t-1), Y(t-1) = y(t-1))$ is simplified to $(\mathbf{x}_{\text{treat}}, \mathbf{x}_{\text{cfd}}, a(t-1), y(t-1))$. Subscript i is suppressed in this section. Then, the conditional treatment probability in the denominator is estimated by regressing the treatment on the corresponding covariates through a generalized linear model with logistic link. Therefore, the (unstabilized) weights are calculated by the inverse of the estimated probabilities. Finally, a weighted pooled logistic regression model is fitted with the derived weights:

$$\text{logit}P(Y(t) = 1) = \gamma_0 + \gamma_{\text{AMOR}} A(t). \quad (2.13)$$

Suppose that the coefficient estimate for $A(t)$ is $\hat{\gamma}_{\text{AMOR}}$. Then, the estimated AMOR is $\exp(\hat{\gamma}_{\text{AMOR}})$.

For the scenario where the treatment is absorbing, the MSMs are again used for estimation, but the process is slightly different. Due to the absorbing nature of the model, only covariates with time up to and including the first treated visit can be used for weight estimation. Since $A(t-1)$ will always be 0 before absorption, $A(t-1)$ will not be included as a predictor in the logistic regression for weight estimation.

However, $A(t-1)$ is still necessary as an implicit condition for the realization of $A(t)$.

Empirical AMOR

Based on the empirical development in Sections 2.2.1 and 2.2.2, the value of the AMOR can be derived via MCMC for AMOR. Suppose that there are D different treatment assignment combinations during the study period $(1, \dots, T)$ for each subject. Then, the corresponding T outcomes of one subject can be simulated by (2.11) for each of its D possible treatment paths. Therefore, each subject would have $D \times T$ rows of information, and each row represents the subject information at one time point in a possible treatment path. For a sample of M subjects, $P(\bar{Y}(T) = \bar{1} | \bar{A}(T) = \bar{1})$ can be empirically represented by the average of the generating probabilities of a subject-time that is treated at the time point in a treatment path, denoted by $\hat{\pi}_{\text{AMOR},1}$. Similarly, we use $\hat{\pi}_{\text{AMOR},0}$ to represent the empirical estimate of $P(\bar{Y}(T) = \bar{1} | \bar{A}(T) = \bar{0})$.

Then, for a sufficiently large sample size (M), the AMOR can be empirically represented by

$$\text{AMOR} = \frac{\hat{\pi}_{\text{AMOR},1}/(1 - \hat{\pi}_{\text{AMOR},1})}{\hat{\pi}_{\text{AMOR},0}/(1 - \hat{\pi}_{\text{AMOR},0})}. \quad (2.14)$$

2.2.4 Variations of IPW Models

For simplicity, only unstabilized weights are mentioned during the discussion of estimation processes. We are interested in the impacts of different weight numerators on the estimation of the proposed causal effects. Note that the numerator of stabilized weights should be conditional on fewer covariates than those appear in the weight denominator, so that the weight will not always be one.

By conditioning on different sets of covariates, there could be more than one stabilized weight model available. Including different covariates in the weight numerator in the IPW model can affect the results of the causal effect estimation. In longitudinal studies where treatment varies over time, it is common to incorporate the recent treatment history into the numerator of a stabilized weight. In addition, it is also possible to include the baseline covariates that are related to the treatment model in the weight numerator. Thus, the following five weight models are proposed in the context and setting of Section 2.2.3. Note that they can also be applied to

time-fixed and partially time-varying settings accordingly. More discussion on the weight model specifications can be found in Appendix A.

Weight Model 0 is an unstabilized weight model with the numerator being one. The conditions of the denominator include all covariates that are associated with treatment: $\mathbf{X}_{\text{treat}}, \mathbf{X}_{\text{cfd}}, A(t-1)$, and $Y(t-1)$. For each of the N subjects in the time-varying setting, the unstabilized weight at time s ($s = 1, \dots, T$), is calculated using Equation (2.12).

Weight Model 1 is a semi-stabilized model with marginal probability of treatment as the numerator, which is the simplest way to treat extreme weights. For each of the N subjects, the semi-stabilized weight at time s is calculated by

$$\text{SW1}(s) = \prod_{t=1}^s \frac{P(A(t) = a(t))}{P(A(t) = a(t) | \mathbf{x}_{\text{treat}}, \mathbf{x}_{\text{cfd}}, a(t-1), y(t-1))}$$

In subsequent weight models, we intend to add different covariates to the numerator of the weight model to compare their potential effect on the AMOR estimation.

Weight Model 2 is a pure stabilized model with the numerator being the probability of treatment given the baseline covariates that are only associated with treatment. For each of N subjects, the stabilized weight at time s is calculated by

$$\text{SW2}(s) = \prod_{t=1}^s \frac{P(A(t) = a(t) | \mathbf{x}_{\text{treat}})}{P(A(t) = a(t) | \mathbf{x}_{\text{treat}}, \mathbf{x}_{\text{cfd}}, a(t-1), y(t-1))}$$

Weight Model 3 is a history stabilized model with the probability of treatment given the treatment at the previous time point as numerator. Then, the stabilized weight at time s is calculated by

$$\text{SW3}(s) = \prod_{t=1}^s \frac{P(A(t) = a(t) | a(t-1))}{P(A(t) = a(t) | \mathbf{x}_{\text{treat}}, \mathbf{x}_{\text{cfd}}, a(t-1), y(t-1))}$$

Weight Model 4 is a mixed stabilized model with baseline covariates that are associated with the treatment as numerator. Under the time-varying setting, the stabilized weight at time s is calculated by

$$\text{SW4}(s) = \prod_{t=1}^s \frac{P(A(t) = a(t) | \mathbf{x}_{\text{treat}}, a(t-1))}{P(A(t) = a(t) | \mathbf{x}_{\text{treat}}, \mathbf{x}_{\text{cfd}}, a(t-1), y(t-1))}$$

In the next section, we present and compare the simulation results of three different time settings with these weight models.

2.3 Simulations

This section covers the results of a series of Monte Carlo simulations. The five weight models proposed in Section 2.2.4 are used to estimate the corresponding causal effect measures under the three scenarios discussed in Sections 2.2.1, 2.2.2, and 2.2.3. The goal is to demonstrate the estimation methods for the proposed estimands in various scenarios.

In Section 2.2, we have described the empirical methods to determine the measure of causal effect of a known population for the three scenarios. With a sample size of $M = 100,000$, we first select the simulated populations that return the causal effect measures (MOR, MOR_{int.} and AMOR) of values 1, 2, 4, and 6 for time-fixed, partially time-varying, and time-varying settings, respectively. Then 500 data sets of $N = 10,000$ subjects are simulated for each of the time settings and weight models. The marginal treatment effects of interest are estimated using the regression models given by Equations (2.6), (2.9), and (2.13) for each time setting, respectively. For time-varying treatment and recurrent outcome, the effect of treatment is characterized by AMOR. Suppose that the estimated AMOR on the logarithmic scale (log AMOR) in the corresponding regression model for the j -th simulation is $\hat{\beta}_j$. In each of the 500 simulated data sets, the mean marginal treatment effect is recorded as the average of $500 \exp(\hat{\beta})$'s. The bias is calculated by subtracting the empirical AMOR (AMOR_{empirical} derived from Equation (2.5)) from the estimated AMOR, which is defined as $\text{bias} = \frac{1}{500} \sum_{j=1}^{500} \exp(\hat{\beta}_j) - \text{AMOR}$. The average standard error (ASE) of the estimated effect of treatment on the logarithmic scale of the 500 data sets is given by $\text{ASE} = \frac{1}{500} \sum_{j=1}^{500} \hat{\sigma}_{\beta_j}$, where $\hat{\sigma}_{\beta_j}$ is the estimated robust standard error of $\hat{\beta}_j$. The empirical standard error (ESE) is given by $\text{ESE} = \sqrt{\frac{\sum_{j=1}^{500} (\frac{1}{500} \sum_{i=1}^{500} \hat{\beta}_i - \hat{\beta}_j)^2}{500-1}}$. The results of the time-fixed and partially time-varying settings are recorded in a similar manner. Simulations were carried out using R version 3.6.0 [70]. Details about the simulation process, coefficient choices, and weight model specifications for each scenario can be found in the Appendix.

2.3.1 Results for Time-Fixed Setting

In this section, we follow Austin’s simulation design[4] based on the models described in Section 2.2.1 to examine the estimation results for the MOR. IPW models are used to adjust for confounding effects. Notice that in the time-fixed scenario, the history stabilized model and the mixed stabilized model (Weight Models 3 and 4) are identical to the semi-stabilized model and the baseline stabilized model (Weight Models 1 and 2), respectively. For simplicity, only Weight Models 0, 1, and 2 will appear in the results. The marginal treatment effect in logarithmic scale ($\log \text{MOR}$) is then estimated by GEE using IPW Weight Models 0-2. The R function “geeglm” from the “geepack” package is used to fit the weighted logistic regression given by Equation (2.6). The robust standard error of the estimated treatment coefficient is recorded accordingly. Note that it is recommended to use an independent working correlation structure to avoid bias, especially for estimations involving repeated measures [97].

We use a box plot to demonstrate the distribution of the estimation bias in the MORs for Weight Models 0-2 based on the 500 simulated data sets in Figure 2.4. The box plot shows that the mean estimation bias remains approximately zero across different weight models and MOR values. As the MOR value increases, the sampling distribution of the estimation bias becomes wider, while the range of bias distribution by Weight Model 2 is the narrowest among all weight models at each MOR levels.

The MORs, estimation bias, ASE and ESE in the 500 simulated data sets are reported in Table 2.1. Note that the ASE and ESE are standard errors for MOR on the logarithmic scale. Weight Models 0, 1, and 2 produce less biased estimations of the marginal treatment effect for different values of the MOR. We notice from Table 2.1 that Weight Models 0 and 1 have produced identical results (with the same 500 sets of simulated data). This is because the time-fixed weights estimated by Weight Models 0 and 1 only differ proportionally by treatment groups. This also implies that Weight Model 1 has only a minimal stabilization effect on the weights for the time-fixed setting. In terms of standard errors, the values of ASE and ESE of the estimated $\log \text{MOR}$ are very close. The estimation bias and the robust standard error of Weight Model 0 also match the results derived by Forbes and Shortreed [23] based Austin’s design. Although all three weight models perform well in estimating the MOR, Weight Model 2 is preferred because it produces the least variability even when the strength of MOR grows larger.




Figure 2.4: Simulation-Based Estimation for Different Weight Models (WM): Time-fixed Treatment and Time-fixed Outcome

MOR ^{<i>a</i>}	Weight Model	Bias ^{<i>b</i>}	ASE ^{<i>c</i>}	ESE ^{<i>d</i>}
1	WM0	0.0044	0.0654	0.0604
1	WM1	0.0044	0.0654	0.0604
1	WM2	0.0036	0.057	0.0513
2	WM0	0.0109	0.0652	0.0601
2	WM1	0.0109	0.0652	0.0601
2	WM2	0.007	0.0564	0.051
4	WM0	0.0012	0.0679	0.0605
4	WM1	0.0012	0.0679	0.0605
4	WM2	-0.0017	0.0584	0.0512
6	WM0	0.0295	0.0711	0.0573
6	WM1	0.0295	0.0711	0.0573
6	WM2	0.0152	0.0611	0.0501

Table 2.1: Simulation-Based Estimation of MOR with Different Weight Models: the Time-fixed Setting

^{*a*}(Empirically derived) MOR

^{*b*}Bias in the estimated MOR

^{*c*}Average standard error of log MOR

^{*d*}Empirical standard error of log MOR

2.3.2 Results for Partially Time-varying Setting

Next, the simulation is implemented with a two-course treatment and a time-fixed outcome using the setting specified in Section 2.2.2. Recall the relationship implied by Figure 2.2, the treatment and outcome statuses are generated with Monte Carlo simulations. The cumulative weights are then estimated with the five weight models for a time-varying treatment and a time-fixed outcome. Next, the marginal treatment effect on the logarithmic scale ($\log \text{MOR}_{\text{int.}}$) is estimated using the GEE method with Weight Models 0-4. The weighted logistic regression is fitted using Equation (2.9) with the R function “geeglm”. Finally, each estimated value and the corresponding robust standard error are recorded. The bias distribution in the estimated $\text{MOR}_{\text{int.}}$ ’s in the 500 simulated data sets for each weight model is presented in Figure 2.5. In general, the bias distribution has more fluctuations for larger $\text{MOR}_{\text{int.}}$ values, where the bias distributions of Weight Models 2 and 4 have narrower ranges than the other three weight models.

The $\text{MOR}_{\text{int.}}$ ’s, estimation bias, ASE, and ESE in the 500 simulated data sets are reported in Table 2.2. Note that ASE and ESE are standard errors for $\text{MOR}_{\text{int.}}$ on the logarithmic scale. Table 2.2 shows that when $\text{MOR}_{\text{int.}}$ is 1, all five weight models have less biased estimates of the marginal treatment effect. As the strength of $\text{MOR}_{\text{int.}}$ increases, Weight Models 0, 1, and 3 in the partially time-varying scenario produce more estimation bias in average, while the estimation bias produced by Weight Models 2 and 4 maintains an average around zero. In terms of estimated robust standard errors, all five weight models produce ASE and ESE that are similar to each other, while Weight Models 2 and 4 have lower standard errors than the other three weight models. Considering that Weight Models 2 and 4 are equivalent in the time-fixed setting, the conclusion that they outperform other weight models is consistent with that seen in Section 2.3.1.




Figure 2.5: Simulation-Based Estimation for Different Weight Models (WM): Partially Time-varying Setting

MOR _{2trt} ^a	Weight Model	Bias ^b	ASE ^c	ESE ^d
1	0	0.0159	0.056	0.0515
1	1	0.0039	0.0566	0.0511
1	2	0.0014	0.0502	0.0462
1	3	0.0039	0.0566	0.0511
1	4	0.0014	0.0502	0.0462
2	0	0.102	0.0588	0.055
2	1	-0.014	0.066	0.062
2	2	-0.0088	0.0534	0.0497
2	3	-0.0141	0.066	0.062
2	4	-0.0088	0.0534	0.0497
4	0	0.4508	0.0627	0.0614
4	1	0.0287	0.0797	0.0746
4	2	0.0111	0.0596	0.0576
4	3	0.0287	0.0796	0.0746
4	4	0.011	0.0596	0.0576
6	0	0.8373	0.0656	0.0586
6	1	0.1633	0.0851	0.0769
6	2	0.0616	0.0629	0.0572
6	3	0.1631	0.0851	0.077
6	4	0.0617	0.0629	0.0572

Table 2.2: Simulation-Based Estimation of Time-varying Treatment MOR_{2trt} Using Different Weight Models: the Partially Time-varying Setting

^a(Empirically derived) time-varying treatment marginal causal OR

^bBias in the estimated time-varying treatment marginal causal OR

^cAverage standard error of log MOR_{2trt}

^dEmpirical standard error of log MOR_{2trt}

2.3.3 Results for Time-varying Treatment Setting

Following the setting specified in Section 2.2.3, simulations are implemented for a time-varying treatment and a recurrent outcome. Simulations are designed to examine the estimation results when AMOR takes the values of 1, 2, 4, and 6, and when $T = 2, 4, 6$, and 10, respectively. For different time and experimental settings, we have performed separate simulations to find the unique relationship between $\exp(\hat{\beta}_3)$ and the (empirical) AMOR, which is included in the “Simulation Details for Section 3.3” of the appendix. Note that the result tables for $T = 10$ are included in Appendix B, since the trend is similar with smaller values of T . We also give a few box plots to demonstrate the distribution of the estimated weights produced by different weight models in Appendix B.

(Original) Nonabsorbing Treatment

Assume first that both $A(t)$ and $Y(t)$ are binary and nonabsorbing. Equations A.1 and A.2 in the appendix specify their generating process. The IPW models discussed in Section 2.2.4 are used to estimate marginal treatment effects. Then, the natural logarithm of the AMOR is estimated using the Weight Models 0 to 4 respectively (see Section 2.2.3 for details). Weighted logistic regression is fitted given Equation (2.13) with the R function “geeglm”, and each estimated log AMOR and the corresponding robust standard error are recorded.

For the time-varying setting with a nonabsorbing treatment, the bias distribution in the estimated AMORs in the 500 simulated data sets for each weight model is presented by a box plot in Figure 2.6.

Note that the bias scale (y-axis) in Figure 2.6 when $T = 6, 10$ is larger than when $T = 2, 4$ due to an increase in estimation bias as T increases. The values of AMOR, estimation bias, ASE, and ESE in the 500 simulated data sets when $T = 2, 4$, and 6 are reported in Table 2.3. Note that the ASE and ESE are standard errors for log AMOR. Compared to time-fixed and partially time-varying settings, the time-varying setting with nonabsorbing treatment has produced a greater fluctuation in mean estimation bias for all weight models, especially with larger T and AMOR values. The mean bias for the weight model 0 increases significantly with the growth of T and AMOR values. Weight Models 2 and 4 slightly outperform Weight Models 1 and 3, respectively. In terms of standard errors, they are correctly estimated with the five weight models ($\text{ASE} = \text{ESE}$). Overall, Weight Model 2 performs the best




Figure 2.6: Simulation-Based Estimation for Different Weight Models (WM): the original Nonabsorbing Treatment Time-varying Setting Specified in Section 2.3.3

among all weight models in terms of the mean and spread of the estimation bias, regardless of the T and AMOR values.

Absorbing Treatment

Suppose now the treatment is absorbing. Denote the treated state as state 1, then state 1 is the absorbing state for $A(t)$. We assume that all subjects are untreated at time 0 ($A(0) = 0$). The rest of the setting is the same as the nonabsorbing treatment scenario specified in Section 2.2.3. Since $A(t-1)$ will not be included as a predictor, Weight Models 3 and 4 become identical to Weight Models 1 and 2. To avoid repetition, only Weight Models 0, 1, and 2 will appear in the results.

The marginal treatment effect is then estimated using the weighted data. The rest of the estimation process is identical to the original nonabsorbing treatment scenario in Section 2.3.3. The bias in the estimated AMORs in the 500 simulated data sets for all weight models is presented in Figure 2.7.

The values of AMORs, estimation bias, ASE, and ESE in the 500 simulated data sets are reported in Table 2.4 for $T = 2, 4$, and 6. When $T = 2$ and $T = 4$, the mean bias of AMOR estimation for Weight Models 0, 1 and 2 is approximately zero, while




Figure 2.7: Simulation-Based Estimation for Different Weight Models (WM): the Absorbing Treatment Time-varying Setting Specified in Section 2.3.3

the variability in bias distribution grows as the strength of AMOR increases. As T increases further to 6 and 10, the mean of estimation bias has increased especially for larger AMOR values. One possible explanation for the increasing trend in bias is the curse of dimensionality. In other words, when T increases, the amount of data required to make a valid inference grows significantly due to the increasing complexity in covariate patterns, while we only have a fixed number of data points. In addition, the strength of AMOR has a greater effect on the estimation bias when treatment is absorbing than when it is not. The difference could also be explained by an increase in the sparsity of the data with absorbed treatment. In general, the estimation bias produced by Weight Model 2 is the closest to zero compared to other weight models. From Table 2.4, it can be observed that all three weight models produce similar ASE and ESE on the logarithmic scale, implying that the standard errors are correctly estimated. On the other hand, Weight Model 2 produces the least estimation bias, especially as T increases.

Nonabsorbing Treatment with Comparative Settings

It is difficult to compare the results of the original nonabsorbing treatment scenario and absorbing treatment cases, because they differ in settings and estimation models.

Thus, we perform an additional simulation for the nonabsorbing treatment case, so that the comparisons between the absorbing and nonabsorbing treatment cases are more straightforward. The comparative nonabsorbing treatment simulation is almost identical to the original nonabsorbing treatment simulation setting, except that now the initial treatment status, $A(0)$, will be strictly zero. This setting represents that the population was never treated prior to the study period.

Unlike the case of absorbing treatment, $A(t - 1)$ is still informative after initial time $t = 0$. Thus, even though the initial status of treatment is fixed, $A(t - 1)$ is included as a predictor in the logistic regression for weight estimation. The same IPW models for the original non-absorbing treatment scenario are used to estimate the marginal treatment effects, so the model details are skipped. Note that Table A.3 in the appendix compares all five different simulation settings specified in Sections 2.3.1, 2.3.2, and 2.3.3. Figure 2.8 shows the box plot of estimation bias in AMOR based on the 500 simulated data sets for the five weight models with different T 's. The AMOR, estimation bias, ASE, and ESE values in the 500 simulated data sets are reported in Table 2.5 for $T = 2, 4$, and 6, respectively. Note that ASE and ESE are standard errors on the logarithmic scale. Based on Figure 2.8, the comparative setting of the nonabsorbing treatment scenario behaves similarly to the original nonabsorbing treatment setting specified in Section 2.3.3. As T increases, more estimation bias is produced, especially by Weight Model 0. Weight Models 2 has the best performance among all the weight models by producing the mean estimation bias closer to zero and a relatively narrower spread of the bias distribution. Robust standard errors are correctly estimated for all weight models according to Table 2.5. With moderately valued T and AMOR, Weight Model 2 produces the least bias with a relatively smaller robust standard error.

The five weight models are compared based on time-fixed, partially time-varying, and time-varying settings with absorbing and nonabsorbing treatments, respectively. We find that the estimation results produced by Weight Model 0 is well improved in terms of the estimation bias by various weight stabilized models, namely Weight Models 1 to 4. Overall, estimation based on the pure stabilized model (Weight Model 2) produces the best results across various T , AMOR strength, and treatment settings. Compared to Weight Model 2, Weight Model 1 does not consider any treatment predictor in the weight numerator, which likely reduces the stabilization effect. Weight Models 3 and 4 generally have similar performance in the time-varying settings.




Figure 2.8: Simulation-Based Estimation for Different Weight Models (WM): the Nonabsorbing Treatment Time-varying Scenario with Comparative Settings Specified in Section 2.3.3

AMOR ¹	Weight Model	$T = 2$			$T = 4$			$T = 6$		
		Bias ²	ASE ³	ESE ⁴	Bias	ASE	ESE	Bias	ASE	ESE
1	0	0.0282	0.0399	0.0375	0.1303	0.0301	0.028	0.2438	0.0238	0.024
1	1	0.0066	0.0393	0.0362	0.012	0.0312	0.0283	0.0218	0.0275	0.0249
1	2	0.0056	0.038	0.0353	0.0076	0.0301	0.0273	0.0161	0.0267	0.0245
1	3	0.0082	0.0403	0.0374	0.0185	0.0326	0.0292	0.0307	0.0289	0.0257
1	4	0.0065	0.0388	0.0362	0.0127	0.0314	0.0282	0.0232	0.0279	0.0249
2	0	0.0098	0.0456	0.0427	0.2143	0.0343	0.0339	0.4381	0.0269	0.0264
2	1	-0.0201	0.0452	0.0434	-0.0187	0.0361	0.0354	-0.0011	0.0321	0.0315
2	2	-0.0198	0.0437	0.0412	-0.0198	0.0349	0.0339	-0.0041	0.0311	0.031
2	3	0.0015	0.0462	0.0445	0.0295	0.0376	0.0361	0.058	0.0334	0.0329
2	4	-0.001	0.0445	0.0423	0.0207	0.0362	0.0349	0.0469	0.0322	0.032
4	0	0.0035	0.0556	0.0529	0.3652	0.0418	0.0413	0.772	0.0326	0.0338
4	1	-0.0553	0.0555	0.0513	-0.0835	0.0447	0.044	-0.082	0.0394	0.0393
4	2	-0.0483	0.0535	0.0492	-0.0801	0.0432	0.0425	-0.0847	0.0382	0.0375
4	3	0.0167	0.0565	0.0528	0.0717	0.0459	0.0447	0.1083	0.0404	0.0405
4	4	0.0148	0.0543	0.0504	0.0556	0.0443	0.0432	0.0834	0.0391	0.0388
6	0	0.073	0.0644	0.061	0.5093	0.0482	0.0477	1.1637	0.0374	0.0371
6	1	-0.0111	0.0641	0.0604	-0.1613	0.052	0.051	-0.1294	0.0461	0.0442
6	2	0.0009	0.0618	0.0578	-0.1501	0.0501	0.0491	-0.1199	0.0446	0.0419
6	3	0.1234	0.0651	0.0613	0.1121	0.0531	0.0521	0.1966	0.0469	0.0448
6	4	0.1204	0.0627	0.058	0.091	0.051	0.0498	0.1684	0.0452	0.0429

Table 2.3: Simulation-Based Estimation of AMOR with Different Weight Models: the original Nonabsorbing Treatment Time-varying Setting Specified in Section 2.3.3

^b(Empirically derived) AMOR

^cBias in the estimated AMOR

^dAverage standard error of log AMOR

^eEmpirical standard error of log AMOR

AMOR ¹	Weight Model	$T = 2$			$T = 4$			$T = 6$		
		Bias ²	ASE ³	ESE ⁴	Bias	ASE	ESE	Bias	ASE	ESE
1	0	-0.0094	0.0403	0.0378	0.0073	0.0324	0.0278	0.0302	0.0285	0.0268
1	1	0	0.0394	0.0368	0.0088	0.0322	0.0272	0.0138	0.0296	0.0288
1	2	-0.0147	0.0393	0.0366	-0.0118	0.0323	0.0272	-0.0077	0.0291	0.0278
2	0	-0.0023	0.0456	0.0442	-0.0077	0.0362	0.0335	0.1198	0.0318	0.0281
2	1	0.0167	0.045	0.0434	-0.0002	0.037	0.0357	0.0904	0.0342	0.0308
2	2	-0.0026	0.0446	0.0436	-0.0369	0.0366	0.0339	0.0437	0.0332	0.0301
4	0	0.0252	0.0552	0.0523	-0.036	0.0428	0.0404	0.3521	0.0374	0.0344
4	1	0.0683	0.055	0.053	-0.0086	0.0447	0.043	0.3063	0.0414	0.0386
4	2	0.0522	0.0541	0.052	-0.074	0.0437	0.0416	0.2017	0.0398	0.0362
6	0	0.0723	0.0632	0.063	-0.0604	0.0484	0.0451	0.6216	0.0421	0.0396
6	1	0.1345	0.0634	0.0627	-0.0112	0.0509	0.048	0.5623	0.0471	0.0466
6	2	0.1272	0.0619	0.0606	-0.1145	0.0495	0.0454	0.3974	0.045	0.0438

Table 2.4: Simulation-Based Estimation of AMOR with Different Weight Models: the Absorbing Treatment Time-varying Setting Specified in Section 2.3.3

^b(Empirically derived) AMOR

^cBias in the estimated AMOR

^dAverage standard error of log AMOR

^eEmpirical standard error of log AMOR

AMOR ¹	Weight Model	$T = 2$			$T = 4$			$T = 6$		
		Bias ²	ASE ³	ESE ⁴	Bias	ASE	ESE	Bias	ASE	ESE
1	0	0.0279	0.0399	0.0391	0.1312	0.0301	0.0295	0.2439	0.0238	0.0237
1	1	0.0063	0.0393	0.0381	0.0113	0.0313	0.0304	0.0229	0.0276	0.0275
1	2	0.0046	0.038	0.0372	0.0073	0.0302	0.0296	0.0182	0.0268	0.0265
1	3	0.0189	0.0407	0.039	0.0294	0.0327	0.0322	0.0413	0.0288	0.0283
1	4	0.0158	0.0392	0.0379	0.0227	0.0315	0.0309	0.0339	0.0278	0.0272
2	0	0.0172	0.0455	0.043	0.2197	0.0344	0.0329	0.4469	0.027	0.0255
2	1	-0.0157	0.0452	0.0432	-0.0124	0.0365	0.0347	0.0047	0.0325	0.0305
2	2	-0.0177	0.0437	0.0414	-0.0158	0.0353	0.033	0	0.0315	0.0296
2	3	0.0409	0.0467	0.0444	0.0727	0.038	0.036	0.0955	0.0336	0.0307
2	4	0.0332	0.045	0.0424	0.06	0.0366	0.0343	0.0804	0.0325	0.0297
4	0	-0.0058	0.055	0.0535	0.3748	0.0418	0.0429	0.7739	0.0326	0.03
4	1	-0.0522	0.0548	0.0533	-0.0841	0.0452	0.0456	-0.1011	0.0404	0.0368
4	2	-0.0495	0.053	0.052	-0.0817	0.0437	0.044	-0.1003	0.0391	0.0358
4	3	0.1151	0.0564	0.0552	0.1699	0.0465	0.0466	0.1639	0.0411	0.0388
4	4	0.1039	0.0543	0.0536	0.146	0.0448	0.045	0.1351	0.0398	0.0375
6	0	0.0531	0.0634	0.0613	0.5501	0.0481	0.0476	1.1788	0.0375	0.0384
6	1	-0.0212	0.0638	0.0593	-0.1219	0.0523	0.0505	-0.1311	0.0469	0.047
6	2	-0.0101	0.0614	0.0567	-0.1142	0.0505	0.0492	-0.1246	0.0455	0.0452
6	3	0.2742	0.0654	0.0603	0.3029	0.0533	0.0506	0.3129	0.0471	0.0478
6	4	0.2599	0.0628	0.0577	0.2676	0.0513	0.0491	0.2756	0.0456	0.0463

Table 2.5: Simulation-Based Estimation of the AMOR with Different Weight Models: the Nonabsorbing Treatment Time-varying Scenario with Comparative Setting specified in Section 2.3.3

^b(Empirically derived) AMOR

^cBias in the estimated AMOR

^dAverage standard error of log AMOR

^eEmpirical standard error of log AMOR

2.4 Application

As an illustrative example, the AMOR estimators outlined in Sections 2.2 and 2.3 are applied to a subset of participants from the 1997 National Longitudinal Study of Youth (NLSY97) [13] to estimate the causal effect of smoking on school enrollment over five years ($T = 5$). The table in Appendix B.4 gives descriptions of the selected variables included in the analysis. Although the precise causal relationship between smoking and school enrollment for young people has not been convincingly established, previous studies have suggested a negative association between smoking behavior and educational experience for adults [89, 26, 98]. One possible causal explanation is that smoking may influence students' social behavior, leading to social isolation from school and consequently increasing the risk of school dropout. Therefore, it is of interest to investigate the possible causal effect of smoking behavior on school enrollment using AMOR estimators. There are 6,748 cross-sectional participants and 2,236 over-sampled participants in NLSY97. Participants were born between 1980 and 1984 and have been interviewed annually up to 19 times since 1997 [57].

The time-varying variables of interest are smoking status $A(t)$ and school enrollment $Y(t)$ during 1998-2002 ($t = 1, \dots, 5$) with baseline measures taken in 1997 ($t = 0$). Smoking status is defined as an indicator of self-reported smoking within the last 30 days before each interview. This is the time-varying treatment variable. School enrollment status is defined as an indicator of current school (high school or college) enrollment or previous attainment of a college degree. Since youths can leave and re-enter school, and since very few participants in NLSY97 had achieved a college degree by 2002, we treat $Y(t)$ as a recurrent event. The percentage of smokers in our cohort ranged from 37% to 43%, with an increasing trend over time. Meanwhile, the percentage of school enrollment defined by $Y(t)$ decreased from 93% to 58% during our study period. Other time-fixed baseline covariates were measured at the initial interview in 1997 ($t = 0$).

We first perform a naïve stepwise selection process on the basis of minimizing the QIC index [20] to classify the baseline variables into vectors $\mathbf{X}_{\text{treat}}$, $\mathbf{X}_{\text{cfd}}$ and $\mathbf{X}_{\text{out}}$ (based on the definitions in 2.2.1). Figure 2.9 shows the corresponding causal graph for this application. The vector of treatment-only related covariates, $\mathbf{X}_{\text{treat}}$, contains race and mother's highest degree; the outcome-only related covariate vector, $\mathbf{X}_{\text{out}}$, includes gender, time-varying age and mother's age at birth; and the vector of baseline confounders, $\mathbf{X}_{\text{cfd}}$, is composed of the indicator of marijuana use at baseline, father's highest degree, and a relative poverty level. The prior event, $Y(t - 1)$, is the




Figure 2.9: Causal graph illustrating the relationship among exposures ($A(t)$), events ($Y(t)$), and baseline covariates classified by X_{trt} , X_{out} and X_{cfd} for $t = 0, 1, \dots, 5$

only time-varying confounder. We have collapsed the levels of some covariates for simplicity and to ensure that participants reside in all combinations of treatment and confounder levels to meet the positivity assumption of causal inference. Furthermore, only participants from the cross-sectional sample without missing data over the entire study period are included in the analysis, and it is assumed that the untestable causal assumptions hold.

The proposed methods were applied to a cohort of 3,131 participants in NLSY97 using the pure stabilized model. In a sensitivity analysis, an insignificant estimated AMOR was found regardless of the grouping of baseline covariates, except if all baseline covariates are considered treatment-related. The estimated AMOR is valued at 0.93 with a 95% Confidence Interval (CI) of 0.78 - 1.10. In other words, the odds of consecutive annual enrollment if a participant actively smokes every year are 93% that of if that participant does not smoke at all, but the result is not significant. Therefore, we conclude that smoking behavior has only a potential minor negative causal impact on school enrollment. Note that we carried out the same estimation

process except using pseudo-expert information to classify $\mathbf{X}_{\text{treat}}$, $\mathbf{X}_{\text{out}}$, and $\mathbf{X}_{\text{cfd}}$. The corresponding results are consistent with our current analysis (AMOR: 0.88; 95% CI: 0.74- 1.05).

2.5 Discussion

MORs are frequently used to describe the relationship between a binary treatment and a time-fixed binary outcome. Following Austin’s design [4], the conventional MOR is consistently estimated by unstabilized IPW and time-invariant logistic regression. We extend the concept of MOR to apply to the more complex setting with a treatment that varies over time. Similar to the Monte Carlo simulation results for conventional MOR in time-fixed settings, the marginal causal treatment effect for a time-varying treatment and a time-fixed outcome ($\text{MOR}_{\text{int.}}$) can be estimated with low bias and stable variance.

Next, we proposed the AMOR as a new marginal causal treatment effect when both treatment and outcome are discrete and time-varying. Under the counterfactual framework discussed in Section 2.1, the AMOR is defined as the ratio between the odds of the outcome if all subjects are always exposed to the treatment, and the odds of the outcome if all are never exposed. In general, the AMOR can be interpreted as a measure of consecutive treatments over discrete periods of time. For instance, it can quantify the causal effect of consecutive salbutamol treatment on daily asthma attacks.

We consider two different scenarios when estimating the time-varying AMOR. In the original nonabsorbing treatment scenario, both treatment and outcome vary over time and their initial statuses in $T = 0$ are random; while in the absorbing treatment scenario, all subjects begin untreated, since we assume that once treatment begins, it is continued indefinitely. For a better comparison, a variant of the nonabsorbing treatment setting is introduced in the nonabsorbing treatment scenario with comparative setting. It specifies the initial state to be untreated and preserves all other settings from the nonabsorbing scenario. Monte Carlo simulation results suggest that the AMOR can be estimated with moderately low bias and stable variance for both absorbing and nonabsorbing treatment scenarios. Its comparative settings produce very similar results. Compared to the absorbing scenario, the estimation bias in the AMOR increases for longer study periods in the nonabsorbing treatment scenario with comparative settings. The reason could be that when T increases, the absorbing treatment setting generates data with fewer combinations of treatment

and outcome levels. In other words, the increasing bias may be due to the curse of dimensionality [28].

According to our simulation study in Section 2.3, the MSMs with appropriate IPW models estimate the proposed AMOR with low bias and reasonable standard error, when the study period is moderately short (when $T = 2, 4, 6$, and 10). Though stabilized weight models generally outperform the unstabilized weight model, some stabilized weight models have their own limitations. The semi-stabilized model does not provide enough stabilization to the weights especially when the treatment is absorbing. The history stabilized model and the mix stabilized model do not work well with longer study period or larger strength of the AMOR, as their weight numerators use the prior treatment as a predictor for the current treatment but the model is confounded. Overall, the pure stabilized weight model (Weight Model 2) works best in both absorbing and absorbing treatment scenarios among other weight models.

Some recent work has focused on time-to-event data and estimated the corresponding causal measures (e.g., hazard ratio) for treatment and outcome that vary over time [53, 2, 96]. However, for clinical data with a relatively short study period, discrete-time dichotomous treatments and recurrent outcomes are more natural and easier to work with. Given this observation, we have introduced a new causal effect measure, AMOR, specifically for time-varying treatment and recurrent outcome. Through Monte Carlo simulations, we demonstrate that AMOR can be estimated with low bias and stable variance by weighted logistic regression and robust variance estimation for study periods of moderate length and given that AMOR is of moderate strength. Meanwhile, including only the baseline covariates that is related solely to the treatment in the weight numerator of a stabilized weight model is demonstrated to work better for adjusting both time-fixed and time-varying confounders than some other types of stabilized weight models. It should also be noted that the AMOR is suitable for both absorbing and nonabsorbing treatments.

Currently, AMOR is defined to measure the causal effect of time-varying treatment on recurrent event as a whole. A possible extension in future studies is to relax the probability statements in the AMOR definition to account for more treatment paths other than fully or never treated over the entire study period. This allows AMOR to measure the causal effect of various treatment regimens. Meanwhile, since loss to follow-up or death of a participant is a common problem in longitudinal studies, it will also be of interest to incorporate missing data and competing event problems into the future study.

Chapter 3

Targeted Maximum Likelihood Estimation with Static Point Treatment and Longitudinal Follow-ups

3.1 Background

In this chapter, we investigate the average causal effect of static point treatments on recurrent events with the presence of censoring in a longitudinal setting. In the Systolic Blood Pressure Intervention Trial (SPRINT), participants were randomly assigned to one of two treatment levels: an intensive treatment group, where the target systolic blood pressure was below 120 mm Hg, and a standard treatment group, with a target of less than 140 mm Hg. This static assignment to distinct treatment regimens provides a clear framework for comparing outcomes directly attributable to the intensity of blood pressure management.

Subjects that are lost to follow-up are considered to be censored and remain so until the end of the study period. Given that censoring plays a pivotal role in determining whether the outcomes and time-varying covariates can be observed for each subject over time, it can be viewed as an implicit time-varying “exposure” [112, 107]. In reality, censoring is typically non-random and may be influenced by various environmental variables or covariate history. Thus, to account for these non-random censoring factors, we assume a dependent censoring mechanism rather than indepen-

dent censoring and incorporate time-varying confounders into our methodology. The inclusion of time-varying confounders enables us to better capture the real-life intricacies of observational data and strengthen the robustness of our method [48, 114]. Our interest lies in understanding the effect of different blood pressure lowering therapies on patient outcomes, while appropriately handling the complications introduced by censoring.

The objective of this chapter is to evaluate various estimators for estimating the ACE of point treatment processes on recurrent causal events subject to dependent censoring. Our novel contribution is a detailed exploration of the targeted maximum likelihood estimation (TMLE) in this context. We provide a comprehensive overview, discuss its theoretical foundations, and detail the implementation procedures specifically tailored to this problem setting — a perspective that has not been previously explored. We also demonstrate the benefits of the TMLE and compare it with other robust estimators that are also based on the so-called efficient influence function of our target estimand in a simulation study with different scenarios of parameter and model specifications. As an illustrative example, we apply these methods to the SPRINT dataset to estimate the ACE of blood pressure lowering therapies on the recurrences of acute kidney injury episodes in a target population. Through this implementation, we aim to draw reliable conclusions that can inform future treatment guidelines and patient care strategies.

3.2 Methodology

In this section, we introduce the observed data structure, counterfactual contrasts, identifying assumptions, as well as various estimators for estimating the ACE of static point treatments on recurrent events that are subject to dependent censoring. Additional proofs and derivations supporting our methodology are available in Appendix C.3.

3.2.1 Notation

With a static point treatment and longitudinal follow-up, the data structure and necessary notations are defined as follows.

The static point treatment, denoted by A , is binary with two possible realizations $a \in \{0, 1\}$. Let k represent the index of time intervals, $k \in \{0, 1, \dots, K + 1\}$, where

interval $K + 1$ is the end of follow-up. Note that each interval k begins at and includes the time point k , and ends just before the time point $k + 1$, that is, the interval $[k, k + 1)$. Each interval may have 0 or 1 events, and the length of the interval is fixed and defined a priori. Throughout, we let L_0 represent a vector of baseline variables measured at the start of the study period at $k = 0$ that might affect the treatment and/or an outcome of interest. Furthermore, let L_k denote a vector of time-varying covariates measured at each time interval $k \in \{1, \dots, K\}$ that may affect censoring as well as the outcome during that interval. Note that L_0 and L_k could be different (sets of) variables. The censoring indicator at interval k is denoted as C_k . If $C_k = 0$ and $C_{k+1} = 1$, then the time-varying covariates and outcome processes are only fully observed until time interval k . In this section, methodologies will be introduced in general for observational studies that assume a working model for treatment. For randomized controlled trials (RCT), on the other hand, one only requires the censoring model as the probability of treatment is fixed or known.

Suppose we can define the intervals fine enough such that the chances of experiencing more than one recurrent event during a single interval is very small. Let Y_k denote an observed binary outcome variable (i.e., $Y_k \in \{0, 1\}$) indicating whether an event occurred during time interval k . Then a corresponding observed cumulative sum is defined as R_{K+1} , where $R_{K+1} = \sum_{j=1}^{K+1} Y_k$ represents the cumulative count of recurrent events over the study period. Note that $C_0 = 0$, $Y_0 = R_0 = 0$ by definition.

We assume a discretized ordering of our observed data

$$O = (L_0, A, C_1, Y_1, L_1, \dots, C_{K+1}, Y_{K+1}).$$

throughout this chapter. For conciseness of representation, define that for any random variable B_k , $\bar{B}_k$ represents the history up to and including the interval k . i.e., $\bar{B}_k = (B_0, \dots, B_{k-1}, B_k)$; while $\underline{B}_k = (B_k, B_{k+1}, \dots, B_{K+1})$ (or $(B_k, B_{k+1}, \dots, B_K)$ if $B_{K+1} = \emptyset$) describes the future starting with the interval k . Note that for notational simplicity, we assume all variables to be discrete throughout; if there are continuous variables, then one would replace all summations with integrals accordingly.

Furthermore, the probability of treatment is abbreviated as $\pi_A^a = P(A = a)$ in trial setting and $\pi_A^a = P(A = a | L_0)$ in observational study setting where L_0 denotes potential confounders. For dependent censoring, define $\bar{\pi}_{C_k}^0 = \prod_{h=1}^k P(C_h = 0 | \bar{Y}_{h-1}, \bar{L}_{h-1}, \bar{C}_{h-1} = \bar{0}, A = a)$ as the cumulative product of the conditional probabilities of being uncensored in each time interval up to and including the interval k . In addition, we denote the sample average of the observed data of size n as $\mathbb{P}_n(\cdot)$. The definition of these notations and abbreviations are necessary for the following derivations and statements.

3.2.2 Definition and Identification of Causal Estimands

Let P denote the true population distribution of the observed data. The probabilities of treatment, i.e., $P(A = a), a \in \{0, 1\}$ in a randomized controlled trial are fixed or known a priori, but at times can depend on some baseline covariates. Our (uncensored) causal contrast is defined as

$$\Psi_I(P) = E_P \left(R_{K+1}^{a=1, \bar{c}_{K+1}=\bar{0}} - R_{K+1}^{a=0, \bar{c}_{K+1}=\bar{0}} \right),$$

which denotes the expected difference in uncensored cumulative counts of events up to a given interval K under two levels of treatment intervention $a = 1$ and $a = 0$. Recall that this type of causal effect estimand is commonly referred to as the average causal effect (ACE).

We introduce an intermediate causal estimand of $E_P(Y_k^{a, \bar{c}_k=\bar{0}})$, denoted by $\psi_{I,k}$. Based on this notation, our estimation of $\Psi_I(P)$ can be streamlined by first estimating $\psi_{I,k}$ for $k = 1, \dots, K + 1$. Subsequently, taking a summation over all $\psi_{I,k}$'s allows us to obtain the expected cumulative count of events until the end of follow-up period under a static treatment intervention $A = a$. Finally, $\Psi_I(P)$ can be estimated by taking the difference between the expected cumulative event counts under the treated ($a = 1$) and untreated ($a = 0$) conditions. To identify our causal estimand, we will make the following assumptions:

1. Exchangeability: Both treatment and censoring are conditionally unconfounded with the potential outcome. Specifically, assume:

$$Y_k^{a, \bar{c}_k=\bar{0}} \perp A | L_0, \quad (3.1)$$

$$\underline{Y}_k^{a, \bar{c}_k=\bar{0}} \perp C_k^{a, \bar{c}_{k-1}=\bar{0}} | \bar{L}_{k-1}^{a, \bar{c}_{k-1}=\bar{0}}, \bar{Y}_{k-1}^{a, \bar{c}_{k-1}=\bar{0}}, \bar{C}_{k-1}^{a, \bar{c}_{k-2}=\bar{0}}, A \quad (3.2)$$

2. Positivity 1: Each individual has a positive probability of receiving the treatments. In trial setting, $P(A = a) > 0$; in observational study setting, $P(A = a | \bar{L}_0 = \bar{l}_0) > 0$ for each stratum of $\bar{l}_0$.
3. Positivity 2: During each time interval, the probability of being uncensored is positive for all subjects uncensored till the last time interval. Specifically, let $P(\bar{C}_{k+1} = \bar{0}) > 0$ for any $C_k = 0$, $\bar{L}_k = \bar{l}_k$ and $\bar{Y}_k = \bar{y}_k$, so $P(C_{k+1} = 0 | \bar{L}_k = \bar{l}_k, \bar{Y}_k = \bar{y}_k, \bar{C}_k = \bar{0}) > 0$.
4. Consistency: The treatment and censoring variables are well-defined such that if $A = a$ and $\bar{C}_k = \bar{0}$, then $\bar{Y}_k = \bar{Y}_k^{a, \bar{c}_k=\bar{0}}$, $\bar{L}_k = \bar{L}_k^{a, \bar{c}_k=\bar{0}}$, and $\bar{R}_k = \bar{R}_k^{a, \bar{c}_k=\bar{0}}$

Figure 3.1 shows a single world intervention graph (SWIG) [72] to illustrate the relationships and identification of our causal estimand. To examine the exchangeability assumption in the SWIG framework, we ensure that, under the static intervention of $A = a, \bar{c}_k = \bar{0}$, all paths connecting A and C_k to the potential outcome $Y_k^{a, \bar{C}_k = \bar{0}}$ are blocked. In trial setting, A is independent of $Y_k^{a, \bar{C}_k = \bar{0}}$ by design; $C_k^{a, \bar{c}_{k-1} = \bar{0}}$ and $Y_k^{a, \bar{c}_k = \bar{0}}$ are d-separated by conditioning on the time-fixed time-varying confounder $\bar{L}_{k-1}^{a, \bar{c}_{k-1} = \bar{0}}$ [72]. Note that in a non-trial setting, there may be a directed causal path from L_0 to A . Then A and $Y_k^{a, \bar{C}_k = \bar{0}}$ is conditionally independent given L_0 .

In SPRINT, the positivity assumption for treatment assignment A is inherently met because each participant has been randomly assigned with either intensive or standard blood pressure lowering therapy; while the positivity for censoring process and consistency is simply assumed for this chapter. It is worthy noting that the consistency assumption might be violated if the treatment is administered differently across participants, such as variations by practitioner or changes in treatment formulation during the trial. These factors can lead to different effects from what is nominally the same treatment.

Using the aforementioned causal assumptions and induction, we can represent the identifying formula (or g-formula) for the uncensored expected outcomes with static point treatment a at time interval k as an iterative conditional expectation form:

$$\psi_{I,k} = E(Y_k^{a, \bar{c}_k = \bar{0}}) = E\{E[E(\dots E[E(Y_k | \bar{Y}_{k-1}, \bar{L}_{k-1}, a, \bar{C}_k = \bar{0}) | \bar{Y}_{k-2}, \bar{L}_{k-2}, a, \bar{C}_{k-1} = \bar{0}] \dots | Y_1, \bar{L}_1, a, C_1 = 0) | L_0]\}. \quad (3.3)$$

Recall the probability of treatment assignment is denoted by π_A^a and let the probability of a subject uncensored at time k be denoted by $\bar{\pi}_{C_k}^0$. An alternative representation of the identifying formula is an expectation inversely weighted by the probabilities of treatment assignment and censoring status, i.e.

$$\psi_{I,k} = E \left[\frac{\mathbb{I}(A = a) \prod_{h=1}^k (1 - C_h) Y_k}{\pi_A^a \bar{\pi}_{C_k}^0} \right]. \quad (3.4)$$

Derivation of the two algebraically equivalent representations of the identifying formula specified by Equations 3.3 and 3.4, respectively, is provided in Appendix C.1.2.

Finally, our final causal contrast of interest Ψ_I , can be obtained by summing up

the estimates for the $\psi_{I,k}$'s based on either 3.3 or 3.4 for time intervals $k = 1, \dots, K + 1$ and then calculating the difference between the sum of estimates under treatment ($a = 1$) minus the sum under no treatment ($a = 0$).




Figure 3.1: A single world intervention graph (SWIG) that identifies the average causal effect and illustrates the relationship between the measured baseline variables L_0 , the static treatment A , the censoring indicator $\bar{C}_{k+1}$, the time-varying variable $\bar{L}_k$ and the (intermediate) recurrent outcome $\bar{Y}_{k+1}$, assuming no unmeasured confounding. Note that for any variable X_l in the counterfactual setting, the superscript $\bar{c}_k = \bar{0}$ in $X_l^{\bar{c}_k = \bar{0}}$ is suppressed as $\bar{c} = 0$, implying X_l under history up to and include the most recent c_k before X_l ($k \leq l$). e.g., $Y_k^{\bar{c}=0} = Y_k^{\bar{c}_k = \bar{0}}$.

As the intermediate estimand $\psi_{I,k}$ can be identified from the observed data under the causal assumptions 1–4 (see also the SWIG shown in Figure 3.1), we will proceed

to discuss various estimators for estimating $\psi_{I,k}$, including those used conventionally and those with attractive asymptotic properties.

3.2.3 Inverse probability weighted (IPW) estimator: $\hat{\psi}_{\text{IPW}}$

Inverse probability weighted (IPW) estimators have been widely used in causal effect estimation, especially when the treatment and/or outcomes are confounded. For instance, IPW has been used to estimate the effect of seatbelts on heads, face and neck injuries during rollover crashes [46]; the population effect of insurance enrollment on institutional delivery and access to skilled birth attendants in Kenya [108]; and the effect of smoking on normalized adipose expression levels for different genes [71].

Recall the IPW representation of the g-formula specified by Equation 3.4, which gives rise to IPW estimators that weighs the observed outcomes by the inverse probabilities of treatment and censoring in those who followed the treatment of interest and remain uncensored. The first step of estimation is to compute the observed weights. For the treatment process in a non-randomized study, fit a model for the empirical probability of receiving treatment a , denoted as $\hat{\pi}_A^a$; and for the censoring process, fit a model for the conditional probability of being uncensored for each time point j ($j = 1, \dots, k$) or $\pi_{C_j}^0 = P(C_j = 0 | \bar{Y}_{j-1}, \bar{L}_{j-1}, \bar{C}_{j-1} = \bar{0}, A = a)$ and denote the estimate by $\hat{\pi}_{C_j}^0$.

One IPW estimator (a Horvitz-Thompson estimator) for the intermediate estimand is given as follows:

$$\hat{\psi}_{\text{IPW}} = \mathbb{P}_n \left[\frac{\mathbb{I}(A = a) \prod_{h=1}^k (1 - C_h)}{\hat{\pi}_A^a \hat{\pi}_{C_k}^0} Y_k \right]. \quad (3.5)$$

A main drawback of IPW estimators is that the estimated weights can be highly variable [81]. In particular, the weights can take extreme values if the denominators (e.g., estimated probability of treatment and being uncensored) are close to zero, which may lead to large standard errors for the estimated treatment effect. As a result, the (Horvitz-Thompson) IPW estimator does not always guarantee that the estimates will be bounded within the parameter space. This is especially problematic for our intermediate binary outcome Y_k , since the (Horvitz-Thompson) IPW estimate for $E(Y_k^{a, \bar{c}_k = \bar{0}})$ may be outside the bound of $[0, 1]$ [107]. For instance, consider a study of 2,000 participants with one individual that has an extremely low chance of treatment and a high chance of being censored. Specifically, the product of its

probability of treatment and its probability of remaining uncensored during the study period is 0.0002. This individual would then be assigned a weight of 5,000, which is significantly larger than the total number of participants in the study. Such a disproportionate weighting can skew the distribution of the weighted population, potentially causing the IPW estimator to yield values that fall outside the expected bounded range of zero to one [106].

Instead, we introduce another IPW estimator that ensures that the estimates of the intermediate causal estimand are constrained within plausible probabilistic bounds, or in other words, are bounded within the interval $[0,1]$.

3.2.4 Hájek inverse probability weighted (HIPW) estimator:

$$\hat{\psi}_{\text{HIPW}}$$

To address the limitation of the unbounded weights in the IPW estimator, we apply the Hájek estimator [30] and consider solving for ψ in the following set of estimating equations:

$$\mathbb{P}_n \left[\frac{\mathbb{I}(A=a) \prod_{h=1}^k (1-C_h)}{\hat{\pi}_A^a \hat{\pi}_{C_k}^0} (Y_k - \psi) \right] = 0,$$

where the estimated cumulative probability $\hat{\pi}_A^a \hat{\pi}_{C_k}^0$ can be derived as in Section 3.2.3. Unlike the (Horvitz-Thompson) IPW estimator, the Hájek inverse probability weighted (HIPW) estimator is guaranteed to be bounded within $[0,1]$. To see this, we note that the HIPW estimator for $\psi_{I,k}$, denoted by $\hat{\psi}_{\text{HIPW}}$, can be written as a ratio estimator:

$$\hat{\psi}_{\text{HIPW}} = \frac{\mathbb{P}_n \left[\frac{\mathbb{I}(A=a) \prod_{h=1}^k (1-C_h)}{\hat{\pi}_A^a \hat{\pi}_{C_k}^0} Y_k \right]}{\mathbb{P}_n \left[\frac{\mathbb{I}(A=a) \prod_{h=1}^k (1-C_h)}{\hat{\pi}_A^a \hat{\pi}_{C_k}^0} \right]}. \quad (3.6)$$

Since $\hat{\psi}_{\text{HIPW}}$ can be considered as a convex combination of the observed Y_k 's, it is indeed sample-bounded by $[0,1]$.

3.2.5 Iterative conditional expectation (ICE) estimator: $\hat{\psi}_{\text{ICE}}$

We now describe an iterative conditional expectation (ICE) estimator, first introduced in Section 1.2.5, within the context of our recurrent event setting. As before,

we handle dependent censoring by implicitly including in our treatment intervention a special intervention that eliminates censoring throughout follow-up. In other words, we are interested in the ACE where there is no censoring. Note that even though one would need to posit a working model for the interim outcome from Equation 3.3, the specification of the probability distribution of the covariates is not necessary for the ICE estimator. This also applies to the weighted ICE estimator, augmented inverse probability weighted estimator and the targeted maximum likelihood estimator that are extended from the ICE estimator. These estimators will be discussed in Sections 3.2.6, 3.2.7, and 3.2.8, respectively.

denote outcomes defined iteratively backwards in time. Due the complex nature of these iterative procedures, the detailed estimating process will be given in $k = 2$ for conciseness from here on. The general estimation processes will be provided in Appendix C.3. We describe an ICE estimator for $\psi_{I,2}$, denoted by $\hat{\psi}_{\text{ICE}}$ as follows:

1. Model Y_2 on past measured variables using observations with treatment level a which are uncensored until time 2.
2. Predict $\hat{Q}_{2,1}$ using the fitted model with observations with treatment level a and are uncensored until time 1.
3. Model $\hat{Q}_{2,1}$ on past measured variables using observations with treatment level a which are uncensored until time 1.
4. Predict $\hat{Q}_{2,0}$ using the fitted model with observations receiving treatment level a .
5. Calculate $\hat{\psi}_{\text{ICE}} = \mathbb{P}_n(\hat{Q}_{2,0})$ for all individuals of the study using sample average.

In the above algorithm, $\hat{Q}_{k,h}$ denotes the corresponding estimates for $Q_{k,h}$ (for $0 \leq h < k$).

Note that the IPW and ICE estimators we have introduced so far are called singly robust estimators in that they only provide one chance at consistent estimation of the (intermediate) causal estimand $\psi_{I,k}$. For instance, the ICE estimator would be biased under misspecification of the iterative working model for the conditional expected outcomes.

Next, we wish to improve estimation through the utilization of efficient influence functions (EIF). By taking the pathwise derivative along a parametric submodel of

the observed data distribution, the explicit form of the EIF for $\psi_{I,2}$, denoted by $D_{I,2}^*$, is given as:

$$D_{I,2}^* = \left\{ \frac{\mathbb{I}(A=a) \prod_{j=1}^2 (1-C_j)}{\pi_A^a \bar{\pi}_{C_2}^0} [Y_2 - Q_{2,1}] + \sum_{m=1}^1 \frac{\mathbb{I}(A=a) \prod_{h=1}^{m+1} (1-C_h)}{\pi_A^a \bar{\pi}_{C_{m+1}}^0} [Q_{2,m} - Q_{2,m-1}] + Q_{2,0} - \psi_{I,2} \right\}, \quad (3.7)$$

where $Q_{k,h}$ (for $0 \leq h < k$) has been defined above.

A detailed derivation of $D_{I,k}^*$ ($1 \leq k \leq K+1$) is provided in Appendix C.2. Based on the EIF, we are able to construct doubly robust (DR) estimators that give consistent results even when either one of the treatment (and/or censoring) or outcome models are misspecified. Note that the last component of $D_{I,2}^*$ is the ICE representation of the g-formula for our intermediate causal estimand.

3.2.6 Augmented inverse probability weighted (AIPW) estimator: $\hat{\psi}_{\text{AIPW}}$

The augmented inverse probability weighted (AIPW) estimator [77, 8], denoted by $\hat{\psi}_{\text{AIPW}}$, can be constructed by solving the estimating equation of the EIF's empirical form in our setting. As such, the AIPW estimator for our intermediate estimand can be written as:

$$\hat{\psi}_{\text{AIPW}} = \mathbb{P}_n \left\{ \frac{\mathbb{I}(A=a) \mathbb{I}(C_2=0)}{\hat{\pi}_A^a \bar{\hat{\pi}}_{C_2}^0} [Y_2 - \hat{Q}_{2,1}] + \sum_{m=1}^1 \frac{\mathbb{I}(A=a) \mathbb{I}(C_{m+1}=0)}{\hat{\pi}_A^a \bar{\hat{\pi}}_{C_{m+1}}^0} [\hat{Q}_{2,m} - \hat{Q}_{2,m-1}] + \hat{Q}_{2,0}^a \right\}, \quad (3.8)$$

where $\hat{\pi}_A^a$, $\bar{\hat{\pi}}_{C_{h+1}}^0$, and $\hat{Q}_{2,h}$ ($0 \leq h < 2$) are derived in the same way as described

in Section 3.2.3 of the IPW estimator and in Section 3.2.5, the ICE estimators, respectively, with suitable working models.

The AIPW estimator is DR in that if either the models for the treatment (censoring) or the models for the outcomes are correctly specified (not necessarily both), then ψ_{AIPW} will still be consistent.

3.2.7 Weighted iterative conditional expectation (WICE) estimator: $\hat{\psi}_{\text{WICE}}$

Based on the EIF of $\psi_{I,2}$, we can construct yet another DR estimator that is different than the AIPW estimator, referred to as the weighted iterative conditional expectation (WICE) estimator. The estimated outcomes $\hat{Q}_{k,h}$ defined in Section 3.2.5 for the ICE estimator also appear in the EIF expression given by Equation 3.7. Therefore, the WICE estimator can be regarded as a natural improvement to the ICE estimator by incorporating information about the treatment and censoring processes into the models for the iterative expected outcomes [8]. The estimation process follows the steps below.

1. Obtain $\bar{\pi}_{C_1}^0$ and $\hat{\pi}_A^a$ as in Section 3.2.3.
2. Model Y_2 on past measured variables, where each observation in the score function is weighted by inverse of the estimated cumulative probability of treatment and censoring, or $[\bar{\pi}_{C_1}^0 \hat{\pi}_A^a]^{-1}$.
3. Predict $\hat{Q}_{2,1}$ with derived model on observations with treatment level $A = a$ and are uncensored until time 1.
4. Model $\hat{Q}_{2,1}$ on past measured variables, where each observation in the score function is weighted by inverse of the estimated cumulative probability of treatment and censoring, or $[\hat{\pi}_A^a]^{-1}$.
5. Predict $\hat{Q}_{2,0}$ with derived model on observations with treatment level $A = a$.
6. Calculate $\hat{\psi}_{\text{WICE}} = \mathbb{P}_n(\hat{Q}_{2,0})$ for all individuals of the study.

3.2.8 Targeted maximum likelihood (TML) estimator: $\hat{\psi}_{\text{TML}}$

Based on EIF, the targeted maximum likelihood estimation (TMLE) is yet another solution for estimating the [101]. To the best of our knowledge, TMLE in the recurrent event setting has not been implemented before. Though closely resembling the WICE estimator and also has the DR property, a targeted maximum likelihood (TML) estimator incorporates an additional update process for the ease of incorporating machine learning algorithms. That is, we update the outcome regression models to account for the residual bias due to potential model misspecification by e.g., fitting a weighted regression to estimate a fluctuation parameter. The detailed estimation steps are given as follows.

1. Obtain $\bar{\pi}_{C_1}^0$ and $\hat{\pi}_A^a$ as in Section 3.2.3.
2. Select individuals that are uncensored until time interval 2 with treatment level $A = a$:
 - (a) Regress Y_2 on past measured variables $\bar{L}_1$ for selected individuals and obtain the initial estimate $\hat{Q}_{2,1}$ by fitting the regression model; denote the model as F_a .
 - (b) Regress Y_2 on 1 with $\hat{Q}_{2,1}$ as offset (fixed intercept) where each observation in the score function is weighed by $[\hat{\pi}_{C_1}^0 \hat{\pi}_A^a]^{-1}$; denote the new model as F_a^* .
3. Select individuals that are uncensored until time interval 1 with treatment level $A = a$:
 - (a) Obtain a new offset term with the fitted model F_a on newly selected data, denoted Q^* .
 - (b) Obtain the updated estimate $\hat{Q}_{2,1}^*$ by applying F_a^* on new offset Q^* and the selected data.
4. In the same individuals:
 - (a) Model $\hat{Q}_{2,1}^*$ on L_0 to obtain the initial estimate $\hat{Q}_{2,0}$; denote the model as F_b .
 - (b) Regress $\hat{Q}_{2,1}^*$ on 1 with $\hat{Q}_{2,0}$ as offset where each observation in the score function is weighed by $[\hat{\pi}_A^a]^{-1}$; denote the new model as F_b^* .
5. For all individuals:

- (a) Obtain a new offset term with the fitted model F_b on all individuals, denoted as Q^{**} .
 - (b) Obtain the updated estimate $\hat{Q}_{2,0}^*$ by fitting F_b^* using new offset Q^{**} .
6. Calculate $\hat{\psi}_{\text{TML}} = \mathbb{P}_n(\hat{Q}_{2,0}^*)$ for all individuals in the study.

With each of the six aforementioned estimation methods for $\psi_{I,k}$ ($k = 1, \dots, K+1$), we obtain the corresponding estimate for our causal contrast of interest. Hence Ψ_I , by plugging in summation of the estimated $\psi_{I,k}$ over all intervals of our study period under the two levels of treatment, can be derived using:

$$\Psi_I = \sum_{k=1}^{K+1} \psi_{I,k}^{a=1, \bar{c}_{K+1}=\bar{0}} - \psi_{I,k}^{a=0, \bar{c}_{K+1}=\bar{0}}.$$

3.3 Simulation

In this section, we illustrate the methodology outlined in Section 3.2 using a series of simulations. For simplicity, we assume a lag-1 variable dependence, which means that the time-varying covariates causally affect only the values of time-varying covariates that immediately follow in the next time interval, but not beyond. The mechanism for the time-varying variable generation can be summarized as follows: L_0 depends on (L_0, A, C_1, Y_1) , L_j ($j > 1$) depends on (A, L_{j-1}, Y_j, C_j) ; C_1 depends on (A, L_0) , C_j ($j > 1$) depends on $(L_{j-1}, C_{j-1}, Y_{j-1})$; and Y_1 depend on (A, L_0, C_1) and Y_j ($j > 1$) depend on $(A, L_{j-1}, C_j, Y_{j-1})$.

Consider a setting, similar to the application to be discussed in Section 3.4 with two time intervals ($K + 1 = 2$) and static treatment randomized at baseline. For a single observation, the observed data is denoted by $O = (L_0, A, C_1, Y_1, L_1, C_2, Y_2)$. Figure 3.2 illustrates an example of a causal graph for two time intervals.

Note that this setting can be generalized for any number of time intervals. Based on the theoretical results provided in Section 3.2, our intermediate estimand of interest can be written as

$$\psi_{I,2} = E\{E[E(Y_2|\bar{C}_2 = \bar{0}, \bar{L}_1, A = a)|L_0]\} + E[E(Y_1|L_0)],$$

In this study, the baseline variables L_0 consists of two components: $L_{01} \sim \text{Binomial}(0.2)$ and $L_{02} \sim \text{Binomial}(0.7)$. The binary treatment is randomly assigned to




Figure 3.2: A single world intervention graph (SWIG) that identifies the average causal effect and illustrates the relationship between the measured baseline variables L_0 , the initial treatment assignment (A), the censoring status $\bar{C}_2$, the confounders $\bar{L}_1$, and the (intermediate) recurrent outcome $\bar{Y}_2$, assuming there is no unmeasured confounding.

each of N simulated subjects with a probability of $P[A = 1] = 0.5$. Next, as we have two time intervals, we generate $(C_1, Y_1, L_1, C_2, Y_2)$ as follows:

$$\begin{aligned}
P(C_1|L_0, A = a) &= \text{expit}(\gamma_0 + \gamma_1 A + \gamma_2 L_{01} + \gamma_5 L_{02} + \gamma_6 L_{01} L_{02}) \\
P(Y_1|C_1 = 0, L_0, A = a) &= \text{expit}(\theta_0 + \theta_1 A + \theta_2 L_{01} + \theta_5 L_{02} + \theta_6 L_{01} L_{02}) \\
P(L_1|Y_1, C_1 = 0, L_0, A = a) &= \text{expit}(\beta_0 + \beta_1 A + \beta_2 L_{01} + \beta_4 Y_1 + \beta_5 L_{02} + \beta_6 L_{01} L_{02}) \\
P(C_2|Y_1, \bar{L}_1, C_1 = 0, A = a) &= \text{expit}(\gamma_0 + \gamma_1 A + \gamma_2 L_{01} + \gamma_3 L_1 + \gamma_4 Y_1 + \gamma_5 L_{02} + \gamma_6 L_{01} L_{02}) \\
P(Y_2|Y_1, \bar{L}_1, \bar{C}_2 = \bar{0}, A = a) &= \text{expit}(\theta_0 + \theta_1 A + \theta_2 L_{01} + \theta_3 L_1 + \theta_4 Y_1 + \theta_5 L_{02} + \theta_6 L_{01} L_{02})
\end{aligned}$$

We investigate the performance of the six estimators under different strengths of causal effect by varying the corresponding coefficients. For instance, γ_1 represents the conditional log odds ratio (logOR) of censoring under different levels of treatment A , i.e., treated vs. untreated, given the other variables are fixed. We will discuss the simulation results under different coefficient specifications in Section 3.3.2.

3.3.1 Model Specifications

To examine the robustness of the six estimators, we consider four scenarios which vary in the correctness of the working models. In Scenario 0 all models f are correctly specified.

Scenario 0: Correct censoring and outcome models

$$\begin{aligned}
f_{C_1} &= \text{logit } P(C_1|L_0, A = a) \\
&= \gamma_0 + \gamma_1 A + \gamma_2 L_{01} + \gamma_5 L_{02} + \gamma_6 L_{01} L_{02} \\
f_{Y_1} &= \text{logit } P(Y_1|C_1 = 0, L_0, A = a) \\
&= \theta_0 + \theta_1 A + \theta_2 L_{01} + \theta_5 L_{02} + \theta_6 L_{01} L_{02} \\
f_{L_1} &= \text{logit } P(L_1|Y_1, C_1 = 0, L_0, A = a) \\
&= \beta_0 + \beta_1 A + \beta_2 L_{01} + \beta_3 L_1 + \beta_4 Y_1 + \beta_5 L_{02} + \beta_6 L_{01} L_{02} \\
f_{C_2} &= \text{logit } P(C_2|Y_1, \bar{L}_1, C_1 = 0, A = a) \\
&= \gamma_0 + \gamma_1 A + \gamma_2 L_{01} + \gamma_3 L_1 + \gamma_4 Y_1 + \gamma_5 L_{02} + \gamma_6 L_{01} L_{02} \\
f_{Y_2} &= \text{logit } P(Y_2|Y_1, \bar{L}_1, \bar{C}_2 = \bar{0}, A = a) \\
&= \theta_0 + \theta_1 A + \theta_2 L_{01} + \theta_3 L_1 + \theta_4 Y_1 + \theta_5 L_{02} + \theta_6 L_{01} L_{02}
\end{aligned}$$

In the remaining scenarios, some (or all) of the models are misspecified through the omission of important variables and/or interaction terms. The misspecified model are denoted as f' . In the specifications given below, any models not mentioned are correctly specified and given by f above.

Scenario 1: Misspecified censoring models

$$\begin{aligned}
f'_{C_1} &= \gamma_0 + \gamma_2 L_{01} \\
f'_{C_2} &= \gamma_0 + \gamma_2 L_{01}
\end{aligned}$$

Scenario 2: Misspecified outcome models

$$\begin{aligned}
f'_{Y_1} &= \theta_0 \\
f'_{Y_2} &= \theta_0
\end{aligned}$$

Scenario 3: Misspecified censoring and outcome models

$$f'_{C_1} = \gamma_0 + \gamma_2 L_{01}$$

$$\begin{aligned}
f'_{C_2} &= \gamma_0 + \gamma_2 L_{01} \\
f'_{Y_1} &= \theta_0 \\
f'_{Y_2} &= \theta_0
\end{aligned}$$

In addition to the four model specification scenarios, we considered three treatment effect levels: null, moderate and strong. In each case, the true value of the ACE, $\Psi_I(P) = E_P(R_2^{a=1} - R_2^{a=0})$ is empirically derived using 1,000,000 simulated subjects. A table giving empirical distributions of the key variables will be presented for each treatment effect level.

For estimation within each of 12 case (4 scenarios $\times$ 3 treatment effects), we generate 1,000 different sets of 5,000 individuals. The results are presented for (mis)specification scenarios 0-3 in Section 3.3.1. The “Bias” column denotes the average of 1,000 estimated ACEs minus the empirical true ACE and the “ESE” column represents the empirical standard error of the 1,000 estimated ACEs. For each of the datasets, we take 1,000 sets of Bootstrap samples of size 5,000, perform all five estimation methods accordingly and record the bootstrap standard error calculated based on the 1,000 samples. The “Boot ASE” column denotes the average of the 1,000 bootstrap standard errors, and the “Cov. Prob.” denotes the The 95% confidence interval coverage probability (Cov Prob) based on the the estimated ACE and the bootstrap standard error within each simulated dataset.

3.3.2 Results

In this section, we will present the results of our simulation study with respect to three levels of treatment effect strengths: null, moderate, and strong.

Null Effect In order to create a setting in which the treatment has no effect on the recurrent event outcomes, we set $\gamma_1 = \beta_1 = \theta_1 = 0$. The remaining coefficient values are given below.

$$\begin{aligned}
\beta_0 &= -3; \beta_2 = 3; \beta_3 = 3; \beta_4 = 3; \beta_5 = 3; \beta_6 = 3; \\
\gamma_0 &= -1; \gamma_2 = -3; \gamma_3 = 1; \gamma_4 = -2; \gamma_5 = -2; \gamma_6 = 4; \\
\theta_0 &= -4.1; \theta_2 = 2; \theta_3 = -1; \theta_4 = 2; \theta_5 = 2; \theta_6 = 1;
\end{aligned}$$

As expected, the ‘true’ ACE is 0, confirming there is no causal treatment effect. The marginal distributions of variables Y_1, Y_2, L_1 , and $R_2 = Y_1 + Y_2$ with respect to

treatment levels are shown in Table 3.1 (without censoring). Note that there is an overall censoring rate of 12% and an event rate of 17% after censoring. Table 3.4 shows the estimated results under a null treatment effect. A discussion of the results from all three treatment levels is given at the end of this section.

	A=1	A=0	Overall
Y_1	0.17089	0.17089	0.17089
Y_2	0.16928	0.16928	0.16928
L_1	0.49411	0.49411	0.49411
R_2	0.34016	0.34016	0.34016
N=	1000000	ACE =	0.00000

Table 3.1: Average of key time-varying variables (Y_1 , Y_2 , L_1 and R_2) under null effect

Moderate Effect Next we create a simulated population setting in which the treatment has a moderate effect on the recurrent event outcomes. For completeness, the full set of coefficient values are given below:

$$\begin{aligned}\beta_0 &= -3; \beta_1 = -2; \beta_2 = 3; \beta_3 = 3; \beta_4 = 3; \beta_5 = 3; \beta_6 = 3; \\ \gamma_0 &= -2; \gamma_1 = 2; \gamma_2 = -3; \gamma_3 = 1; \gamma_4 = -2; \gamma_5 = -2; \gamma_6 = 4; \\ \theta_0 &= -3.5; \theta_1 = -1.5; \theta_2 = 2; \theta_3 = -1; \theta_4 = 2; \theta_5 = 2; \theta_6 = 1;\end{aligned}$$

The ‘true’ ACE is -0.27, which represents a moderate causal treatment effect. The marginal distributions Y_1 , Y_2 , L_1 , R_2 with respect to treatment levels are shown in Table 3.2 (without censoring). Note that after censoring, the simulated population achieves approximately an overall censoring rate of 13% and an event rate of 17%. Table 3.5 shows the estimated results under moderate treatment effect.

	A=1	A=0	Overall
Y_1	0.10137	0.23471	0.16804
Y_2	0.10229	0.23830	0.17029
L_1	0.23228	0.51622	0.37425
R_2	0.20366	0.47301	0.33833
N =	1000000	ACE =	-0.26935

Table 3.2: Average of key time-varying variables (Y_1 , Y_2 , L_1 and R_2) under moderate effect

Strong Effect Finally, we create a simulated population setting with a strong effect where the values of the coefficients are given below.

$$\begin{aligned}\beta_0 &= -3; \beta_1 = -2; \beta_2 = 3; \beta_3 = 3; \beta_4 = 3; \beta_5 = 3; \beta_6 = 3; \\ \gamma_0 &= -2; \gamma_1 = 2; \gamma_2 = -3; \gamma_3 = 1; \gamma_4 = -2; \gamma_5 = -2; \gamma_6 = 4; \\ \theta_0 &= -3; \theta_1 = -3.5; \theta_2 = 2; \theta_3 = -1; \theta_4 = 2; \theta_5 = 2; \theta_6 = 1;\end{aligned}$$

The ‘true’ ACE is -0.56, which represents a strong causal treatment effect. The marginal distributions of variables Y_1, Y_2, L_1, R_2 with respect to treatment levels are shown in Table 3.3 (without censoring). Note that after censoring, the simulated population achieves approximately an overall censoring rate of 14% and an event rate of 17%. Table 3.6 shows the estimated results under strong treatment effect.

	A=1	A=0	Overall
Y_1	0.032722	0.30130	0.16701
Y_2	0.025411	0.31535	0.17038
L_1	0.217530	0.54194	0.37974
R_2	0.058133	0.61665	0.33739
$N =$	10000000	ACE =	-0.55852

Table 3.3: Average of key time-varying variables (Y_1, Y_2, L_1 and R_2) under strong effect

Overall, the simulated estimation results reveal a generally similar pattern in the performance of each estimator type across varying degrees of model (mis)specification and under different treatment effect strengths. The Bootstrap ASE and ESE are reasonably similar, indicating a stable estimation.

When all models are correctly specified (Scenario 0), all estimators show no bias, which confirms with the theoretical results that all six estimators are consistent. However, as the level of misspecification increases, singly robust estimators like IPW and HIPW suffer in terms of bias and standard error. In Scenario 1, with a misspecified censoring model, both estimators show significant weaknesses. Similarly, in Scenario 2, the ICE estimator struggles with the largest bias and the lowest coverage probability due to an incorrectly specified outcome model.

In Scenario 3 where both the treatment and outcome models are misspecified, all estimators tend to be biased, but the doubly robust estimators still outperform their singly robust counterparts, particularly in strong treatment effect scenarios

where they exhibit performance closer to the ideal, maintaining bias near zero and coverage probabilities around 95%.

It is worth mentioning that the estimates for WICE and TMLE methods are identical in scenarios 3 and 4. This is likely because both methods collapse to solving the same estimating equations when the outcome working models are misspecified as nonparametric. With more complex model settings and misspecifications, such as non-binary variables and parametric misspecifications, we expect the WICE and TML estimators to produce different estimates.

When just one of the treatment and outcome models is correctly specified (Scenarios 1 and 2), the AIPW, WICE, and TML estimators consistently demonstrate minimal bias and maintain coverage probabilities between 94% and 95% across null, moderate and strong treatment effect cases. This double robustness ensures that they perform reliably regardless of the accuracy of the underlying models, confirming theoretical expectations. Conversely, the singly robust IPW and HIPW estimators exhibit increased bias, particularly when the treatment model is misspecified. The IPW estimator, in particular, struggles with higher estimated variances and reduced coverage probabilities, indicating it is less reliable under these conditions. The ICE estimator shows the most pronounced bias when the outcome model is incorrectly specified, affecting its overall efficacy.

Overall, the findings from these simulations across various model (mis)specification scenarios and under different levels of treatment effect strength demonstrate the superiority of AIPW, WICE, and TML estimators in handling biases and achieving closer adherence to theoretical coverage probabilities. While these three DR estimators perform similarly in simulation, one might prefer the TML estimator as it allows for incorporating machine learning algorithms in the estimation process. In the following section, we implement the TML estimator in estimating the ACE of a randomized treatment on recurrent event outcomes using the SPRINT dataset.

Scenarios		Bias(x100)	ESE(x100)	Boot ASE(x100)	Cov. Prob.
Scenario0	IPW	-0.098633	2.0322	1.9983	0.948
	HIPW	-0.096509	2.0423	2.0048	0.948
	ICE	-0.061587	1.4567	1.4186	0.947
	AIPW	-0.061348	1.4563	1.4192	0.945
	WICE	-0.061312	1.4563	1.4191	0.945
	TML	-0.061336	1.4563	1.4192	0.945
Scenario1	IPW	-0.089288	2.0822	2.0461	0.938
	HIPW	-0.101284	2.0529	2.0076	0.946
	ICE	-0.061587	1.4567	1.4157	0.943
	AIPW	-0.061587	1.4567	1.4157	0.943
	WICE	-0.061531	1.4568	1.4157	0.944
	TML	-0.061587	1.4567	1.4157	0.943
Scenario2	IPW	-0.098633	2.0322	1.9964	0.944
	HIPW	-0.096509	2.0423	2.0029	0.941
	ICE	-0.102463	2.0922	2.0517	0.948
	AIPW	-0.096448	2.0428	2.0033	0.941
	WICE	-0.096509	2.0423	2.0029	0.941
	TML	-0.096509	2.0423	2.0029	0.941
Scenario3	IPW	-0.089288	2.0822	2.0480	0.936
	HIPW	-0.101284	2.0529	2.0097	0.945
	ICE	-0.102463	2.0922	2.0521	0.944
	AIPW	-0.101579	2.0531	2.0097	0.945
	WICE	-0.101284	2.0529	2.0097	0.945
	TML	-0.101284	2.0529	2.0097	0.945

Bias(x100): 100 times the estimation bias

ESE(x100): 100 times the empirical standard error

Boot ASE(x100): 100 times the Bootstrap Average Standard Error

Cov. Prob.: Coverage Probability

IPW: Inverse Probability Weighting (estimator)

HIPW: Hájek Inverse Probability Weighting (estimator)

ICE: Iterative Conditional Expectation (estimator)

AIPW: Augmented Inverse Probability Weighting (estimator)

WICE: Weighted Iterative Conditional Expectation (estimator)

TML: Targeted Maximum Likelihood (estimator)

Table 3.4: Table of estimation results under null effect based on 1,000 sets of 5,000 simulated individuals with six estimators: IPW, HIPW, ICE, AIPW, WICE and TML respectively

Scenarios		Bias(x100)	ESE(x100)	Boot ASE(x100)	Cov. Prob.
Scenario0	IPW	-0.030865	1.9640	1.9610	0.942
	HIPW	-0.027486	1.9683	1.9670	0.940
	ICE	0.028112	1.5353	1.5126	0.947
	AIPW	0.026905	1.5375	1.5143	0.946
	WICE	0.027077	1.5375	1.5143	0.946
	TML	0.026953	1.5375	1.5143	0.946
Scenario1	IPW	-10.859699	2.0941	2.0947	0.000
	HIPW	-0.598052	1.9594	1.9667	0.936
	ICE	0.028112	1.5353	1.5120	0.949
	AIPW	0.028112	1.5353	1.5120	0.949
	WICE	0.028082	1.5353	1.5120	0.948
	TML	0.028112	1.5353	1.5120	0.949
Scenario2	IPW	-0.030865	1.9640	1.9581	0.940
	HIPW	-0.027486	1.9683	1.9640	0.940
	ICE	-0.792206	1.9851	1.9904	0.932
	AIPW	-0.027713	1.9684	1.9642	0.940
	WICE	-0.027486	1.9683	1.9640	0.940
	TML	-0.027486	1.9683	1.9640	0.940
Scenario3	IPW	-10.859699	2.0941	2.0913	0.000
	HIPW	-0.598052	1.9594	1.9648	0.935
	ICE	-0.792206	1.9851	1.9909	0.933
	AIPW	-0.478602	1.9592	1.9647	0.939
	WICE	-0.598052	1.9594	1.9648	0.935
	TML	-0.598052	1.9594	1.9648	0.935

Bias(x100): 100 times the estimation bias

ESE(x100): 100 times the empirical standard error

Boot ASE(x100): 100 times the Bootstrap Average Standard Error

Cov. Prob.: Coverage Probability

IPW: Inverse Probability Weighting (estimator)

HIPW: Hájek Inverse Probability Weighting (estimator)

ICE: Iterative Conditional Expectation (estimator)

AIPW: Augmented Inverse Probability Weighting (estimator)

WICE: Weighted Iterative Conditional Expectation (estimator)

TML: Targeted Maximum Likelihood (estimator)

Table 3.5: Table of estimation results under moderate effect based on 1,000 sets of 5,000 simulated individuals with six estimators: IPW, HIPW, ICE, AIPW, WICE and TML respectively

Scenarios		Bias(x100)	ESE(x100)	Boot ASE(x100)	Cov. Prob.
Scenario0	IPW	-0.0370788	1.7451	1.7804	0.945
	HIPW	-0.0366672	1.7487	1.7829	0.944
	ICE	-0.0067178	1.5172	1.5259	0.953
	AIPW	-0.0045085	1.5178	1.5261	0.951
	WICE	-0.0044450	1.5179	1.5262	0.951
	TML	-0.0044939	1.5178	1.5278	0.952
Scenario1	IPW	-12.3598006	2.0185	2.0425	0.000
	HIPW	-1.5787882	1.7415	1.7694	0.871
	ICE	-0.0067178	1.5172	1.5236	0.952
	AIPW	-0.0067178	1.5172	1.5236	0.952
	WICE	-0.0067285	1.5172	1.5236	0.952
	TML	-0.0067178	1.5172	1.5251	0.952
Scenario2	IPW	-0.0370788	1.7451	1.7809	0.946
	HIPW	-0.0366672	1.7487	1.7834	0.944
	ICE	-1.7194646	1.7542	1.7918	0.849
	AIPW	-0.0367220	1.7487	1.7835	0.944
	WICE	-0.0366672	1.7487	1.7834	0.944
	TML	-0.0366672	1.7487	1.7834	0.944
Scenario3	IPW	-12.3598006	2.0185	2.0438	0.000
	HIPW	-1.5787882	1.7415	1.7709	0.869
	ICE	-1.7194646	1.7542	1.7923	0.847
	AIPW	-1.5560783	1.7436	1.7703	0.869
	WICE	-1.5787883	1.7415	1.7709	0.869
	TML	-1.5787883	1.7415	1.7709	0.869

Bias(x100): 100 times the estimation bias

ESE(x100): 100 times the empirical standard error

Boot ASE(x100): 100 times the Bootstrap Average Standard Error

Cov. Prob.: Coverage Probability

IPW: Inverse Probability Weighting (estimator)

HIPW: Hájek Inverse Probability Weighting (estimator)

ICE: Iterative Conditional Expectation (estimator)

AIPW: Augmented Inverse Probability Weighting (estimator)

WICE: Weighted Iterative Conditional Expectation (estimator)

TML: Targeted Maximum Likelihood (estimator)

Table 3.6: Table of estimation results under strong effect based on 1,000 sets of 5,000 simulated individuals with six estimators: IPW, HIPW, ICE, AIPW, WICE and TML respectively

3.4 Application

3.4.1 SPRINT Data

Based on the SPRINT study introduced in Section 1.1.2, we consider the acute kidney injury episodes as our recurrent outcome of interest and study the effect of intensive treatment on the acute kidney injury recurrences. We consider an analysis to obtain an estimate of the intention-to-treat (ITT) effect. In our analysis, 9,323 subjects with complete baseline variables (treatment assignment and potential confounders) are included (0.4% omitted due to missing variables) and thereby we adopt a missing completely at random assumption for the incompletely observed baseline data. Table 3.7 summarizes preliminary data information on baseline variables with respect to treatment assignment. Note that we eliminated one out of the 9,323 subjects with complete baseline information because they had no information on follow-up visits.

Similar to Section 3.3.2, we define key variables as follows, with their corresponding notation in parentheses. The treatment has two levels: the standard blood pressure lowering therapy ($A = 0$) and the intensive blood pressure lowering therapy ($A = 1$). We define Y_k as the indicator for whether a subject has experienced an acute kidney injury episode in time interval k . As before, we assume that the probability of having more than one acute kidney injury episode within the same time interval remains negligible. Then, our outcome of interest is defined to be the total number of acute kidney injury episodes until the end of study period K ($R_K = \sum_{k=1}^K Y_k$) and $Y_k = R_k - R_{k-1}$. Subject's smoking status, history of clinical or subclinical cardiovascular disease, clinical or subclinical chronic kidney disease, statin use, and sex are considered as the baseline confounders (L_0). The subject's recent measurement of mean arterial pressure (MAP) over time k (L_k) is considered as the time-varying confounder that affects both the subject's censoring status and the recurrent outcomes. Note that the time-varying MAP is estimated by the equation $\text{MAP} = \text{DP} + 1/3(\text{SP} - \text{DP})$ where SP and DP denote systolic and diastolic blood pressure, respectively, which is a standard practice in clinical studies [55, 21].

In terms of censoring, subjects are defined to be “censored” at the time of their death or at the end of four-years follow-up in the SPRINT studies [10]. For simplicity, the first four years of the SPRINT will be considered as the study period of this secondary analysis. By the definition of censoring given by SPRINT, we will treat death as censoring in this chapter, but will formalize different estimands that can arise from complications due to death in the next chapter.

3.4.2 Effect Estimand

Recall our causal estimand of interest given in Section 3.2. Specifically, interest lies in the estimation of $\Psi_I = E_P \left(R_2^{a=1, \bar{c}_2=\bar{0}} - R_2^{a=0, \bar{c}_2=\bar{0}} \right)$, defined as the expected difference in the total number of acute kidney injury events if all the uncensored participants were treated with intensive therapy during the first four years of the SPRINT study versus that if they were treated with standard therapy.

3.4.3 Data Coarsening

For ease of implementation, our study period of 1,460 days (four years) is divided into two time intervals. Specifically, let $t = 1$ denote time interval from day 1 to day 730 and $t = 2$ denote time interval from day 731 to day 1,460.

In longitudinal studies, data can be presented in different formats. A long format data arranges each subject's repeated observations as separate rows per time point, whereas wide format data compiles all observations for each subject into a single row, showcasing different variables across time points [52]. For SPRINT data, we transform the original data from long format to wide format by discretizing the continuous-time variables Y_k and C_k into binary indicators through a threshold process. That is, for $t \in \{1, 2\}$, $Y_t = 1$ if there is acute kidney injury event incurred during time interval t for a subject, and 0 otherwise. Let $C_t = 1$ if the subject is censored during time interval t , and 0 otherwise. Table 3.8 shows that the majority of subjects encountered at most one event during each of the two time intervals. In addition, L_t is transformed to be the average of all MAP recordings during the time interval t .

Table 3.9 shows the average of the time-varying variables after data coarsening. Although some information is inevitably lost through this transformation process, the simplification enables a more straightforward analysis and interpretation while preserving the main characteristics of the data.

Characteristic	Treatment Groups	Intensive Therapy on 4661 subjects(%)	Standard Therapy on 4661 subjects (%)
Gender	female	1676(36.0)	1641(35.2)
Race	white	3039(65.2)	3033(65.0)
	black	1448(31.1)	1485(31.9)
Cardiovascular disease		939(20.1)	931(20.0)
Chronic kidney disease*		1327(28.4)	1313(28.1)
Statin use		1917(41.0)	2001(42.6)
Smoking status	never	2053(43.9)	2066(44.2)
	former	1970(42.2)	1994(42.6)
	current	638(13.7)	601(12.8)

*Chronic kidney disease was defined as an estimated glomerular filtration rate of less than 60 ml per minute per 1.73 m^2 of body-surface area. Race and ethnic group were self-reported; the race category of “black” includes Hispanic black and black as part of a multiracial identification; the race category of “other” includes Asian, Native and other ethnic groups.

Table 3.7: Baseline Characteristics for the SPRINT Study Population of 9,322 Subjects Included in the Analysis

	Intensive ($A = 1$)		Standard ($A = 0$)	
Event Count	$t = 1$	$t = 2$	$t = 1$	$t = 2$
0	4534	4564	4592	4601
1	117	86	61	51
2	8	5	6	7
3	2	4	2	2
4	0	1	0	0
5	0	1	1	1

Table 3.8: Event Count of Acute Kidney Injury Episode within Time Interval 1 and 2 for the Intensive and Standard Therapy Arms ($A = 1$ and $A = 0$ Respectively)

Time-varying Variables	Intensive Therapy	Standard Therapy	Overall
Dead at $t = 1$ (C_1)	0.017	0.020	0.019
Dead at $t = 2$ (C_2)	0.025	0.032	0.029
Event status at $t = 1$ (Y_1)	0.027	0.015	0.021
Event status at $t = 2$ (Y_2)	0.020	0.013	0.017
Mean arterial pressure at $t = 1$ (L_1)	85.142	95.321	90.232
Total events (R_2)	0.047	0.028	0.038

Table 3.9: Average of Key Time-varying Variables in the first interval ($t = 1$) or in the second interval ($t = 2$) for the SPRINT Study Population Included in the Analysis

3.4.4 Results

With the assumptions 1–4 listed in Section 3.2, our causal contrast of Ψ_I can be identified from the observed data of SPRINT, as discussed in Section 3.2. We adopt the TMLE algorithm specified in Section 3.2.8 to our specific setting to obtain an estimated ITT effect of 0.021 with a 95% confidence interval of [0.012,0.030]. The standard error for the the confidence interval construction was obtained using 1,000 nonparametric bootstrap samples. Therefore, individuals would have more acute kidney injury events had they survived the first four years of the SPRINT and been treated with intensive therapy than had they survived with standard therapy. Note that the estimated result aligns with the conclusions of other related analysis [48, 111].

3.5 Conclusion

In this chapter, we described and compared the (Horvitz-Thompson) IPW estimator, HIPW estimator, ICE estimator, WICE estimator, AIPW estimator, and TML estimator for the ACE of a static point treatment on recurrent event outcomes with censoring. The theoretical insights for these proposed estimators were validated through simulations, revealing that the TML, WICE, and AIPW estimators perform the best due to their double robustness. Among them, the TML estimator has the additional advantage of incorporating machine learning mechanisms in the update step. This approach requires fewer data assumptions and is more suitable for complex data structures. Therefore, we implemented the TML method on data from the SPRINT study. Through an ITT analysis, we concluded that the intensive blood pressure lowering therapy on uncensored participants caused slightly more acute kidney injury recurrences than the standard therapy in the first four years of the SPRINT study. This finding prompts further investigation into the trade-offs between blood pressure management and kidney health, which will be explored in depth in the subsequent chapter.

Chapter 4

Targeted Maximum Likelihood Estimation with Time-varying Treatments

In this chapter, we extend the estimation methods used in Chapter 3 to handle time-varying treatments. Proofs and derivations supporting our methodology are available in Appendix D.

4.1 Background

Building on the discussion from Chapter 3 on static point treatments, this chapter delves deeper into the setting of time-varying treatment strategies. For instance, in the SPRINT study, subjects are randomly assigned to target a systolic blood pressure of less than 120 mm Hg or 140 mm Hg. Doctors prescribe blood pressure lowering medications to their patients to achieve these targets, but patients' adherence to these prescribed medications may vary over time and can depend on various factors. Therefore, the time-varying adherence level to the assigned treatment can also affect the recurrence of participants' acute kidney injuries and should be considered as part of the treatment process.

In this chapter, we aim to utilize targeted maximum likelihood estimation (TMLE) to assess the average causal effect (ACE) of static, time-varying treatments strategies on recurrent event outcomes, with simulation to confirm its theoretical properties

along with other comparative estimation methods. In addition, we will apply our methods to the SPRINT dataset, where we address various estimands that arise in a semi-competing risk setting.

4.2 Methodology

4.2.1 Notation

Based on notation introduced in Section 3.2.1, we define additional notation necessary for our setting as follows.

We introduce a time-varying treatment or exposure variable A_k^* for time interval k ($k \in \{0, \dots, K\}$). For notational simplicity, we let the vector A_k denote an individual's joint exposure status during time interval k that is composed of two binary variables that indicate treatment status A_k^* in interval k and censoring status C_{k+1} in interval $k+1$, i.e., $A_k = (A_k^*, C_{k+1})$.

Similarly, we let the vector L_{k+1} denote an individual's joint time-varying covariates during time interval $k+1$ that is composed of a binary (intermediate) outcome variable Y_{k+1} and a vector of covariates L_{k+1}^* in interval $k+1$, i.e., $L_{k+1} = (Y_{k+1}, L_{k+1}^*)$. For completeness, define $L_0 = L_0^*$ and $L_{K+1} = Y_{K+1}$. Therefore, the observed data for the time-varying treatment and recurrent outcome processes is denoted by:

$$O = (L_0, A_0^*, C_1, Y_1, L_1^*, A_1^*, \dots, C_K, Y_K, L_K^*, A_K^*, C_{K+1}, Y_{K+1}),$$

or more simply as:

$$O = (L_0, A_0, L_1, A_1, \dots, L_K, A_K, Y_{K+1}).$$

Suppose the true population distribution of the observed data is denoted by P . Our primary causal contrast of interest, $\Psi_{II}(P)$, is defined as $E_P(R_{K+1}^{\bar{a}_K^\dagger} - R_{K+1}^{\bar{a}_K^{\dagger\dagger}})$. This represents the expected difference in cumulative counts of events up to interval $K+1$ under two distinct treatment intervention paths, $\bar{a}^\dagger$ and $\bar{a}^{\dagger\dagger}$. Both intervention paths include the component where censoring is completely eliminated, specifically, $\bar{C}_{K+1} = \bar{0}$. Then, we define the intermediate estimand as the expected count of event in interval k with a treatment path $\bar{a}_k^*$ where censoring is abolished, denoted

by $\psi_{II,k}(P) = E_P(Y_k^{\bar{a}_{k-1}})$. To identify our causal estimand, we will make the following assumptions:

1. Exchangeability: The joint treatment is conditionally unconfounded with the next potential outcome. Specifically, assume:

$$Y_k^{\bar{a}_{k-1}} \perp A_{k-1}^{\bar{a}_{k-2}} | \bar{L}_{k-1}, \bar{A}_{k-2}^{\bar{a}_{k-3}}.$$

2. Positivity 1: Each subject has a positive probability of receiving the treatments at each time interval of study, given past covariates. That is, for any possible paths of $\bar{a}_k^*$ and all observed $\bar{l}_k$, $P(\bar{A}_k^* = \bar{a}_k^* | \bar{L}_k = \bar{l}_k, \bar{A}_{k-1}^* = \bar{a}_{k-1}^*) > 0$.
3. Positivity 2: During each time interval, the probability of being uncensored is positive for all uncensored subjects. Specifically, let $P(C_{k+1} = 0 | \bar{L}_k = \bar{l}_k, \bar{C}_k = \bar{0}, \bar{A}_k^* = \bar{a}_k^*) > 0$ for any possible paths of $\bar{a}_k^*$ and all observed $\bar{l}_k$.
4. Consistency: If $\bar{A}_{k-1} = \bar{a}_{k-1}$, then $\bar{Y}_k = \bar{Y}_k^{\bar{a}_{k-1}}$, $\bar{L}_k = \bar{L}_k^{\bar{a}_{k-1}}$.

Define the joint probabilities of the time-varying treatment of uncensored subjects in each time interval until Y_k as $\bar{\pi}_{A_{k-1}}^{a_{k-1}}$, where $\bar{\pi}_{A_{k-1}}^{a_{k-1}} = \prod_{j=0}^{k-1} P(A_j = a_j | \bar{L}_{j-1}, \bar{A}_{j-1} = \bar{a}_{j-1})$. Then, the intermediate causal estimand can be represented in the inverse probability weighted (IPW) form given by:

$$\psi_{II,k} = E \left(\frac{\prod_{h=1}^{k-1} \mathbb{I}(A_h = a_h)}{\bar{\pi}_{A_{k-1}}^{a_{k-1}}} Y_k \right); \quad (4.1)$$

or the iteratively conditional expectation (ICE) form given by:

$$\psi_{II,k}(P) = E\{E[E(\dots E[E(Y_k | \bar{L}_{k-1}, \bar{a}_{k-1}) | \bar{L}_{k-2}, \bar{a}_{k-2}] \dots | \bar{L}_1, \bar{a}_1) | L_0]\}. \quad (4.2)$$

In Appendix D.2, we provide the proof of identification for these two representations based on the causal assumptions of interest. Figure 4.1 provides a graphical approach to assess the exchangeability conditions that are used to identify $\psi_{II,k}(P)$, as explained in Section 3.2.2 of the previous chapter. As the intermediate estimand $\psi_{II,k}$ can be identified from the observed data under the causal assumptions 1–4, we proceed to describe a suite of estimators to estimate $\psi_{II,k}$ for $k = 1, \dots, K+1$. These estimators are similar to those specified in Section 3.2, and the general estimation procedures for some of these estimators will be provided in Appendix D.3. Note that we will alternate between the composite notations and the full notations (e.g. A_k vs

(A_k^*, C_{k+1}) , respectively) to ensure a balance in conciseness and clarity in explaining the estimation procedures.




Figure 4.1: A single world intervention graph (SWIG) that illustrates the connections and dependencies among the key variables: L_0 (the baseline covariates), A_{k-1} (the time-varying treatment and the censoring indicator in interval k), L_{k-1} (the time-varying covariate and the intermediate recurrent outcome in interval k), and Y_{K+1} (the last intermediate outcome), where $k \in \{1, \dots, K+1\}$

4.2.2 Inverse probability weighted (IPW) estimator: $\hat{\psi}_{\text{IPW}}$

The IPW representation of the g-formula specified by Equation 4.1 motivates the Horvitz-Thompson IPW estimator (simplified as the IPW estimator for conciseness) defined by Equation (4.3) below. The first step of this IPW estimator is to model the treatment and censoring processes. For each time interval $j \in \{0, 1, \dots, k-1\}$, the composite treatment process can be written as

$$\begin{aligned} \pi_{A_j}^{a_j} &= P(A_j^* = a_j^*, C_{j+1} = 0 | \bar{L}_{j-1}, \bar{a}_{j-1}) \\ &= P(C_{j+1} = 0 | a_j^*, \bar{L}_{j-1}, \bar{a}_{j-1}) P(A_j^* = a_j^* | \bar{L}_{j-1}, \bar{a}_{j-1}). \end{aligned}$$

Therefore, the censoring process $P(C_{j+1} = 0 | a_j^*, \bar{L}_{j-1}, \bar{a}_{j-1})$ and the treatment process $P(A_j^* = a_j^* | \bar{L}_{j-1}, \bar{a}_{j-1})$ can be estimated separately, and the estimate of the joint probability of treatment in interval j and being uncensored in interval $j+1$, denoted as $\bar{\pi}_{A_{k-1}}^{a_{k-1}}$, can be derived by multiplying the estimated probabilities of treatment and

being uncensored for $j \in \{0, 1, \dots, k-1\}$. The Horvitz-Thompson IPW estimator, later simplified as the IPW estimator, is given by:

$$\hat{\psi}_{\text{IPW}} = \mathbb{P}_n \left[\frac{\prod_{j=0}^{k-1} \mathbb{I}(A_j^* = a_j^*, C_{j+1} = 0)}{\hat{\pi}_{A_{k-1}}^{a_{k-1}}} Y_k \right], \quad (4.3)$$

where $\mathbb{P}_n(\cdot)$ denotes the sample average over the observed data. As described in Section 3.2.3, the IPW estimator is singly robust, and the estimator may be unstable if some of the estimated probabilities of treatment and/or being uncensored are very small.

4.2.3 Hájek inverse probability weighted (HIPW) estimator:

$$\hat{\psi}_{\text{HIPW}}$$

Alternatively, we can also consider a Hájek IPW estimator for estimating the expected intermediate outcome given by Equation 4.1 [30] (referred to as the HIPW estimator). We solve for the HIPW estimator ψ_{HIPW} in the following set of equations:

$$\mathbb{P}_n \left[\frac{\prod_{j=0}^{k-1} \mathbb{I}(A_j^* = a_j^*, C_{j+1} = 0)}{\hat{\pi}_{A_{k-1}}^{a_{k-1}}} (Y_k - \psi_{\text{HIPW}}) \right] = 0,$$

where $\hat{\pi}_{A_{k-1}}^{a_{k-1}}$ is derived as in Section 4.2.2. Thus, the HIPW estimator for $\psi_{II,k}$, denoted by $\hat{\psi}_{\text{HIPW}}$, can also be written as an ratio estimator that will always be bounded between $[0,1]$ for binary outcomes, analogous to the argument given in Section 3.2.4.

4.2.4 Iterative conditional expectation (ICE) estimator: $\hat{\psi}_{\text{ICE}}$

Similarly as in Section 3.2.5, an iterative conditional expectation (ICE) estimator can be formulated as follows. Based on the g-formula represented by Equation 4.2, the ICE estimator iteratively fits a series of outcome models conditional on the corresponding variable of uncensored individuals. Define $Q_{k,k} = \hat{Q}_{k,k} = Y_k$, and iteratively for $h \in \{0, 1, \dots, k-1\}$, let

$$Q_{k,h} = E(Q_{k,h+1} | \bar{a}_h^*, \bar{C}_{h+1} = \bar{0}, \bar{L}_h)$$

denote outcomes defined iteratively backwards in time, and let $\hat{Q}_{k,h}$ denote its corresponding estimates. Furthermore, let $\hat{Q}_{k,j}$ denote a corresponding estimate for $Q_{k,j}$. Due the complex nature of these iterative procedures, the detailed estimating process will be given in $k = 2$ for conciseness from here on. The general estimation processes will be provided in Appendix D.3. We now describe an ICE estimator for $\psi_{II,2}$, denoted by $\hat{\psi}_{\text{ICE}}$, as follows:

1. Model Y_2 on past measured variables $\bar{L}_1$ using observations with treatment path $\bar{a}_1$ that are uncensored until time 2, denote the model as F_a .
2. Predict $\hat{Q}_{2,1}$ with model F_a using observations with treatment path $\bar{a}_0$ that are uncensored until time 1.
3. Model $\hat{Q}_{2,1}$ on baseline variables $\bar{L}_0$ using observations with treatment path $\bar{a}_0$ that are uncensored until time 1, denote the model as F_b .
4. Predict $\hat{Q}_{2,0}$ with model F_b using all observations.
5. Calculate $\hat{\psi}_{\text{ICE}} = \mathbb{P}_n(\hat{Q}_{2,0})$ for all observations using a sample average.

Building on our discussion of singly robust estimators for point treatments in Section 3.2.5, the ICE estimator for time-varying treatments also falls into this category. It is singly robust because any misspecification of the working models for the iterated conditional expected outcomes would result in inconsistent estimation of the ACE. Similarly, the IPW and HIPW estimators are singly robust in that they produce consistent estimation of the ACE only if the working models for the treatment and censoring processes are correctly specified.

To further improve estimation, we derive the efficient influence function (EIF) for $\psi_{II,2}$, denoted by $D_{II,2}^*$, which is given as:

$$\begin{aligned}
D_{II,2}^* = & \left\{ \frac{\mathbb{I}(\bar{A}_1 = \bar{a}_1)}{\bar{\pi}_{A_1}^{a_1}} [Y_2 - Q_{2,1}] \right. \\
& + \frac{\mathbb{I}(A_0^* = a_0^*, C_1 = 0)}{\bar{\pi}_{A_0}^{a_0}} [Q_{2,1} - Q_{2,0}] \\
& \left. + Q_{2,0} - \psi_{II,2} \right\}, \tag{4.4}
\end{aligned}$$

where $Q_{2,2} = Y_2$, and for $h = 0, 1$, $Q_{2,h}$ denotes the conditional outcome defined by $E(Q_{2,h+1}|\bar{a}_h^*, \bar{C}_{h+1}, \bar{L}_h)$. Note that the detailed derivation process of $D_{II,k}^*$ for an arbitrary k is given in Appendix D.2.

4.2.5 Augmented inverse probability weighted (AIPW) estimator: $\hat{\psi}_{\text{AIPW}}$

The Augmented IPW (AIPW) estimator is the simplest DR estimator in our recurrent event setting, and can be constructed by solving the empirical form of the EIF given as follows:

$$\begin{aligned} \hat{\psi}_{\text{AIPW}} = \mathbb{P}_n \Big\{ & \frac{\mathbb{I}(\bar{A}_1^* = \bar{a}_1^*, \bar{C}_2 = \bar{0})}{\bar{\pi}_{A_1}^{a_1}} [Y_2 - \hat{Q}_{2,1}] \\ & + \frac{\mathbb{I}(A_0^* = a_0^*, C_1 = 0)}{\bar{\pi}_{A_0}^{a_0}} [\hat{Q}_{2,1} - \hat{Q}_{2,0}] \\ & + \hat{Q}_{2,0} \Big\}, \end{aligned} \quad (4.5)$$

where $\bar{\pi}_{A_h}^{a_h}$ and $\hat{Q}_{k,h}$ ($h = 0, 1$) are derived in the same way as described in Section 4.2.2 and 4.2.4, respectively, with suitable working models.

4.2.6 Weighted iterative conditional expectation (WICE) estimator: $\hat{\psi}_{\text{WICE}}$

Another estimator based on the EIF is the weighted ICE (WICE) estimator, which is constructed in a similar way as the ICE estimator:

1. Obtain $\bar{\pi}_{A_h}^{a_h}$ ($h \in \{0, 1\}$) as described in Section 4.2.2.
2. Model Y_2 on past measured variables $\bar{L}_1$ using observations with treatment path $\bar{a}_1$ that are uncensored until time 2, while each observation is weighted by $[\bar{\pi}_{A_1}^{a_1}]^{-1}$ in the score function, denote the model as F_a .
3. Predict $\hat{Q}_{2,1}$ with model F_a that are uncensored until time 1 with treatment $A_0^* = a_0^*$.

4. Model $\hat{Q}_{2,1}$ on baseline variables L_0 using observations with treatment path $\bar{a}_0$ that are uncensored until time 1, while each observation in the score function is weighted by $[\hat{\pi}_{A_0}^{a_0}]^{-1}$; denote the model as F_b .
5. Predict $\hat{Q}_{2,0}$ with model F_b using all observations.
6. Calculate $\hat{\psi}_{\text{ICE}} = \mathbb{P}_n(\hat{Q}_{2,0})$ for all observations using a sample average.

4.2.7 Targeted maximum likelihood (TML) estimator: $\hat{\psi}_{\text{TML}}$

The targeted maximum likelihood (TML) estimator is also constructed based on the EIF and has the property of doubly robustness. With an additional update step, it allows for the incorporation of machine learning algorithms. For demonstration purpose, the estimation process with the simplest update step is shown as follows.

1. Obtain $\bar{\pi}_{A_h}^{a_h} (h \in \{0, 1\})$ as described in Section 4.2.2.
2. Select individuals that are uncensored until time interval 2 with treatment path $\bar{a}_1$:
 - (a) Regress Y_2 on past measured variables $\bar{L}_1$ using the selected data, denote the model as F_a .
 - (b) Obtain (the initial estimate of) $\hat{Q}_{2,1}$ by fitting F_a .
 - (c) Regress Y_2 on 1 with an offset (fixed intercept) of the initial estimate logit $\hat{Q}_{2,1}$, where each observation in the score function is weighted by $[\bar{\pi}_{A_1}^{a_1}]^{-1}$; denote the new model as F_a^* .
3. Select individuals that are uncensored until time interval 1 with treatment a_0 :
 - (a) Obtain the fitted values, denoted as Q^* , using model F_a with the newly selected data.
 - (b) Obtain the updated estimate $\hat{Q}_{2,1}^*$ by fitting F_a^* using the new offset logit Q^* with the selected data.
4. In the same individuals:
 - (a) Regress $\hat{Q}_{2,1}^*$ on L_0 , denote the model as F_b .
 - (b) Obtain the fitted values with model F_b , denoted as $\hat{Q}_{2,0}$.

- (c) Regress $\hat{Q}_{2,1}^*$ on 1 with the offset of logit $\hat{Q}_{2,0}$ where each observation in the score function is weighted by $[\hat{\pi}_{A_0}^{a_0}]^{-1}$; denote the new model as F_b^* .
- 5. For all individuals (unstratified),
 - (a) Obtain a new offset term using model F_b , denoted as Q^{**} .
 - (b) Obtain the updated estimate $\hat{Q}_{2,0}^*$ by fitting F_b^* using new offset logit Q^{**} .
- 6. Calculate $\hat{\psi}_{\text{TML}} = \mathbb{P}_n(\hat{Q}_{2,0}^*)$ for all observations in the study.

After estimating $\psi_{II,k}$ with each of the six estimators, for $k = 1, \dots, K + 1$, we can obtain their corresponding estimates for our causal contrast of interest, Ψ_{II} , by summing them up under the (pre-defined) treatment paths of interest, $\bar{a}_k^{\dagger\dagger}$ and $\bar{a}_k^{\dagger\dagger}$, respectively, as

$$\Psi_{II} = \sum_{k=1}^{K+1} \psi_{II,k}^{\bar{a}_k^{\dagger}} - \psi_{II,k}^{\bar{a}_k^{\dagger\dagger}}.$$

4.3 Simulation

Based on the methodology outlined in Section 4.2, we implement a series of simulations to illustrate some theoretical properties of the estimators. Assuming a lag-1 variable dependence, then the time-varying variables have the following generation mechanism. Censoring C_j depends on $(A_{j-1}^*, L_{j-1}^*, Y_{j-1}, C_{j-1})$; outcome Y_j depends on $(A_{j-1}^*, L_{j-1}^*, Y_{j-1}, C_j)$; covariate L_j^* depends on $(A_{j-1}^*, L_{j-1}^*, Y_j, C_j)$; and treatment A_j^* depends on $(A_{j-1}^*, L_j^*, Y_j, C_j)$.

Consider a setting, similar to the application to be discussed in Section 4.4 with two time intervals. We aim to estimate the effect difference between two treatment paths with uncensored subjects: $E_P(R_2^{\bar{a}_1^{\dagger}} - R_2^{\bar{a}_1^{\dagger\dagger}})$, where $\bar{a}_1^{\dagger} = ((a_0^* = 1, \bar{c}_1 = \bar{0}), (a_1^* = 1, \bar{c}_2 = \bar{0}))$ and $\bar{a}_1^{\dagger\dagger} = ((a_0^* = 0, \bar{c}_1 = \bar{0}), (a_1^* = 0,))$. Note that this setting can be generalized for any $K \geq 2$ and any treatment paths. Followed from the theoretical results provided in Section 4.2, our intermediate estimand of interest can be written as

$$\Psi_{II} = E\{E[E(Y_2|\bar{A}_1 = \bar{a}_1, \bar{L}_1)|L_0]\} + E[E(Y_1|L_0)].$$

For each individual, the observed data is denoted by $O = (L_0, A_0, L_1, A_1, Y_2)$. Figure 4.2 illustrates an example of a causal graph for two time intervals, but it can be generalized to longer study periods ($K \geq 2$).




Figure 4.2: A single world intervention graph (SWIG) that illustrates the relationship between the key variables and identifies the causal estimand of $Y_2^{\bar{a}_1}$ with $k = 2$ time intervals

Similar to the setting in Section 3.3.2, the baseline variables L_0 consists of two components: $L_{01} \sim \text{Binomial}(0.2)$ and $L_{02} \sim \text{Binomial}(0.7)$. Next, we generate $(A_0^*, C_1, Y_1, L_1^*, A_1^*, C_2, Y_2)$ as follows:

$$\begin{aligned}
P(A_0^*|L_0) &= \text{expit}(\alpha_0 + \alpha_2 L_{01} + \alpha_5 L_{02} + \alpha_6 L_{01} L_{02}) \\
P(C_1|L_0, a_0^*) &= \text{expit}(\gamma_0 + \gamma_1 A_0^* + \gamma_2 L_{01} + \gamma_5 L_{02} + \gamma_6 L_{01} L_{02}) \\
P(Y_1|C_1 = 0, L_0, a_0^*) &= \text{expit}(\theta_0 + \theta_1 A_0^* + \theta_2 L_{01} + \theta_5 L_{02} + \theta_6 L_{01} L_{02}) \\
P(L_1^*|Y_1, C_1 = 0, L_0, a_0^*) &= \text{expit}(\beta_0 + \beta_1 A_0^* + \beta_2 L_{01} + \beta_4 Y_1 + \beta_5 L_{02} + \beta_6 L_{01} L_{02}) \\
P(A_1^*|\bar{L}_1, \bar{A}_0 = \bar{a}_0) &= \text{expit}(\alpha_0 + \alpha_1 A_0^* + \alpha_2 L_{01} + \alpha_3 L_1^* + \alpha_4 Y_1 + \alpha_5 L_{02} + \alpha_6 L_{01} L_{02}) \\
P(C_2|\bar{L}_1, C_1 = 0, \bar{a}_1^*) &= \text{expit}(\gamma_0 + \gamma_1 A_1^* + \gamma_2 L_{01} + \gamma_3 L_1^* + \gamma_4 Y_1 + \gamma_5 L_{02} + \gamma_6 L_{01} L_{02}) \\
P(Y_2|\bar{L}_1, \bar{a}_1^*) &= \text{expit}(\theta_0 + \theta_1 A_1^* + \theta_2 L_{01} + \theta_3 L_1^* + \theta_4 Y_1 + \theta_5 L_{02} + \theta_6 L_{01} L_{02}).
\end{aligned}$$

We investigate the performance of the six estimators under different strengths of causal effect by varying the corresponding coefficients. For each estimated scenario, we generate 2,000 different sets of 2,000 individuals. In the results, the “Bias” column denotes the average of 2,000 estimated ACEs minus the empirical true ACE and the “ESE” column represents the empirical standard error of the 2,000 estimated ACEs. For each of the 2,000 datasets, we take 2,000 sets of Bootstrap samples of size 2,000, perform all five estimation methods accordingly and record the bootstrap standard error calculated based on the 2,000 samples. The “Boot ASE” column denotes the average of the 2,000 bootstrap standard errors, and the “Cov. Prob.” denotes the The 95% confidence interval coverage probability (Cov Prob) based on the the estimated ACE and the bootstrap standard error within each simulated dataset.

In the next section, we discuss the various working models that are used to assess robustness of different estimators discussed herein.

4.3.1 Model Specifications

To examine the robustness of the six estimators, we consider four scenarios which vary in the correctness of the working models. In Scenario 0, all models f are correctly specified.

Scenario 0: Correct censoring and outcome models

$$\begin{aligned}
f_{A_0^*} &= \text{logitP}(A_0^*|L_0) \\
&= \alpha_0 + \alpha_2 L_{01} + \alpha_5 L_{02} + \alpha_6 L_{01} L_{02} \\
f_{C_1} &= \text{logit P}(C_1|L_0, a_0^*) \\
&= \gamma_0 + \gamma_1 A_0^* + \gamma_2 L_{01} + \gamma_5 L_{02} + \gamma_6 L_{01} L_{02} \\
f_{Y_1} &= \text{logit P}(Y_1|C_1 = 0, L_0, A_0^* = a_0^*) \\
&= \theta_0 + \theta_1 A_0^* + \theta_2 L_{01} + \theta_5 L_{02} + \theta_6 L_{01} L_{02} \\
f_{L_1^*} &= \text{logit P}(L_1^*|Y_1, C_1 = 0, L_0, A_0^* = a_0^*) \\
&= \beta_0 + \beta_1 A_0^* + \beta_2 L_{01} + \beta_3 L_1^* + \beta_4 Y_1 + \beta_5 L_{02} + \beta_6 L_{01} L_{02} \\
f_{A_1^*} &= \text{logitP}(A_1^*|\bar{L}_1, \bar{A}_0 = \bar{a}_0) \\
&= \alpha_0 + \alpha_1 A_0^* + \alpha_2 L_{01} + \alpha_3 L_1^* + \alpha_4 Y_1 + \alpha_5 L_{02} + \alpha_6 L_{01} L_{02} \\
f_{C_2} &= \text{logit P}(C_2|\bar{L}_1, A_1^* = a_1^*, \bar{A}_0 = \bar{a}_0) \\
&= \gamma_0 + \gamma_1 A_1^* + \gamma_2 L_{01} + \gamma_3 L_1^* + \gamma_4 Y_1 + \gamma_5 L_{02} + \gamma_6 L_{01} L_{02} \\
f_{Y_2} &= \text{logit P}(Y_2|\bar{L}_1, \bar{A}_1 = \bar{a}_1) \\
&= \theta_0 + \theta_1 A_1^* + \theta_2 L_{01} + \theta_3 L_1^* + \theta_4 Y_1 + \theta_5 L_{02} + \theta_6 L_{01} L_{02}
\end{aligned}$$

In the remaining scenarios, some or all of the models f' are misspecified through the omission of important variables. In the specifications given below, any models not mentioned are correctly specified (i.e., given by f above).

Scenario 1: Misspecified censoring models

$$\begin{aligned}
f'_{A_0^*} &= \alpha_0 + \alpha_2 L_{01} \\
f'_{C_1} &= \gamma_0 + \gamma_2 L_{01} \\
f'_{A_1^*} &= \alpha_0 + \alpha_2 L_{01}
\end{aligned}$$

$$f'_{C_2} = \gamma_0 + \gamma_2 L_{01}$$

Scenario 2: Misspecified outcome models

$$\begin{aligned} f'_{Y_1} &= \theta_0 \\ f'_{Y_2} &= \theta_0 \end{aligned}$$

Scenario 3: Misspecified censoring and outcome models

$$\begin{aligned} f'_{A_0^*} &= \alpha_0 + \alpha_2 L_{01} \\ f'_{C_1} &= \gamma_0 + \gamma_2 L_{01} \\ f'_{A_1^*} &= \alpha_0 + \alpha_2 L_{01} \\ f'_{C_2} &= \gamma_0 + \gamma_2 L_{01} \\ f'_{C_2} &= \gamma_0 + \gamma_2 L_{01} \\ f'_{Y_1} &= \theta_0 \\ f'_{Y_2} &= \theta_0 \end{aligned}$$

4.3.2 Results

We considered three treatment effect levels: null, moderate and strong. In each case, the true value of the ACE,

$$\Psi_{II}(P) = E_P(R_2^{((a_0^*=1, \bar{c}_1=\bar{0}), (a_1^*=1, \bar{c}_2=\bar{0}))}) - R_2^{((a_0^*=0, \bar{c}_1=\bar{0}), (a_1^*=0, \bar{c}_2=\bar{0}))},$$

is empirically derived using 1,000,000 simulated subjects. A table for empirical distributions of the key variables will be presented for each treatment effect level.

Null Effect In order to create a setting in which the treatment has no effect on the recurrent event we set $\alpha_1 = \gamma_1 = \beta_1 = \theta_1 = 0$. The remaining coefficient values are given below.

$$\begin{aligned} \alpha_0 &= -1.6; \alpha_2 = 2; \alpha_3 = -1; \alpha_4 = -1; \alpha_5 = 2; \alpha_6 = -2; \\ \beta_0 &= -3; \beta_2 = 3; \beta_3 = 3; \beta_4 = 3; \beta_5 = 3; \beta_6 = 3; \\ \gamma_0 &= -0.7; \gamma_2 = -2; \gamma_3 = 1; \gamma_4 = -2; \gamma_5 = -2; \gamma_6 = 2; \\ \theta_0 &= -3.8; \theta_2 = 2; \theta_3 = -1; \theta_4 = 2; \theta_5 = 2; \theta_6 = 1. \end{aligned}$$

As expected, the ‘true’ ACE is zero using the empirical method, confirming there is no causal treatment effect. The marginal distributions of variables Y_1, Y_2, L_1 , and $R_2 = Y_1 + Y_2$ with respect to treatment levels are shown in Table 4.1 (without censoring) and the censoring rate is approximately 14%. Table 4.2 shows the estimated results under a null treatment effect.

	At(1,1)	At(0,0)	Overall
Y1	0.20043	0.20043	0.20043
Y2	0.20098	0.20098	0.20098
L1	0.50346	0.50346	0.50346
R2	0.40141	0.40141	0.40141
N=	1000000	ACE =	0.00000

Table 4.1: Average of key time-varying variables (Y_1, Y_2, L_1 and R_2) under null effect

Moderate Effect For a moderate treatment effect, the coefficient values are given below.

$$\begin{aligned} \alpha_0 &= -1.6; \alpha_1 = 2; \alpha_2 = 2; \alpha_3 = -2; \alpha_4 = -2; \alpha_5 = 2; \alpha_6 = -2; \\ \beta_0 &= -3; \beta_1 = -2; \beta_2 = 3; \beta_3 = 3; \beta_4 = 3; \beta_5 = 3; \beta_6 = 3; \\ \gamma_0 &= -2; \gamma_1 = 2; \gamma_2 = -3; \gamma_3 = 1; \gamma_4 = -2; \gamma_5 = -2; \gamma_6 = 4; \\ \theta_0 &= -3.5; \theta_1 = -1.5; \theta_2 = 2; \theta_3 = -1; \theta_4 = 2; \theta_5 = 2; \theta_6 = 1. \end{aligned}$$

The ‘true’ ACE is -0.27, indicating a moderate treatment effect. The marginal distributions of variables Y_1, Y_2, L_1 , and $R_2 = Y_1 + Y_2$ with respect to treatment levels are shown in Table 3.2 (without censoring), and we set the parameters such that approximately 13% of all subjects are censored. Table 4.4 shows the estimated results under a moderate treatment effect.

Strong Effect For a strong treatment effect, the coefficient values are given below.

$$\begin{aligned} \alpha_0 &= -1.6; \alpha_1 = 2; \alpha_2 = 2; \alpha_3 = -2; \alpha_4 = -2; \alpha_5 = 2; \alpha_6 = -2; \\ \beta_0 &= -3; \beta_1 = -2; \beta_2 = 3; \beta_3 = 3; \beta_4 = 3; \beta_5 = 3; \beta_6 = 3; \\ \gamma_0 &= -2; \gamma_1 = 2; \gamma_2 = -3; \gamma_3 = 1; \gamma_4 = -2; \gamma_5 = -2; \gamma_6 = 4; \\ \theta_0 &= -2.9; \theta_1 = -3; \theta_2 = 2; \theta_3 = -1; \theta_4 = 2; \theta_5 = 2; \theta_6 = 1. \end{aligned}$$

The ‘true’ ACE is -0.55, suggesting a strong treatment effect. The marginal distributions of variables Y_1 , Y_2 , L_1 , and $R_2 = Y_1 + Y_2$ with respect to treatment levels are shown in Table 4.5 (without censoring), and approximately 12% of all subjects are censored. Table 4.6 shows the estimated results under a strong treatment effect.

Overall, the Bootstrap ASE and ESE align with each other for all six estimators, confirming stable estimation across varying degrees of model (mis)specification and under different treatment effect strengths. Notably, when the models are correctly specified (scenario 0), all six estimators exhibit near-zero bias, mirroring the findings from earlier point treatment settings discussed in Section 3.3.2.

In scenarios where at most one working model is misspecified, the AIPW, WICE, and TML estimators continue to show minimal bias, underscoring their doubly robustness. The singly robust estimators display larger bias and increased standard errors, particularly, the IPW estimate in Scenario 1 and the ICE estimate in Scenario 2. Despite these challenges, coverage probabilities remain around 95% except when estimators exhibit significant bias due to model misspecification.

Overall, The AIPW, WICE, and TML estimators prove to be more reliable across all scenarios, offering unbiased results even when at least one of the treatment or outcome models is correct. In contrast, the singly robust estimators, including IPW, HIPW, and ICE, show limited efficacy when their respective models are misspecified.

Similar to the previous chapter, we recommend using the TML estimator in complex settings involving time-varying treatments for its doubly robustness and the flexibility of using machine learning algorithms in estimation. Therefore, the TML estimator will be implemented on the SPRINT in the following section.

Scenarios		Bias(x100)	ESE(x100)	Boot ASE(x100)	Cov. Prob.
Scenario0	IPW	0.083173	3.5670	3.4998	0.9585
	HIPW	0.257175	3.8625	3.8712	0.9540
	ICE	0.039320	3.3007	3.2284	0.9530
	AIPW	0.044065	3.3245	3.2395	0.9530
	WICE	0.039917	3.2964	3.2278	0.9515
	TML	0.042017	3.3055	3.2325	0.9535
Scenario 1	IPW	5.087815	3.8055	3.7889	0.7375
	HIPW	-1.581814	3.4888	3.4778	0.9280
	ICE	0.039320	3.3007	3.2290	0.9530
	AIPW	0.039320	3.3007	3.2290	0.9530
	WICE	0.037668	3.3023	3.2294	0.9525
	TML	0.039320	3.3007	3.2306	0.9530
Scenario 2	IPW	0.083173	3.5670	3.4995	0.9575
	HIPW	0.257175	3.8625	3.8716	0.9545
	ICE	2.689321	4.0695	4.0214	0.8970
	AIPW	0.118008	3.8654	3.9050	0.9570
	WICE	0.257175	3.8625	3.8716	0.9545
	TML	0.257175	3.8625	3.8716	0.9545
Scenario 3	IPW	5.087815	3.8055	3.7863	0.7345
	HIPW	-1.581814	3.4888	3.4761	0.9285
	ICE	2.689321	4.0695	4.0211	0.8965
	AIPW	-1.812276	3.4851	3.4762	0.9205
	WICE	-1.581814	3.4888	3.4761	0.9285
	TML	-1.581814	3.4888	3.4761	0.9285

Bias(x100): 100 times the estimation bias

ESE(x100): 100 times the empirical standard error

Boot ASE(x100): 100 times the Bootstrap Average Standard Error

Cov. Prob.: Coverage Probability

IPW: Inverse Probability Weighting (estimator)

HIPW: Hájek Inverse Probability Weighting (estimator)

ICE: Iterative Conditional Expectation (estimator)

AIPW: Augmented Inverse Probability Weighting (estimator)

WICE: Weighted Iterative Conditional Expectation (estimator)

TML: Targeted Maximum Likelihood (estimator)

Table 4.2: Table of estimation results under null effect based on 2,000 sets of 2,000 simulated individuals and 1,000 Bootstrap samples with six estimators: IPW, HIPW, ICE, AIPW, WICE and TML respectively

	At(1,1)	At(0,0)	Overall
Y1	0.10137	0.23471	0.16804
Y2	0.10229	0.23830	0.17029
L1	0.23228	0.51622	0.37425
R2	0.20366	0.47301	0.33833
N=	1000000	ACE =	-0.26935

Table 4.3: Average of key time-varying variables (Y_1 , Y_2 , L_1 and R_2) under moderate effect

Scenarios		Bias(x100)	ESE(x100)	Boot ASE(x100)	Cov. Prob.
Scenario0	IPW	-0.067915	3.1628	3.1638	0.9510
	HIPW	-0.045805	3.1752	3.1909	0.9525
	ICE	-0.086489	2.9067	2.9001	0.9530
	AIPW	-0.048017	2.9258	2.9183	0.9510
	WICE	-0.043261	2.9262	2.9185	0.9535
	TML	-0.047393	2.9258	2.9177	0.9510
Scenario1	IPW	-15.425191	3.9904	4.0647	0.0245
	HIPW	-0.943788	2.7551	2.8158	0.9460
	ICE	-0.086489	2.9067	2.9015	0.9535
	AIPW	-0.086489	2.9067	2.9015	0.9535
	WICE	-0.082665	2.9082	2.9016	0.9535
	TML	-0.086489	2.9067	2.9018	0.9535
Scenario2	IPW	-0.067915	3.1628	3.1625	0.9535
	HIPW	-0.045805	3.1752	3.1905	0.9535
	ICE	2.444013	2.9554	3.0267	0.8755
	AIPW	-0.070261	3.1600	3.1720	0.9530
	WICE	-0.045805	3.1752	3.1905	0.9535
	TML	-0.045805	3.1752	3.1905	0.9535
Scenario3	IPW	-15.425191	3.9904	4.0624	0.0225
	HIPW	-0.943788	2.7551	2.8140	0.9480
	ICE	2.444013	2.9554	3.0276	0.8715
	AIPW	-1.079232	3.0070	3.0697	0.9455
	WICE	-0.943788	2.7551	2.8140	0.9480
	TML	-0.943788	2.7551	2.8140	0.9480

Bias(x100): 100 times the estimation bias

ESE(x100): 100 times the empirical standard error

Boot ASE(x100): 100 times the Bootstrap Average Standard Error

Cov. Prob.: Coverage Probability

IPW: Inverse Probability Weighting (estimator)

HIPW: Hájek Inverse Probability Weighting (estimator)

ICE: Iterative Conditional Expectation (estimator)

AIPW: Augmented Inverse Probability Weighting (estimator)

WICE: Weighted Iterative Conditional Expectation (estimator)

TML: Targeted Maximum Likelihood (estimator)

Table 4.4: Table of estimation results under moderate effect based on 2,000 sets of 2,000 simulated individuals with six estimators: IPW, HIPW, ICE, AIPW, WICE and TML respectively

	$A_t(1,1)$	$A_t(0,0)$	Overall
Y1	0.053355	0.31641	0.18488
Y2	0.047096	0.33336	0.19023
L1	0.221106	0.54803	0.38457
R2	0.100451	0.64977	0.37511
N=	10000000	ACE =	-0.54932

Table 4.5: Average of key time-varying variables (Y_1 , Y_2 , L_1 and R_2) under strong effect

Scenarios		Bias(x100)	ESE(x100)	Boot ASE(x100)	Cov. Prob.
Scenario0	IPW	0.077611	3.2012	3.1299	0.9465
	HIPW	0.088231	3.1820	3.1166	0.9495
	ICE	0.044488	3.0664	2.9898	0.9475
	AIPW	0.085410	3.0848	3.0065	0.9465
	WICE	0.092769	3.0887	3.0104	0.9465
	TML	0.085705	3.0850	3.0177	0.9480
Scenario1	IPW	-22.789661	4.4405	4.4353	0.0000
	HIPW	1.712201	2.8762	2.8651	0.9060
	ICE	0.044488	3.0664	2.9896	0.9445
	AIPW	0.044488	3.0664	2.9896	0.9445
	WICE	0.047182	3.0673	2.9899	0.9455
	TML	0.044488	3.0664	3.0000	0.9455
Scenario2	IPW	0.077611	3.2012	3.1291	0.9450
	HIPW	0.088231	3.1820	3.1163	0.9480
	ICE	5.585180	3.0115	3.0096	0.5275
	AIPW	0.077063	3.1788	3.1127	0.9490
	WICE	0.088231	3.1820	3.1163	0.9480
	TML	0.088231	3.1820	3.1163	0.9480
Scenario3	IPW	-22.789661	4.4405	4.4356	0.0000
	HIPW	1.712201	2.8762	2.8644	0.9055
	ICE	5.585180	3.0115	3.0114	0.5265
	AIPW	0.949211	3.1578	3.1612	0.9435
	WICE	1.712201	2.8762	2.8644	0.9055
	TML	1.712201	2.8762	2.8644	0.9055

Bias(x100): 100 times the estimation bias

ESE(x100): 100 times the empirical standard error

Boot ASE(x100): 100 times the Bootstrap Average Standard Error

Cov. Prob.: Coverage Probability

IPW: Inverse Probability Weighting (estimator)

HIPW: Hájek Inverse Probability Weighting (estimator)

ICE: Iterative Conditional Expectation (estimator)

AIPW: Augmented Inverse Probability Weighting (estimator)

WICE: Weighted Iterative Conditional Expectation (estimator)

TML: Targeted Maximum Likelihood (estimator)

Table 4.6: Table of estimation results under strong effect based on 2,000 sets of 2,000 simulated individuals with six estimators: IPW, HIPW, ICE, AIPW, WICE and TML respectively

4.4 Application

In this section, we revisit the SPRINT dataset used for point treatment estimation in Chapter 3.4, but delve deeper into its analysis. After properly transforming the data, we consider an intervention on both the initial assignment of the intensive or standard blood pressure lowering therapy, and participants’ adherence to the assignment over time. The intervention considered in this application can be regarded as an extension to that of Chapter 3.4, whereby in this intervention, participants are forced to strictly adhere to their assigned blood pressure lowering therapy. We will then compare the results from these analysis and draw conclusions accordingly.

4.4.1 Data Description

In Chapter 3, we implemented an intention-to-treat (ITT) analysis, where all participants were analyzed in the groups to which they were originally assigned at randomization. However, patients and clinicians may be more interested in what would happen to some outcome of interest if the treatment was strictly followed as recommended or prescribed, i.e., the per-protocol (PP) effect. As such, in this section, we consider time-varying nature of ‘treatment’, where its definition is extended to encompass patients’ adherence level to their assigned treatments. The goal is to estimate the causal effect on counterfactual recurrent events if individuals had fully adhered to their prescribed blood pressure-lowering medications for an intensive treatment compared to standard treatment, as per-protocol. Our new time-varying exposure variable is defined as $(Z, \bar{A}_k)$ at time interval k for $k = 0, \dots, K$ within the study period, where Z refers to the assignment of intensive or standard blood pressure lowering therapy and $\bar{A}_K$ represents the time-varying adherence status of uncensored participants throughout the study period.

In Section 3.4 of the previous chapter, we estimated the ACE of the static treatment assignment on outcome recurrences while controlling for participants’ censoring status. It is worth noting that during the study period (i.e. the first four years of the trial) all censored participants in SPRINT were due to death. Thus, in the original analysis of the SPRINT data completed in Section 3.4, death is not differentiated from censoring events. As such, the causal estimand of interest in Section 3.4, which eliminates potential mediators (i.e., death) by directly controlling for them, is referred to as the controlled direct effect (CDE) [78].

However, it is difficult to envision an intervention that prevents mortality in all individuals. In other words, incorporating death as a component of the intervention

may not be practical or feasible. Alternatively, one could instead consider death as a separate time-varying variable, which opens an additional causal path from the treatments to the recurrent outcomes [48]. Similar to a “mediator”, death is considered as a semi-competing event to the primary outcome of acute kidney recurrences [1]. In this semi-competing risks setting, our research interest focuses on the recurrent outcome of acute kidney injuries, which is non-terminal; while the acute kidney injuries are subject to the terminal event of death [33]. We are interested in the total effect of treatments, which includes both the direct treatment effect and the indirect semi-competing effect through death.

Define D_k as the indicator of whether a participant is dead at interval k . Since there is no lost to follow-ups during the study period and D_k is regarded as an event that prevents Y_k from occurring, we do not need to consider the censoring process in our estimation procedure for the total effects estimands. That is, we have $\bar{A}_K = (\bar{A}_K^*, \bar{C}_{K+1} = \text{NULL}) = \bar{A}_K^*$. We are still interested in the difference in the expected total acute kidney recurrences had all participants strictly adhered to an assigned intensive blood pressure lowering therapy, compared to if they had fully adhered to the standard therapy. However, unlike the approach taken in the previous chapter, we do not treat death as a censoring event. We assume the key variables follow the chronological order as in Figure 4.3.

4.4.2 Data Transformation

As discussed in Section 3.4, we select the first four years (or 1,460 days) of the SPRINT dataset as our study period, and divided them into two time intervals ($t \in \{1, 2\}$). Then, we perform a data transformation process very similar to before by discretizing the continuous-time variables A_k , Y_k and $C_k(D_k)$ ($k \in \{1, \dots, 1460\}$) into binary indicators through the same threshold process. Then, the time-varying covariate L_k is calculated as the average of all MAP recordings. After transformation, A_k becomes the indicator of full compliance to doctor-prescribed blood pressure lowering medications during the transformed interval of a patient that is uncensored at the next interval. Note that participants’ adherence information was originally collected from participants’ self-reported surveys asking the percentage of days they have taken all of blood pressure lowering medications as prescribed since last visit. We assume that participants’ adherence level at the beginning of the coarsened time interval reflects participants’ true adherence of that interval.

Although some information is inevitably lost through the transformation process, the simplification enables a more straightforward analysis and interpretation while




Figure 4.3: A single world intervention graph (SWIG) that identifies the average causal effect and illustrates the relationships between the measured baseline variables L_0 , the joint treatment variable of initial treatment assignment & time-varying intervention of adherence $(Z, \bar{A}_1)$, the semi-competing event of death $\bar{D}_2$, the time-varying confounder L_1 , and the (intermediate) recurrent outcome $\bar{Y}_2$, assuming there is no unmeasured confounding.

preserving the main characteristics of the data. Table 4.7 show the average of the time-varying variables after data coarsening. Recall that R_2 counts the cumulative number of recurrent events that have occurred from the beginning of the study up to interval 2, or $R_2 = Y_1 + Y_2$. Then, our estimand of the total PP effect can be formulated as $E_P(R_2^{\bar{a}_1, z=1}) - E_P(R_2^{\bar{a}_1, z=0})$.

4.4.3 Total Per-Protocol Effect of the Intensive vs. Standard Blood Pressure Lowering Therapy

In Section 3.4, we estimated the ITT CDE of the intensive versus the standard blood pressure lowering therapy, where all participants were analyzed in the originally-assigned (intensive or standard) therapy groups. We additionally estimate the total

Time-varying Variables	Full Adherence	Partial/Non- Adherence	Overall Average
Number of Subjects at $t = 1$	8022	1300	0.961
Dead at $t = 1$ (D_1)	0.017	0.030	0.019
Event status at $t = 1$ (Y_1)	0.021	0.024	0.021
Mean arterial pressure at $t = 1$ (L_1)	89.894	92.332	90.234
Number of Subjects at $t = 2$	6757	2565	0.725
Dead at $t = 2$ (D_2)	0.022	0.113	0.047
Event status at $t = 2$ (Y_2)	0.018	0.013	0.017
Total acute kidney injury events (R_2)	0.037	0.039	0.038

Table 4.7: The Average or Count of Key Time-varying Variables in the First Interval ($t = 1$) or in the Second Interval ($t = 2$) for the SPRINT Study Population Included in the Analysis

ITT effect in Appendix D.4.1 to be comparable with the total per-protocol effect calculated in this chapter.

Under the causal assumptions 1–4 given above, our causal estimand of $E_P(R_2^{\bar{a}_1, z=1}) - E_P(R_2^{\bar{a}_1, z=0})$ can be identified from the observed data of SPRINT and is estimated to be 0.017 with a 95% CI of [0.008, 0.026] utilizing the TML estimator. The standard error for the the confidence interval construction was obtained using 1,000 nonparametric bootstrap samples. This implies that, following a PP analysis, SPRINT participants under full adherence to their assigned intensive blood pressure lowering therapy would experience marginally more recurrent acute kidney injury events within the first four years of the study compared to them fully adhering to the standard therapy. Notably, the total PP effect estimate is practically identical to those from the total ITT analysis discussed in Section D.4.1. This is not surprising given that, when estimating the population intervention effects of adherence to prescribed blood pressure medications in both the intensive and standard therapy groups, as detailed in Appendix D.4.2, the results indicate that participants' adherence does not significantly affect the risk of recurrent acute kidney injuries in SPRINT.

4.5 Conclusion

In contrast to Chapter 3, this chapter focuses on estimating the ACE of time-varying interventions on recurrent event outcomes with censoring. We explore various singly and doubly estimation methods including TMLE, and expand our theoretical validation using simulations. The simulation results demonstrated that the TML, WICE,

and AIPW estimators performed best regarding their double robustness against model misspecifications. Among them, we recommend the TML estimator for its doubly robustness and the potential to incorporate machine learning in the update step.

As an application, we utilized the SPRINT dataset to demonstrate our methods in a time-varying treatments context and extend on more complex estimands and settings. The design of SPRINT allows for comparison between the ITT and PP analyses. The ITT analysis offers insights into the effectiveness of the assigned treatment regimens irrespective of the actual adherence (estimated in Section 3.4), whereas in this chapter, the PP analysis focuses on outcomes for participants who adhered fully to their assigned treatments, thus isolating the effect of the treatment itself.

Additionally, our analyses account for the semi-competing risk of death to the recurrent outcome of acute kidney injuries. We estimated the total effect of treatment which encompasses both the direct treatment effect and the indirect treatment effect mediated by participants' death. However, in estimating the total effect, we face the challenge of distinguishing the effect of treatment on recurrent events, whether mediated or not mediated through the semi-competing risk of death. Addressing this issue may involve exploring the separable effect, which will be discussed as future work in Chapter 5.

Chapter 5

Conclusion and Future Work

5.1 Conclusion

In this dissertation, we developed methods to identify and estimate the causal effect of treatment strategies on recurrent event outcomes in various complex longitudinal settings.

Chapter 1 began with an introduction of two real life datasets, which provided motivation for investigating causal questions involving recurrent event outcomes. In this chapter, we also discussed the counterfactual causal framework, and reviewed essential concepts and terminologies that underpin our analytical approaches. This chapter also offered a comprehensive review of various causal estimands and corresponding estimation methods in non-recurrent event settings, giving us the necessary tools for conducting effective causal analysis on more complex longitudinal studies with recurrent event outcomes.

In Chapter 2, we expanded upon conventional causal measure of the marginal odd ratio (MOR), particularly focusing on scenarios involving time-varying treatments and outcomes. Traditionally, MORs have been employed to explore the relationship between a binary treatment and a time-fixed binary outcome, with estimations being derived through unstabilized inverse probability weighting (IPW) and logistic regression following Austin’s foundational designs [4]. We advanced this concept to handle more complex settings where treatments vary over time, developing a modified MOR ($\text{MOR}_{\text{int.}}$), and demonstrating through Monte Carlo simulations that this could be estimated with low bias and stable variance. Further, we introduced the Aggregated Marginal Odds Ratio (AMOR) for scenarios where both the treatment

and the outcome are discrete and changed over time. Our studies utilized a counterfactual framework to define AMOR as a ratio of outcome odds under constant exposure versus no exposure to the treatment, across several time periods. We examined different treatment scenarios—nonabsorbing and absorbing—and discovered that the AMOR could be reliably estimated with low bias, particularly in shorter study periods and using stabilized weight models. As an illustrative example, the proposed causal measure and estimation methods are applied to the 1997 National Longitudinal Study of Youth (NLSY97). In this study, we found that that smoking behavior potentially had a minor negative causal impact on school enrollment for youths during the study period.

In Chapter 3, we shifted the focus from proposing new causal effect measures to estimating the average causal effect (ACE) of static treatment strategies on recurrent event outcomes that are subject to censoring. Based on the theoretical derivations (provided in Appendix C.2), we presented the estimation details of various estimators with their corresponding estimating formulas, particularly the targeted maximum likelihood (TML) estimator. Simulations were carried out to confirm the theoretical properties of the proposed methods (e.g., to confirm that in this setting, TML estimators are indeed doubly robust against model misspecifications). Furthermore, the implementation of the proposed methodologies on the Systolic Blood Pressure Intervention Trial (SPRINT) data through an intention-to-treat (ITT) analysis utilizing the TMLE algorithm provided insightful results. The analysis suggested a positive ACE of intensive blood pressure lowering therapy on acute kidney injury recurrences compared to the standard treatment for those who survived the first four years of the SPRINT.

Building on the foundational methodologies introduced in the preceding chapter, Chapter 4 extends the analysis to the ACE of interventions on time-varying treatment, again focusing on recurrent event outcomes with censoring. Based on Chapter 3, this chapter explores the singly and doubly robust estimation methods, particularly TMLE, with theoretical derivations in the time-varying treatment context. Theoretical frameworks for these estimators were rigorously validated via simulation studies, demonstrating their applicability and robustness. Utilizing the SPRINT dataset, we conducted a per-protocol (PP) analysis that takes into account participants' time-varying adherence to medications, as opposed to the ITT analysis carried out in Chapter 3. In particular, we estimated the total ACEs of intensive versus standard blood pressure lowering therapies on acute kidney recurrences had all participants fully adhered to their prescribed blood pressure medications during the study period. The results of the PP analysis indicate that, the participants fully

comply to their assigned intensive blood pressure lowering therapy are expected to experience marginally more acute kidney injury recurrences than those adhere to the the standard therapy in the first four years of the SPRINT. Additionally, we noted that the effect estimate of the PP analysis that includes both the controlled direct effect and the effect mediated by semi-competing risk of death, closely align with the ITT effect estimate in Chapter 3. This is consistent with our finding that the adherence to prescribed blood pressure medications in either the intensive or standard therapy arm does not significantly impact the recurrent outcome of interest.

To conclude, the research goal throughout this dissertation is to enhance our understanding and develop methodologies to estimate the complex causal relationships between treatment interventions and recurrent event outcomes. Although these efforts have shed light on these complex relationships, there is more to explore and refine, as will be discussed in the next section.

5.2 Future Work

Moving forward, the objective is to broaden the analytical scope and address key challenges in causal inference involving complex longitudinal settings based on this thesis.

5.2.1 Extensions of AMOR

In Chapter 2, the Aggregated Marginal Odds Ratio (AMOR) is specifically developed to estimate the causal impact of time-varying treatments on recurrent event outcomes. However, the assumption of a binary treatment status—either fully treated or never treated throughout the study—limits the applicability of AMOR in clinical settings where treatment regimens often vary over time. Future work should therefore aim to extend the definition of AMOR to allow for a spectrum of intermittent and dosage-varying treatment strategies, i.e.,

$$\text{Extended AMOR} = \frac{P(\bar{Y}(T)^{(\bar{A}(T)=\bar{a}_1(T))} = \bar{1})}{1 - P(\bar{Y}(T)^{(\bar{A}(T)=\bar{a}_1(T))} = \bar{1})} \bigg/ \frac{P(\bar{Y}(T)^{(\bar{A}(T)=\bar{a}_2(T))} = \bar{1})}{1 - P(\bar{Y}(T)^{(\bar{A}(T)=\bar{a}_2(T))} = \bar{1})},$$

where $P(\bar{Y}(T)^{(\bar{A}(T)=\bar{a}_d(T))} = \bar{1})$ represents the probability of the occurrence of the

recurrent event outcome across all time intervals under a pre-specified treatment path $\bar{a}_d, d = 1, 2$.

In the weight estimation of AMOR, the current analysis only includes covariates from the true treatment model in the weight numerator and denominator, with the goal of reducing bias. However, it is known that including covariates that are related only to the outcome in the propensity score model can reduce variance in estimation [12]. In future work, we will explore additional propensity score models that incorporate $\mathbf{X}_{\text{out}}$.

Furthermore, longitudinal studies with missing data (possibly due to death, lost to follow-up, or sporadic missed visits) pose additional challenges, and may prevent the observation of the primary event of interest. These challenges necessitate the integration of methodologies for handling missing data into the AMOR framework. Techniques such as multiple imputation [88, 86] could potentially improve the accuracy of the estimated effects but also extend the usability of AMOR in more clinically realistic scenarios.

5.2.2 Improvements on Data Pre-processing

To ensure a more accurate depiction of patient behavior and a more accurate estimation of ACEs in the Systolic Blood Pressure Intervention Trial (SPRINT) data, we plan to refine our data processing strategies. In the current data preparation stage, data was coarsened to equally-spaced (bi-yearly) time intervals which inevitably led to information loss. One main concern is that some of the measured variables, such as monthly and tri-monthly blood pressure measurements, may be treated as unmeasured confounders after bi-yearly data coarsening. Moreover, if a recurrent event outcome fails to be captured in an analysis, then its unmeasured effect on other variables may distort and bias the ACE of interest. Figure 5.1 demonstrates a scenario where the causal effect cannot be identified because a recurrent event outcome Y_* was not captured in the analytical data, and thus its unmeasured effect on other variables biases the ACE of interest. In other words, the unmeasured confounder Y_* opens a backdoor paths between treatment A_1 and outcome Y_2 given the measured covariate history L_1 . Though we have carefully examined the data and discretized our time intervals to ensure there is no such unmeasured recurrent outcome in our analysis for Chapters 3 and 4, it is important to address these potential issues in future studies.

Another potential issue that can lead to biased estimation and/or incorrect inference is our definition of adherence in the SPRINT data analysis. In Chapter 4,

we assumed that participants’ adherence level at the beginning of each coarsened time interval reflects participants’ true adherence to blood pressure medication during that interval. In other words, we allowed participants’ adherence levels to carry forward from the start of each coarsened time interval to the end of that interval. This assumption is reasonable for SPRINT as participants’ compliance rate is high in the study. In more general settings, however, if a participant does not fully adhere to the medications midway through a coarsened interval but is recorded as fully adhered in that interval, the coarsened adherence data would be inaccurate. Moreover, failing to consider possible adherence fluctuations within each time interval, similar to Figure 5.1, may open up new backdoor pathways given past covariates.

The direct solution to these issues is skipping the data coarsening process and preserving the original data density. However, in some cases, data coarsening may just be unavoidable due to the structure of the data collection. For example, data were collected monthly for the initial three months, and subsequently, every three months. Note that challenges specifically related to unmeasured confounding will be discussed in the next section.

To address our previous concerns, we could also conduct a sensitivity analysis on data granularity to assess the robustness of our findings. For instance, by considering different coarsening scenarios where adherence may vary within smaller intervals, we can better understand the impact of adherence variations within the coarsened intervals on treatment effects estimation. Furthermore, establishing upper and lower bounds for PP effects can also enhance the interpretation of our results [115]. This approach involves determining the range within which the true effect of treatment adherence lies, accounting for potential misclassification biases. Such bounds provide a more comprehensive view of the potential impacts of treatment, facilitating more informed decision-making in clinical settings.

5.2.3 Dynamic Treatment Strategies

Dynamic Treatment Regimes (DTR) provide a structured approach to evolving treatment strategies over time, allowing for adjustments based on individual patient responses and changing conditions [61]. To account for the evolving nature of individual conditions or external influences that could affect the treatment efficacy over time, we can consider extending our methodology from static treatment settings to a dynamic treatment framework. In the context of SPRINT, a DTR would involve intervening on blood pressure management strategies dynamically based on some stochastic rules or on health indicators collected from previous visits. For instance, a DTR could

their usual manner, including occasional deviation in medication intake.

This approach introduces a degree of realism often absent in rigid clinical trial settings, acknowledging that patients may not always follow treatment protocols precisely due to various personal and contextual factors. Implementing such grace periods can potentially enhance the applicability and generalizability of study results to everyday clinical practices.

5.2.4 Methods for Unmeasured Confounding

In the realm of causal inference, unmeasured confounding remains one of the most challenging obstacles to obtaining valid inference. Traditionally, many analytical approaches assume that all common causes of an exposure and outcome of interest are measured, an assumption which, if violated, can significantly undermine the validity of the study’s findings. Sensitivity analysis is a crucial approach for assessing how sensitive the conclusions are to the presence of unmeasured confounding. This method estimates the degree to which an unmeasured confounder would need to affect both the treatment and the outcome to explain away a treatment effect. Techniques for conducting sensitivity analysis are extensively discussed in the literature, providing tools for researchers to address the potential bias introduced by unmeasured confounders [83].

The unmeasured confounding can also be pronounced in the SPRINT, where variables such as stress and other psychological factors that affect adherence to prescribed treatments medication, are not measured. These factors can also influence the recurrence of acute kidney injuries, which classifies them as potential unmeasured confounders in assessing PP effects. Specifically, we assume that psychological stress can influence participants in two ways: one way is to affect their adherence to the intensive blood pressure treatment regimen (through stress-related mental issues); and the other way is to potentially increase the risk of acute kidney injury recurrences (through physiological changes affecting kidney function).

Various methods have been proposed to adjust for unmeasured confounders. Proxy variables can serve as stand-ins for unmeasured confounding. For instance, if psychological stress is an unobserved confounder affecting both the intervention (adherence) and the outcome (acute kidney injuries), a related measure such as the frequency of healthcare visits or prescription of anti-anxiety medications might serve as a proxy.

While proxy variables provide a straightforward method to account for unseen

confounders, they often fall short in fully capturing the complex relationship between confounders, treatment, and outcomes, potentially introducing new biases. Proximal learning, on the other hand, leverages structured sets of proxy variables, known as negative control exposures and outcomes, to better correcting the unmeasured confounding effects [56]. In the SPRINT, a suitable negative control exposure could be the prescription of medications unrelated to blood pressure but related to stress management, such as anxiolytics. This prescription data, which is available in SPRINT, can be used as it is likely influenced by the participant’s stress levels but theoretically should not directly affect the recurrent event outcomes, except through its impact on stress. An appropriate negative control outcome might be the occurrence of sleep disturbances, which are also recorded in the SPRINT. These outcomes are expected to be influenced by stress but are generally not a direct consequence of blood pressure management strategies. The simultaneous use of these controls allows for identification of the ACE even with the presence of unmeasured confounders.

The front door criterion [65] provides yet another methodological framework when a mediating variable can be placed between the treatment and the outcome. This mediating variable is influenced only by the treatment and also mediates the effect of the treatment on the outcome. In the SPRINT context, it may be reasonable to assume that adherence to medication affects recurrent kidney injuries through a mediating variable, such as clinical health metrics. In this context, the mediating variable that mediates the effect of adherence on kidney injury recurrences is blood pressure measurements, which is also measured in the dataset.

5.2.5 Competing events and Separable Effects

Separable effects provides another analytical framework in examining the impact of blood pressure lowering therapy on outcomes related to acute kidney injury. This approach is particularly useful when considering the dual objectives of blood pressure management in patients at high cardiovascular risk: reducing blood pressure to prevent cardiovascular events and managing the risk of inducing or exacerbating kidney injury.

Blood pressure medications are designed primarily to lower blood pressure and, by extension, reduce the risk of cardiovascular complications. However, these medications may also have effects that are independent of their cardiovascular benefits. For instance, while part of the medication’s effect might be directed towards achieving the target blood pressure reduction — thereby lowering the risk of heart attacks

or strokes — another part might inadvertently influence renal function, potentially leading to acute kidney injury or affecting its recurrence.

In this context, the outcome of interest is the recurrence of acute kidney injury, and death serves as a semi-competing effect. This is because death prevents the observation of any further kidney injury episodes, thereby complicating the analysis of the medication’s impact on kidney health. By adopting a separable effects approach, researchers can disentangle the effect of blood pressure lowering therapy on kidney injury from its effect on mortality.

For example, intensive blood pressure lowering therapy might be hypothesized to have a separable effect where one aspect of the treatment effectively reduces blood pressure and the associated risk of death from cardiovascular causes. Simultaneously, another aspect of the same treatment could pose a risk to renal health, possibly increasing the likelihood of acute kidney injury episodes. Understanding these separable effects allows for a more comprehensive evaluation of the treatment’s net impact — identifying the beneficial effects of blood pressure reduction on cardiovascular health while also acknowledging and quantifying any adverse effects on kidney function. Figure 5.2 shows an illustrative SWIG for the separable effects of some blood pressure lowering therapy.

By focusing on the separable effects, clinicians and researchers can better assess the trade-offs involved in the blood pressure management strategies, especially in populations at high risk for both cardiovascular events and kidney complications. This analytical perspective supports more informed decision-making in clinical practice, guiding the optimization of treatment strategies to maximize cardiovascular benefits while minimizing risks to renal health.

5.2.6 Enhancements for the current TML estimator

The current study employs targeted maximum likelihood (TML) estimators, a powerful tool in causal inference designed to provide efficient, doubly robust estimates even in the presence of complex data structures. As part of my future work, we intend to further enhance the capabilities of these estimators by incorporating advanced machine learning techniques. Specifically, the integration of adaptive LASSO and kernel methods into the working models for nuisance functions promises to refine the accuracy of our estimations.

The adaptive LASSO method (Least Absolute Shrinkage and Selection Operator), an extension of the traditional LASSO, provides a robust approach to variable selection and regularization in high-dimensional data scenarios. By assigning different




Figure 5.2: A single world intervention graph (SWIG) identifies the separable effects of some time-varying treatment intervention A_k on the semi-competing event D_{k+1} and the (intermediate) recurrent outcome Y_{k+1} , with the measured time-varying confounder L_k for $k \in \{0, \dots, K\}$, assuming no unmeasured confounding. Note that we use **red** arrows to indicate the hypothetical separation of the treatment effect.

weights to different coefficients, adaptive LASSO improves the selection consistency and is particularly adept at retaining relevant variables while eliminating those that are redundant or irrelevant [119]. Integrating the adaptive LASSO method in the update step for TML estimator would allow us to enhance its capability to deal with complex causal analyses involving numerous predictors.

Kernel methods, including support vector machines and Gaussian processes, are renowned for their ability to capture complex, non-linear relationships between variables. By transforming the data into a higher-dimensional space, these methods allow for the modeling of non-linear interactions that are often present in biomedical data sets [91]. Integrating kernel methods into TML estimators can significantly enhance the model's ability to capture complex relationships within the data.

By integrating advanced machine learning techniques into TML estimators, my future work aims to not only increase the accuracy of ACE estimates but also enhance the ability to tailor interventions to individual patient needs. These enhancements will provide deeper insights into the causal dynamics of complex datasets, paving

the way for more informed clinical decision-making and policy formulation.

References

- [1] Odd Aalen, Ornulf Borgan, and Hakon Gjessing. *Survival and event history analysis: a process point of view*. Springer Science & Business Media, 2008.
- [2] Leila D A F Amorim and Jianwen Cai. Modelling recurrent events: A tutorial for analysis in epidemiology. *International Journal of Epidemiology*, 44:324–333, 2015.
- [3] Joshua Anderson. Hume, causation and counterfactuals. *Humanities Bulletin*, 2:36–49, 2019.
- [4] Peter C Austin. The performance of different propensity score methods for estimating marginal odds ratios. *Statistics in Medicine*, 26:3078–3094, 2007.
- [5] Peter C Austin. An introduction to propensity score methods for reducing the effects of confounding in observational studies. *Multivariate Behavioral Research*, 46:399–424, 2011.
- [6] Peter C Austin, Paul Grootendorst, and Geoffrey M Anderson. A comparison of the ability of different propensity score models to balance measured variables between treated and untreated subjects: A monte carlo study. *Statistics in Medicine*, 26:734–753, 2007.
- [7] Peter C Austin and Elizabeth A Stuart. Moving towards best practice when using inverse probability of treatment weighting (IPTW) using the propensity score to estimate causal treatment effects in observational studies. *Statistics in Medicine*, 34:3661–3679, 2015.
- [8] Heejung Bang and James M Robins. Doubly robust estimation in missing data and causal inference models. *Biometrics*, 61(4):962–973, 2005.

- [9] Helen Beebee and Peter Menzies. Counterfactual theories of causation. In *Stanford Encyclopedia of Philosophy*. 2019.
- [10] Brandon K Bellows, Yiyi Zhang, Zugui Zhang, Donald M Lloyd-Jones, Adam P Bress, Jordan B King, Paul Kolm, William C Cushman, Karen C Johnson, Leonardo Tamariz, et al. Estimating systolic blood pressure intervention trial participant posttrial survival using pooled epidemiologic cohort data. *Journal of the American Heart Association*, 10(10):e020361, 2021.
- [11] Lisa M Bodnar, Marie Davidian, Anna Maria Siega-Riz, and Anastasios A Tsiatis. Marginal structural models for analyzing causal effects of time-dependent treatments: an application in perinatal epidemiology. *American Journal of Epidemiology*, 159:926–934, 2004.
- [12] M Alan Brookhart, Sebastian Schneeweiss, Kenneth J Rothman, Robert J Glynn, Jerry Avorn, and Til Stürmer. Variable selection for propensity score models. *American Journal of Epidemiology*, 163:1149–1156, 2006.
- [13] U.S. Department of Labor Bureau of Labor Statistics. National longitudinal survey of youth 1997 cohort, 1997-2017 (rounds 1-18), The Ohio State University. Columbus, OH: 2019. Produced and distributed by the Center for Human Resource Research (CHRR).
- [14] Stephen R Cole and Constantine E Frangakis. Commentary: The consistency statement in causal inference: A definition or an assumption? *Epidemiology*, 20:3–5, 2009.
- [15] Stephen R Cole and Miguel A Hernán. Constructing inverse probability weights for marginal structural models. *American Journal of Epidemiology*, 168:656–664, 2008.
- [16] John Concato, Nirav Shah, and Ralph I Horwitz. Randomized, controlled trials, observational studies, and the hierarchy of research designs. *New England journal of medicine*, 342(25):1887–1892, 2000.
- [17] Nancy R Cook, Stephen R Cole, and Charles H Hennekens. Use of a marginal structural model to determine the effect of aspirin on cardiovascular mortality in the physicians’ health study. *American Journal of Epidemiology*, 155:1045–1053, 2002.
- [18] Richard J Cook and Jerald F Lawless. *The statistical analysis of recurrent events*. New York: Springer, 2007.

- [19] Cecilia A Cotton and Patrick J Heagerty. A data augmentation method for estimating the causal effect of adherence to treatment regimens targeting control of an intermediate measure. *Statistics in Biosciences*, 3:28–44, 2011.
- [20] James Cui. QIC program and model selection in GEE analyses. *The Stata Journal*, 7:209–220, 2007.
- [21] Daniel DeMers and Daliah Wachs. Physiology, mean arterial pressure. In *StatPearls [Internet]*. StatPearls Publishing, 2021.
- [22] Catherine Fitzmaurice and Ken Pease. *The psychology of judicial sentencing*. Manchester University Press Manchester, 1986.
- [23] Andrew Forbes and Susan Shortreed. Inverse probability weighted estimation of the marginal odds ratio: Correspondence regarding ‘the performance of different propensity score methods for estimating marginal odds ratios’ by p. austin, *statistics in medicine*, 2007; 26:3078–3094. *Statistics in Medicine*, 27:5556–5559, 2008.
- [24] John A Gallis and Elizabeth L Turner. Relative measures of association for binary outcomes: Challenges and recommendations for the global health researcher. *Annals of Global Health*, 85:1–12, 2019.
- [25] Jonathan Gillard et al. *A First Course in Statistical Inference*. Springer, 2020.
- [26] Stephen E Gilman, Laurie T Martin, David B Abrams, Ichiro Kawachi, Laura Kubzansky, Eric B Loucks, Richard Rende, Rima Rudd, and Stephen L Buka. Educational attainment and cigarette smoking: a causal association? *International journal of epidemiology*, 37(3):615–624, 2008.
- [27] Erika Graf and Martin Schumacher. Comments on “the performance of different propensity score methods for estimating marginal odds ratios“. *Statistics in Medicine*, 27:3915–3917, 2008.
- [28] Sander Greenland, Mohammad Ali Mansournia, and Douglas G Altman. Sparse data bias: a problem hiding in plain sight. *BMJ*, 352:i1981, 2016.
- [29] Sandeep K Gupta. Intention-to-treat concept: a review. *Perspectives in clinical research*, 2(3):109–112, 2011.
- [30] Jaroslav Hajek. Asymptotic Theory of Rejective Sampling with Varying Probabilities from a Finite Population. *The Annals of Mathematical Statistics*, 35(4):1491 – 1523, 1964.

- [31] M Elizabeth Halloran and Claudio J Struchiner. Causal inference in infectious diseases. *Epidemiology*, pages 142–151, 1995.
- [32] Frank R Hampel. The influence curve and its role in robust estimation. *Journal of the american statistical association*, 69(346):383–393, 1974.
- [33] Sebastien Haneuse and Kyu Ha Lee. Semi-competing risks data analysis: accounting for death as a competing risk when the outcome of interest is non-terminal. *Circulation: Cardiovascular Quality and Outcomes*, 9(3):322–331, 2016.
- [34] M A Hernán, B A Brumback, and J M Robins. Estimating the causal effect of zidovudine on CD4 count with a marginal structural model for repeated measures. *Statistics in Medicine*, 21(12):1689–1709, 2002.
- [35] Miguel A Hernán. The hazards of hazard ratios. *Epidemiology*, 21(1):13–15, 2010.
- [36] Miguel A Hernán, Babette A Brumback, and James M Robins. Marginal structural models to estimate the joint causal effect of nonrandomized treatments. *Journal of the American Statistical Association*, 96:440–448, 2001.
- [37] Miguel A Hernán, Sonia Hernández-Díaz, and James M Robins. A structural approach to selection bias. *Epidemiology*, 15(5):615–625, 2004.
- [38] Miguel A Hernán, James M Robins, et al. Per-protocol analyses of pragmatic trials. *N Engl J Med*, 377(14):1391–1398, 2017.
- [39] Miguel Angel Hernán, Babette A Brumback, and James M Robins. Marginal structural models to estimate the causal effect of zidovudine on the survival of hiv-positive men. *Epidemiology*, 11:561–570, 2000.
- [40] Miguel A Hernán. Beyond exchangeability: The other conditions for causal inference in medical research. *Statistical Methods in Medical Research*, 21:3–5, 2012.
- [41] Miguel A Hernán and James M. Robins. *Causal Inference: What If*. Boca Raton: Chapman & Hall/CRC, 2020.
- [42] Oliver Hines, Oliver Dukes, Karla Diaz-Ordaz, and Stijn Vansteelandt. Demystifying statistical learning based on efficient influence functions. *The American Statistician*, 76(3):292–304, 2022.

- [43] M Höfler. Causal inference based on counterfactuals. *BMC Medical Research Methodology*, 5:28, 2005.
- [44] Paul W Holland. Statistics and causal inference. *Journal of the American Statistical Association*, 81:945–960, 1986.
- [45] Daniel G Horvitz and Donovan J Thompson. A generalization of sampling without replacement from a finite universe. *Journal of the American statistical Association*, 47(260):663–685, 1952.
- [46] Jingwen Hu, Clifford C Chou, King H Yang, and Albert I King. A weighted logistic regression analysis for predicting the odds of head/face and neck injuries during rollover crashes. In *Annual Proceedings/Association for the Advancement of Automotive Medicine*, volume 51, page 363. Association for the Advancement of Automotive Medicine, 2007.
- [47] Guido W Imbens and Donald B Rubin. *Causal inference in statistics, social, and biomedical sciences*. Cambridge university press, 2015.
- [48] Matias Janvin, Jessica G. Young, Pål C. Ryalen, and Mats J. Stensrud. Causal inference with recurrent and competing events, 2022.
- [49] Aksel K G Jensen, Henrik Ravn, Signe Sørup, and Per Andersen. A marginal structural model for recurrent events in the presence of time-dependent confounding: non-specific effects of vaccines on child hospitalisations. *Statistics in Medicine*, 35:5051–5069, 2016.
- [50] Michael R Kosorok and Eric B Laber. Precision medicine. *Annual review of statistics and its application*, 6:263–286, 2019.
- [51] Jonathan Levy. Tutorial: Deriving the efficient influence curve for large models, 2019.
- [52] Yuelin Li, Jonathan Baron, Yuelin Li, and Jonathan Baron. Reading and transforming data format. *Behavioral research data analysis with r*, pages 19–37, 2012.
- [53] Haodi Liang and Cecilia Cotton. Causal inference with recurrent data via inverse probability treatment weighting method (IPTW). *arXiv preprints arXiv:1911.06868*, 2019.

- [54] Travis M Loux, Christiana Drake, and Julie Smith-Gagen. A comparison of marginal odds ratio estimators. *Statistical Methods in Medical Research*, 26:155–175, 2017.
- [55] Eduardo Meaney, Felix Alva, Rafael Moguel, Alejandra Meaney, JUAN ALVA, and RICHARD WEBEL. Formula and nomogram for the sphygmomanometric calculation of the mean arterial pressure. *Heart*, 84(1):64–64, 2000.
- [56] Wang Miao, Zhi Geng, and Eric J Tchetgen Tchetgen. Identifying causal effects with proxy variables of an unmeasured confounder. *Biometrika*, 105(4):987–993, 2018.
- [57] Robert T Michae and Michael R Pergamit. The National Longitudinal Survey of Youth, 1997 Cohort. *Journal of Human Resources*, 36(4):628–640, 2001.
- [58] Whitney Moore, Steven Pedlow, Parvati Krishnamurty, Kirk Wolter, and IL Chicago. National longitudinal survey of youth 1997 (nlsy97). *National Opinion Research Center, Chicago, IL*, 254, 2000.
- [59] Donald F Moores. One size does not fit all: Individualized instruction in a standardized educational system. *American Annals of the Deaf*, 158(1):98–103, 2013.
- [60] Stephen L. Morgan and Christopher Winship. *Causal Graphs*, page 77–102. Analytical Methods for Social Research. Cambridge University Press, 2014.
- [61] Susan A Murphy. Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 65(2):331–355, 2003.
- [62] David J Newell. Intention-to-treat analysis: implications for quantitative and qualitative research. *International journal of epidemiology*, 21(5):837–841, 1992.
- [63] Elizabeth L Ogburn and Tyler J VanderWeele. Causal diagrams for interference. 2014.
- [64] Judea Pearl. Causal inference in statistics: An overview. *Statistics Surveys*, 3:96–146, 2009.
- [65] Judea Pearl. *Causality*. Cambridge university press, 2009.

- [66] Tim J Peters, S H Richards, C R Bankhead, A E Ades, and J A C Sterne. Comparison of methods for analysing cluster randomized trials: An example involving a factorial design. *International Journal of Epidemiology*, 32:840–846, 2003.
- [67] Maya Petersen, Joshua Schwab, Susan Gruber, Nello Blaser, Michael Schomaker, and Mark van der Laan. Targeted maximum likelihood estimation for dynamic and static longitudinal marginal structural working models. *Journal of causal inference*, 2(2):147–185, 2014.
- [68] Maya L Petersen, Kristin E Porter, Susan Gruber, Yue Wang, and Mark J van der Laan. Diagnosing and responding to violations in the positivity assumption. *Statistical Methods in Medical Research*, 21(1):31–54, 2012.
- [69] Núria Porta, Catalina Bonet, and Erik Cobo. Discordance between reported intention-to-treat and per protocol analyses. *Journal of clinical epidemiology*, 60(7):663–669, 2007.
- [70] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2017.
- [71] Sarah A Reifeis and Michael G Hudgens. On variance of the treatment effect in the treated when estimated by inverse probability weighting. *American Journal of Epidemiology*, 191(6):1092–1097, 2022.
- [72] Thomas S Richardson and James M Robins. Single world intervention graphs (swigs): A unification of the counterfactual and graphical approaches to causality. *Center for the Statistics and the Social Sciences, University of Washington Series. Working Paper*, 128(30):2013, 2013.
- [73] James Robins. A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect. *Mathematical modelling*, 7(9-12):1393–1512, 1986.
- [74] James Robins and Miguel Hernan. Estimation of the causal effects of time-varying exposures. *Chapman & Hall/CRC Handbooks of Modern Statistical Methods*, pages 553–599, 2008.
- [75] James M Robins. Addendum to “a new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect”. *Computers & Mathematics with Applications*, 14(9-12):923–945, 1987.

- [76] James M Robins. Association, causation, and marginal structural models. *Synthese*, 121(1/2):151–179, 1999.
- [77] James M Robins. Robust estimation in sequentially ignorable missing data and causal inference models. In *Proceedings of the American Statistical Association*, volume 1999, pages 6–10. Indianapolis, IN, 2000.
- [78] James M Robins and Sander Greenland. Identifiability and exchangeability for direct and indirect effects. *Epidemiology*, 3(2):143–155, 1992.
- [79] James M Robins, Sander Greenland, and Fu-Chang Hu. Estimation of the causal effect of a time-varying exposure on the marginal mean of a repeated binary outcome. *Journal of the American Statistical Association*, 94:687–700, 1999.
- [80] James M Robins, Miguel Ángel Hernán, and Babette A Brumback. Marginal structural models and causal inference in epidemiology. *Epidemiology*, 11:550–560, 2000.
- [81] James M. Robins, Mariela Sued, Quanhong Lei-Gomez, and Andrea Rotnitzky. Comment: Performance of double-robust estimators when “inverse probability” weights are highly variable. *Statistical Science*, 22:544–559, 2007.
- [82] Elizabeth T Rogawski McQuade, Jade Benjamin-Chung, Daniel Westreich, and Benjamin F Arnold. Population intervention effects in observational studies to emulate target trial results: reconciling the effects of improved sanitation on child growth. *International journal of epidemiology*, 51(1):279–290, 2022.
- [83] Paul R Rosenbaum and Paul R Rosenbaum. *Overt bias in observational studies*. Springer, 2002.
- [84] Paul R Rosenbaum and Donald B Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70:41–55, 1983.
- [85] Donald B Rubin. Estimating causal effects of treatment in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66:688–701, 1974.
- [86] Donald B Rubin. The design of a general and flexible system for handling nonresponse in sample surveys. *The American Statistician*, 58(4):298–302, 2004.

- [87] Donald B Rubin. Causal inference using potential outcomes: Design, modeling, decisions. *Journal of the American Statistical Association*, 100(469):322–331, 2005.
- [88] Donald B Rubin and Nathaniel Schenker. Multiple imputation in health-care databases: An overview and some applications. *Statistics in medicine*, 10(4):585–598, 1991.
- [89] Melanie D Sabado, Denise Haynie, Stephen E Gilman, Bruce Simons-Morton, and Kelvin Choi. High school cigarette smoking and post-secondary education enrollment: Longitudinal findings from the NEXT generation health study. *Preventive Medicine*, 105:250–256, 2017.
- [90] Anja Saso and Beate Kampmann. Vaccine responses in newborns. In *Seminars in immunopathology*, volume 39, pages 627–642. Springer, 2017.
- [91] Bernhard Schölkopf and Alexander J Smola. *Learning with kernels: support vector machines, regularization, optimization, and beyond*. MIT press, 2002.
- [92] Sharon Schwartz, Nicolle M Gatto, and Ulka B Campbell. Extending the sufficient component cause model to describe the stable unit treatment value assumption (SUTVA). *Epidemiologic Perspectives and Innovations*, 9:1–11, 2012.
- [93] Valerie A Smith, Cynthia J Coffman, and Michael G Hudgens. Interpreting the results of intention-to-treat, per-protocol, and as-treated analyses of clinical trials. *Jama*, 326(5):433–434, 2021.
- [94] Mats J Stensrud, James M Robins, Aaron Sarvet, Eric J Tchetgen Tchetgen, and Jessica G Young. Conditional separable effects. *Journal of the American Statistical Association*, 118(544):2671–2683, 2023.
- [95] Jonathan AC Sterne and Kate Tilling. G-estimation of causal effects, allowing for time-varying confounding. *The Stata Journal*, 2(2):164–182, 2002.
- [96] Chien-Lin Su, Russell Steele, and Ian Shrier. Doubly robust estimation and causal inference for recurrent event data. *Statistics in Medicine*, 39:2324–2338, 2020.
- [97] Eric J Tchetgen Tchetgen, M Maria Glymour, Jennifer Weuve, and James M Robins. Specifying the correlation structure in inverse-probability-weighting estimation for repeated measures. *Epidemiology*, 23:644–646, 2012.

- [98] Steven Tenn, Douglas A Herman, and Brett Wendling. The role of education in the production of health: An empirical analysis of smoking behavior. *Journal of Health Economics*, 29(3):404–417, 2010.
- [99] Linh Tran, Constantin Yiannoutsos, Kara Wools-Kaloustian, Abraham Siika, Mark van der Laan, and Maya Petersen. Double Robust Efficient Estimators of Longitudinal Treatment Effects: Comparative Performance in Simulations and a Case Study. *The International Journal of Biostatistics*, 15:20170054, 2019.
- [100] Anastasios A Tsiatis. Semiparametric theory and missing data. 2006.
- [101] Mark J Van der Laan, Sherri Rose, et al. *Targeted learning: causal inference for observational and experimental data*, volume 4. Springer, 2011.
- [102] Mark J Van Der Laan and Daniel Rubin. Targeted maximum likelihood learning. *The international journal of biostatistics*, 2(1), 2006.
- [103] Machteld Vandecandelaere, Stijn Vansteelandt, Bieke De Fraine, and Jan Van Damme. Time-varying treatments in observational studies: Marginal structural models of the effects of early grade retention on math achievement. *Multivariate Behavioral Research*, 51:843–864, 2016.
- [104] Tyler J Vanderweele, Louise C Hawkey, Ronald A Thisted, and John T Cacioppo. A marginal structural model analysis for loneliness: Implications for intervention trials and clinical practice. *Journal of Consulting and Clinical Psychology*, 79:225–235, 2011.
- [105] Kerollos Nashat Wanis, Aaron L Sarvet, Lan Wen, Jason P Block, Sheryl L Rifas-Shiman, James M Robins, and Jessica G Young. Grace periods in comparative effectiveness studies of sustained treatments. *Journal of the Royal Statistical Society Series A: Statistics in Society*, page qnae002, 2024.
- [106] Lan Wen, Miguel A. Hernán, and James M. Robins. Multiply robust estimators of causal effects for survival outcomes. *Scandinavian Journal of Statistics*, 49:1304–1328, 2022.
- [107] Lan Wen, Jessica G Young, James M Robins, and Miguel A Hernán. Parametric g-formula implementations for causal survival analyses. *Biometrics*, 77(2):740–753, 2021.

- [108] Lawrence PO Were, Edwin Were, Richard Wamai, Joseph Hogan, and Omar Galarraga. The association of health insurance with institutional delivery and access to skilled birth attendants: evidence from the kenya demographic and health survey 2008–09. *BMC health services research*, 17(1):1–10, 2017.
- [109] Daniel Westreich. From patients to policy: population intervention effects in epidemiology. *Epidemiology*, 28(4):525–528, 2017.
- [110] Michelle M Wiest, Katherine J Lee, and John B Carlin. Statistics for clinicians: An introduction to logistic regression. *Journal of Paediatrics and Child Health*, 51(7):670–673, 2015.
- [111] Jackson T Wright Jr, Paul K Whelton, and David M Reboussin. A randomized trial of intensive versus standard blood-pressure control. *The New England journal of medicine*, 374(23):2294–2294, 2016.
- [112] Jessica G Young, Lauren E Cain, James M Robins, Eilis J O’Reilly, and Miguel A Hernán. Comparative effectiveness of dynamic treatment regimes: an application of the parametric g-formula. *Statistics in biosciences*, 3:119–143, 2011.
- [113] Jessica G Young, Roger W Logan, James M Robins, and Miguel A Hernán. Inverse probability weighted estimation of risk under representative interventions in observational studies. *Journal of the American Statistical Association*, 2018.
- [114] Jessica G. Young, Mats J. Stensrud, Eric J. Tchetgen Tchetgen, and Miguel A. Hernán. A causal framework for classical statistical estimands in failure-time settings with competing events. *Statistics in Medicine*, 39:1199–1236, 2020.
- [115] Jessica G Young, Rajet Vatsa, Eleanor J Murray, and Miguel A Hernán. Interval-cohort designs and bias in the estimation of per-protocol effects: a simulation study. *Trials*, 20:1–9, 2019.
- [116] Wenling Zhang and Cecilia A Cotton. Causal inference for recurrent events via aggregated marginal odds ratio. *Statistics in Medicine*, 42(18):3208–3235, 2023.
- [117] Zhiwei Zhang. Estimating a marginal causal odds ratio subject to confounding. *Communications in Statistics - Theory and Methods*, 38:309–321, 2009.

- [118] Yaqian Zhu, Rebecca A. Hubbard, Jessica Chubak, Jason Roy, and Nandita Mitra. Core concepts in pharmacoepidemiology: Violations of the positivity assumption in the causal analysis of observational data: Consequences and statistical approaches. *Pharmacoepidemiology and Drug Safety*, 30(11):1471–1485, 2021.
- [119] Hui Zou. The adaptive lasso and its oracle properties. *Journal of the American statistical association*, 101(476):1418–1429, 2006.

APPENDICES

Appendix A

Weight Estimation and Detailed Simulation Process for Chapter 2

In this appendix, we discuss the intuition for the proposed weight models, and give the numerical details for the simulations of time-fixed setting, the partially time-varying setting, and the time-varying settings specified in Section 2.3. Step-by-step instructions for the simulation process are presented, including the specifications for the covariate coefficients.

As specified in Section 2.2.4 and in this appendix, all the treatment models in the weight numerators and denominators all depends on the lag-1 covariate history for simplicity. This is based on the assumption that the impact of covariate history on treatment $A(t)$ can be fully captured by the the outcome at the previous time point $Y(t - 1)$ in the data.

A.1 Simulation Details for Section 2.3.1

Recall the relationship implied by Figure 2.1, let $x_1, x_2, \dots, x_9$ be the baseline covariates, where $\mathbf{x}_{\text{treat}} = (x_7, x_8)$, $\mathbf{x}_{\text{cfd}} = (x_1, x_2, x_4, x_5)$, and $\mathbf{x}_{\text{out}} = (x_3, x_6)$. These baseline covariates are generated from nine independent Bernoulli distributions, each with parameter $p = 0.5$. Next, the treatment and outcome statuses of all N subjects are generated by $A \sim \text{Bernoulli}(p_{\text{treat}})$ and $Y \sim \text{Bernoulli}(p_{\text{out}})$, where

$$\begin{aligned}\text{logit}(p_{\text{treat}}) &= \text{logit}P(A = 1|x_1, x_2, x_4, x_5, x_7, x_8) \\ &= \alpha_0 + \alpha_1 x_1 + \alpha_2 x_2 + \alpha_3 x_4 + \alpha_4 x_5 + \alpha_5 x_7 + \alpha_6 x_8,\end{aligned}$$

and

$$\begin{aligned}\text{logit}(p_{\text{out}}) &= \text{logit}P(Y = 1|x_1, x_2, x_3, x_4, x_5, x_6, A) \\ &= \beta_0 + \beta_{\text{trt}}A + \beta_1x_1 + \beta_2x_2 + \beta_3x_3 + \beta_4x_4 + \beta_5x_5 + \beta_6x_6.\end{aligned}$$

Define $\beta_1, \beta_3, \beta_5, \alpha_1, \alpha_4, \alpha_3 = \log(5)$, and $\beta_2, \beta_4, \beta_6, \alpha_2, \alpha_5, \alpha_6 = \log(2)$ to reflect the strong and weak association of the corresponding baseline covariates. The parameter β_0 is set to -3.5 and α_0 is set to -5.0. In this setting, approximately half of the subjects are exposed to the treatment, half of the treated subjects experience the outcome while less than 20% of the untreated subjects experience the outcome.

A.2 Weight Estimation for Section 2.3.1

In correspondence with the weight models discussed in Section 2.2.4, the following three IPW models are used to estimate the marginal treatment effects. Note that the history stabilized model and the mixed stabilized model are eliminated because they coincide with the semi-stabilized model and the baseline stabilized model in the time-fixed scenario.

Weight Model 0: For each of the N subjects, the unstabilized weight is calculated by

$$\text{UW}_{\text{tf}} = \frac{1}{P(A = a|x_1, x_2, x_4, x_5, x_7, x_8)}.$$

Weight Model 1: The semi-stabilized model specifies the numerator to be the marginal probability of treatment. For each subject, the semi-stabilized weight is calculated by

$$\text{SW1}_{\text{tf}} = \frac{P(A = a)}{P(A = a|x_1, x_2, x_4, x_5, x_7, x_8)}.$$

Weight Model 2: The pure stabilized model has $\mathbf{x}_{\text{treat}} = (x_7, x_8)$ as the condition of the numerator, so the weight for each subject is

$$\text{SW2}_{\text{tf}} = \frac{P(A = a|x_7, x_8)}{P(A = a|x_1, x_2, x_4, x_5, x_7, x_8)}.$$

A.3 Simulation Details for Section 2.3.2

Recall the relationship implied by Figure 2.2, let $x_1, x_2, \dots, x_6$ be the baseline covariates, where $\mathbf{x}_{\text{treat}} = (x_5, x_6)$, $\mathbf{x}_{\text{cfd}} = (x_1, x_2)$, and $\mathbf{x}_{\text{out}} = (x_3, x_4)$. These baseline covariates are generated from six independent Bernoulli distributions, each with $p = 0.5$. Next, the two-course treatment and outcome statuses for all N subjects are generated by $A(1) \sim \text{Bernoulli}(\pi(1)_{\text{treat}})$, $A(2) \sim \text{Bernoulli}(\pi(2)_{\text{treat}})$, and $Y \sim \text{Bernoulli}(p_{\text{out}})$, where

$$\begin{aligned}\text{logit}(\pi(1)_{\text{treat}}) &= \text{logit}P(A(1) = 1|x_1, x_2, x_5, x_6) \\ &= \alpha_{01} + \alpha_{11}x_1 + \alpha_{21}x_2 + \alpha_{31}x_5 + \alpha_{41}x_6,\end{aligned}$$

$$\begin{aligned}\text{logit}(\pi(2)_{\text{treat}}) &= \text{logit}P(A(2) = 1|A(1), x_1, x_2, x_5, x_6) \\ &= \alpha_{02} + \alpha_{12}x_1 + \alpha_{22}x_2 + \alpha_{32}x_5 + \alpha_{42}x_6 + \alpha_{A1}A(1),\end{aligned}$$

and

$$\begin{aligned}\text{logit}(p_{\text{out}}) &= \text{logit}P(Y = 1|A(2), x_1, x_2, x_3, x_4) \\ &= \beta_0 + \beta_{\text{trt}}A(2) + \beta_1x_1 + \beta_2x_2 + \beta_3x_3 + \beta_4x_4,\end{aligned}$$

where $\beta_1, \beta_3, \alpha_{11}, \alpha_{31}, \alpha_{12}, \alpha_{32} = \log(5)$, and $\beta_2, \beta_4, \alpha_{21}, \alpha_{41}, \alpha_{22}, \alpha_{42} = \log(2)$, which reflect a strong association and a weak association of the corresponding baseline covariates, respectively. Note that the effect of treatment on the outcome is only directly through $A(2)$. The parameter β_0 is set to -3.7 so that approximately half of the subjects are exposed to the treatment. Similarly, α_{01} is set at -2.0 and α_{02} is set at -2.3. Under this setting, more than half of the subjects are exposed to the treatment, more than half of the treated subjects experience the outcome, while about 20% of the untreated subjects experience the outcome.

A.4 Weight Estimation for Section 2.3.2

In accordance with Section 2.2.4, The following five IPW models are used to estimate the marginal treatment effects.

Weight Model 0: For each subject, the unstabilized weight is calculated by

$$UW_v = \frac{1}{P(A(1) = a(1)|x_1, x_2, x_5, x_6)} \times \frac{1}{P(A(2) = a(2)|A(1) = a(1), x_1, x_2, x_5, x_6)}.$$

Weight Model 1: For each subject, the semi-stabilized weight is calculated by

$$SW1_v = \frac{P(A(1) = a(1))}{P(A(1) = a(1)|x_1, x_2, x_5, x_6)} \times \frac{P(A(2) = a(2))}{P(A(2) = a(2)|A(1) = a(1), x_1, x_2, x_5, x_6)}.$$

Weight Model 2: For each subject, the stabilized weight is calculated by

$$SW2_v = \frac{P(A(1) = a(1)|x_5, x_6)}{P(A(1) = a(1)|x_1, x_2, x_5, x_6)} \times \frac{P(A(2) = a(2)|x_5, x_6)}{P(A(2) = a(2)|A(1) = a(1), x_1, x_2, x_5, x_6)}.$$

Weight Model 3: The history stabilized model has the treatment history $a(1)$ as a condition in the numerator at the second time point only. For each subject, the weight is calculated by

$$SW3_v = \frac{P(A(1) = a(1))}{P(A(1) = a(1)|x_1, x_2, x_5, x_6)} \times \frac{P(A(2) = a(2)|a(1))}{P(A(2) = a(2)|A(1) = a(1), x_1, x_2, x_5, x_6)}.$$

Weight Model 4: The mixed stabilized model has both $a(1)$ and $\mathbf{x}_{\text{treat}} = (x_5, x_6)$ as a condition of the treatment probability in the numerator at the second time point. For each subject, the weight is calculated by

$$SW4_v = \frac{P(A(1) = a(1)|x_5, x_6)}{P(A(1) = a(1)|x_1, x_2, x_5, x_6)} \times \frac{P(A(2) = a(2)|a(1), x_5, x_6)}{P(A(2) = a(2)|A(1) = a(1), x_1, x_2, x_5, x_6)}.$$

Note that to prevent extreme weights from affecting the estimation process, any weights greater than 20 are truncated to be strictly 20. Since extreme weights are rare (less than 1%) for time-fixed treatments and outcomes (Section 2.3.1, truncation applies only to the weight models of partially time-varying and time-varying cases specified in Sections 2.3.2 and 2.3.3. Please refer to the supplementary material for more details regarding the weight distributions and truncation.

A.5 Simulation Details for Section 2.3.3

When treatment and outcome vary over time, both time-fixed and time-varying covariates should be considered. We followed and simplified the covariate settings defined in the time-fixed case (Section 2.2.1) as our baseline covariates. For simplicity, we consider the recent outcome history ($Y(t-1)$) to be a time-varying confounder in Section 2.3.3. Note that any time-varying covariates will be treated with a structure similar to that of $Y(t-1)$. Recall the relationship implied by Figure 2.3, let x_1 , x_2 , x_3 , and x_4 be the baseline covariates, where $\mathbf{x}_{\text{treat}} = x_3$, $\mathbf{x}_{\text{cfd}} = (x_1, x_2)$, and $\mathbf{x}_{\text{out}} = x_4$. These baseline covariates are generated from discrete and continuous distributions to resemble medical data in real life. In particular, let

$$A(0) \sim \text{Bernoulli}(0.5)$$

$$Y(0) \sim \text{Bernoulli}(0.5),$$

and

$$x_1 \sim \text{Bernoulli}(0.5); x_2 \sim \text{Normal}(0, 1); x_3 \sim \text{Normal}(0, 0.7); x_4 \sim \text{Normal}(0, 0.9).$$

For $t = 1, \dots, T$, the treatment and outcome statuses for N subjects are generated by

$$A(t) \sim \text{Bernoulli}(\pi(t)_{\text{treat}}),$$

where

$$\begin{aligned} \text{logit}(\pi(t)_{\text{treat}}) &= \text{logit}P(A(1) = 1 | x_1, x_2, x_3, A(t-1), Y(t-1)) \\ &= \alpha_0 + \alpha_1 x_1 + \alpha_2 x_2 + \alpha_3 A(t-1) + \alpha_4 Y(t-1) + \alpha_5 x_3, \end{aligned} \quad (\text{A.1})$$

and

$$Y(t) \sim \text{Bernoulli}(\pi_{\text{out}}),$$

where

$$\begin{aligned} \text{logit}(\pi(t)_{\text{out}}) &= \text{logit}P(Y = 1 | A(t), Y(t-1), x_1, x_2, x_4) \\ &= \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \beta_{\text{trt}} A(t) + \beta_4 Y(t-1) + \beta_5 x_4. \end{aligned} \quad (\text{A.2})$$

The next step is to determine the covariate coefficients. We first discuss the coefficient choices in Section 2.3.3. For the original non-absorbing treatment case, the non-intercept covariate coefficients are set to have a moderate effect on the treatment

and outcome.

Specifically, let $\alpha_1, \dots, \alpha_5, \beta_1, \dots, \beta_6$ (except for β_{trt}) have a magnitude between 0 and 1 as follows:

$$\begin{aligned}\alpha_1 &= 0.9, \alpha_2 = 0.7, \alpha_3 = 0.5, \alpha_4 = 0.9, \alpha_5 = 0.6 \\ \beta_1 &= -0.9, \beta_2 = 0.7, \beta_4 = 0.75, \beta_5 = 0.6.\end{aligned}$$

The intercept of the treatment model, α_0 , needs to vary with different values of T and AMOR to ensure that approximately 50% of the subjects-time are treated. For each value of T , a grid search is performed to find the optimal α_0 value that generates a data set with all specifications satisfied. The resulting α_0 values are shown in the first column of Table A.1.

Recalling from Section 2.2.3, the empirical method draws a connection between the desired values of AMOR (1, 2, 4, and 6) and the conditional treatment effect, which is represented by β_{treat} through Equation (2.11). To study the AMOR of values 1, 2, 4, and 6, the first step is to find the corresponding β_{treat} values by grid search. This method requires consideration of all possible treatment assignments. With a fixed study period and a nonabsorbing treatment, each subject has $D = 2^T$ different possible treatment paths. We demonstrate a grid search when $T = 2$ is implemented as follows.

The grid of possible values for β_{treat} is set to be $\beta_{\text{treat}} = \log \text{amor} \pm 1$, where “amor” denotes the value of the corresponding AMOR. For example, for AMOR = 1, the corresponding grid search range for β_{treat} should be $\log(1) \pm 0.5$. Considering the underlying relationship between conditional and marginal treatment effects, this range is expected to cover the true AMOR. For each desired AMOR value, there are 1001 values in the grid search, with a grid precision of 0.001. Then, for each value in the grid, a data set with D treatment paths is randomly generated based on Equation (A.2), and an empirical AMOR is calculated by Equation (2.14). After confirming that the 1001 output empirical AMOR values cover the true AMOR value (i.e., there are empirical AMORs greater and smaller than the desired AMOR), we find the empirical AMOR that is closest to the desired AMOR and record the corresponding β_{treat} value. Alternatively, one could start with a grid with a precision of 0.1 and record both the smallest empirical AMOR that is larger than the desired AMOR and the largest empirical AMOR that is smaller than the desired AMOR. Then, the two recorded empirical AMORs are used as the new grid range for β_{treat} with a precision of 0.01. Next, we gradually reduce the width of the range of β_{treat} and increase the precision of the grid search until we reach the precision of 0.001 and record the

best β_{treat} value. With multiple rounds of grid search, one could achieve the same precision but save more computational power.

The process of finding the corresponding β_{treat} is repeated for $T = 4, 6$, and 10 , and the results are shown in Table A.1. Then, the intercept β_0 is specified as 1.1 to ensure that overall approximately 70% to 82% of the treated subject experience the outcome at least once. The frequency of outcome occurrence under this parameter setting for $T = 4$ and $AMOR = 1$ is as follows: approximately 4.7% of the subjects never experience the outcome, approximately 10.2% of the subjects experience the outcome once, approximately 17.8% of the subjects experience the outcome twice, and about 28.4% of the subjects experience the outcome three times, and about 38.9% of the subjects experience the outcome four times.

The simulation process for absorbing treatments is similar to that for nonabsorbing treatments specified in Section 2.3.3. The methods of choosing the appropriate covariate coefficients are similar for absorbing and nonabsorbing treatments. The main difference is that there are $D = T + 1$ possible treatment paths since the treatment status can only change once over the entire study period. Following the steps from the last section, the grid search returns Table A.2 of β_{treat} in the outcome model given by Equation A.1 for the absorbing treatment scenario. And the α_0 values are shown with respect to T values in the first column of Table A.2. The results are derived again through a grid search of optimal α_0 values that ensures that the specifications in Section 2.3.3 are met and that approximately half of the population are ever treated. The other covariate coefficients, $\alpha_1, \dots, \alpha_5, \beta_0, \beta_1, \dots, \beta_6$ (except for β_{treat}) are set the same as above.

Note that the two settings of the nonabsorbing treatment scenario described in Sections 2.3.3 only differ in the value of $A(0)$. In particular, the original nonabsorbing treatment setting specifies that $A(0) \sim \text{Bernoulli}(0.5)$, while $A(0)$ is set to be zero in the comparative nonabsorbing treatment setting. Since $A(0)$ does not have a direct impact on the outcome model given by Equation (2.11), the coefficients for the outcome model are the same for the two settings, respectively. Recall that the treatment model given by Equation (2.10) depends on $A(t - 1)$. Therefore, the only difference in terms of parameter setting between them are the different α_0 values specified in Table A.1 because they differ in $A(0)$.

A.6 Weight Estimation for Section 2.3.3

For the two non-absorbing treatment scenarios, the following five IPW models are used to estimate the AMOR.

Weight Model 0:

$$\text{UW}(s) = \prod_{t=1}^s \frac{1}{P(A(t) = a(t)|x_1, x_2, x_3, a(t-1), y(t-1))}$$

Weight Model 1:

$$\text{SW1}(s) = \prod_{t=1}^s \frac{P(A(t) = a(t))}{P(A(t) = a(t)|x_1, x_2, x_3, a(t-1), y(t-1))}$$

Weight Model 2:

$$\text{SW2}(s) = \prod_{t=1}^s \frac{P(A(t) = a(t)|x_3)}{P(A(t) = a(t)|x_1, x_2, x_3, a(t-1), y(t-1))}$$

Weight Model 3:

$$\text{SW3}(s) = \prod_{t=1}^s \frac{P(A(t) = a(t)|a(t-1))}{P(A(t) = a(t)|x_1, x_2, x_3, a(t-1), y(t-1))}$$

Weight Model 4 :

$$\text{SW4}(s) = \prod_{t=1}^s \frac{P(A(t) = a(t)|x_3, a(t-1))}{P(A(t) = a(t)|x_1, x_2, x_3, a(t-1), y(t-1))}$$

For the absorbing treatment scenario, the following IPW models are used to estimate the AMOR.

Weight Model 0:

$$\text{UW}(s) = \prod_{t=1}^s \frac{1}{P(A(t) = a(t)|x_1, x_2, x_3, y(t-1))}$$

Weight Model 1:

$$\text{SW1}(s) = \prod_{t=1}^s \frac{P(A(t) = a(t))}{P(A(t) = a(t)|x_1, x_2, x_3, y(t-1))}$$

Weight Model 2:

$$\text{SW2}(s) = \prod_{t=1}^s \frac{P(A(t) = a(t)|x_3)}{P(A(t) = a(t)|x_1, x_2, x_3, y(t-1))}$$

	$\alpha_{0,1}$	$\alpha_{0,2}$	β_{trt} (1)	β_{trt} (2)	β_{trt} (4)	β_{trt} (6)
T=2	-1.25	-1.1	0.001	0.818	1.603	2.06
T=4	-1.30	-1.25	-0.001	0.816	1.593	2.036
T=6	-1.30	-1.30	-0.001	0.818	1.584	2.029
T=10	-1.40	-1.30	-0.001	0.813	1.588	2.038

Table A.1: Table of $\alpha_{0,1}$ (for the original setting), $\alpha_{0,2}$ (for the comparative setting) and β_{trt} (in terms of AMOR) Values for the Time-varying Nonabsorbing Treatment settings

	a_0	β_{trt} (1)	β_{trt} (2)	β_{trt} (4)	β_{trt} (6)
T=2	-1.45	-0.046	0.768	1.548	1.990
T=4	-2.09	-0.047	0.717	1.462	1.889
T=6	-2.48	-0.041	0.752	1.516	1.950
T=10	-3.00	-0.029	0.758	1.517	1.950

Table A.2: Table of α_0 and β_{trt} (in terms of AMOR) values for the Time-varying Absorbing treatment settings

Settings	Treatment(s) ^a	Outcome(s) ^b	Initial States
2.3.1 ^c	A	Y	-
2.3.2 ^d	nonabsorbing $A(t)$	Y	-
2.3.3 ^e	nonabsorbing $A(t)$	nonabsorbing $Y(t)$	$A(0) \sim \text{Bin}(0.5), Y(0) \sim \text{Bin}(0.5)$
2.3.3 ^f	Absorbing $A(t)$	nonabsorbing $Y(t)$	$A(0) = 0, Y(0) \sim \text{Bin}(0.5)$
2.3.3 ^g	nonabsorbing $A(t)$	nonabsorbing $Y(t)$	$A(0) = 0, Y(0) \sim \text{Bin}(0.5)$

Table A.3: Table comparing the five settings specified in Sections 2.3.1, 2.3.2, and 2.3.3.

^aThe treatment can be time-fixed or time-varying within the study period($t = 1, \dots, T$).

^bThe outcome can be time-fixed or time-varying within the study period($t = 1, \dots, T$).

^cThe time-fixed design

^dThe partially time-varying design

^eThe time-varying design with nonabsorbing treatments

^fThe time-varying design with absorbing treatments

^gThe time-varying design with nonabsorbing treatments and comparative setting

Appendix B

Supplementary Information for Chapter 2

B.1 Definition and Identification of Causal Estimands

With the defined

$$\text{AMOR} = \frac{P(\bar{Y}(T)^{(\bar{A}(T)=\bar{1})} = \bar{1})}{1 - P(\bar{Y}(T)^{(\bar{A}(T)=\bar{1})} = \bar{1})} \bigg/ \frac{P(\bar{Y}(T)^{(\bar{A}(T)=\bar{0})} = \bar{1})}{1 - P(\bar{Y}(T)^{(\bar{A}(T)=\bar{0})} = \bar{1})},$$

we introduce an intermediate causal estimand of $P(\bar{Y}(T)^{(\bar{A}(T)=\bar{a}(T))} = \bar{1})$, denoted by $\psi(\bar{a}(T))$, $\bar{a}(T) \in \{\bar{0}, \bar{1}\}$. To identify our causal estimand, we will make the following assumptions:

1. Exchangeability: The potential outcome is conditionally unconfounded with the series of potential treatments so far. Specifically, for any $t \in \{1, \dots, T\}$, assume:

$$Y(t)^{a(t)} \perp A(t)^{a(t-1)}, A(t-1)^{a(t-2)}, \dots, A(1) \mid X_{\text{out}}, X_{\text{cfd}}, Y(t-1)^{a(t-1)}, \dots, Y(1)^{a(1)}. \quad (\text{B.1})$$

2. Positivity: Each individual has a positive probability of receiving the treatments. In trial setting, $P(\bar{A}(T) = \bar{a}(T)) > 0$ for any treatment paths $\bar{a}(T)$; in

observational study setting, $P(\bar{A}(T) = \bar{a}(T) \mid x_{\text{trt}}, x_{\text{cfd}}, \bar{a}(T-1), \bar{y}(T-1)) > 0$ for any combinations of $\bar{a}(T)$, $\bar{y}(T)$, x_{trt} and x_{cfd} .

3. Consistency: The treatments are well-defined such that, for any $t \in \{1, \dots, T\}$, if $\bar{A}(t) = \bar{a}(t)$, then $Y(t) = Y(t)^{\bar{A}(t)}$. Additionally, we assume a lag-1 dependence structure in the covariates, so: $Y(t)^{\bar{a}(t)} = Y(t)^{a(t)}$.

Using the aforementioned causal assumptions and induction, we can identify $\psi(\bar{a}(T))$ and represent the identifying formula (or g-formula) for the expected outcomes with static treatment $\bar{a}(T)$ in an observational study setting as a series of expectations inversely weighted by the probabilities of treatments, i.e.

$$\begin{aligned} \psi(\bar{a}(T)) &= \frac{\mathbb{I}(\bar{A}(T) = \bar{a}(T))}{\bar{\pi}(a(T))} P(\bar{Y}(T) = \bar{1} \mid \bar{a}(T), \bar{y}(T-1), \mathbf{x}_{\text{treat}}, \mathbf{x}_{\text{cfd}}) \\ &= \frac{\mathbb{I}(\bar{A}(T) = \bar{a}(T))}{\bar{\pi}(a(T))} E(\bar{Y}(T) \mid \bar{a}(T), \bar{y}(T-1), \mathbf{x}_{\text{treat}}, \mathbf{x}_{\text{cfd}}), \end{aligned} \tag{B.2}$$

where $\bar{\pi}(a(T)) = \prod_{t=1}^T P(A(t) = a(t) \mid a(t-1), y(t-1), \mathbf{x}_{\text{treat}}, \mathbf{x}_{\text{cfd}})$.


Given no known violation to these assumptions, our simulations are based on variable relationships of realistic settings.

B.2 Weight Distributions for Simulation Studies


To demonstrate the estimated weight distributions for different weight models, we randomly generate a sample with $T = 4$ for each of the scenarios specified in Section 2.3.3, respectively and produce the following box plots. Note that all weights above 20 have been truncated to 20 for the extreme weights elimination, and the number of weights truncated are reported in the title of each plot.

As shown in the box plots, weight distribution for the absorbing treatment case has noticeably less weights being truncated and smaller average weights compare to the other two nonabsorbing treatment cases, which have very similar distribution across different weight models and AMOR values. The contrast between the absorbing and nonabsorbing cases may be due to the nature of absorbing treatment setting, where all treated subjects do not further contribute to the cumulative weights after absorption.


Meanwhile, Weight Model 0 has the most extreme weight distribution with significantly larger average weights than other weight models across different settings



(a) Estimated Weight Distribution with 11454 weights truncated at 20 for non-absorbing treatment



(b) Estimated Weight Distribution with 73 weights truncated at 20 for absorbing treatment



(c) Estimated Weight Distribution with 11663 weights truncated at 20 for non-absorbing treatment with comparative setting

Figure B.1: Estimated Weight Distributions when $T = 4$ Specified in Section 2.3.3

and specifications. This further indicates that the unstabilized IPW model is not appropriate for time-varying weight estimation.

B.3 Additional Results for Simulation Studies

Tables B.1 to B.3 show the estimation results for Section 2.3.3 respectively with $T = 10$.

AMOR ^a	Weight Model	Bias ^b	ASE ^c	ESE ^d
1	0	0.3609	0.0178	0.017
1	1	0.0324	0.024	0.0226
1	2	0.027	0.0235	0.0219
1	3	0.0444	0.0251	0.0245
1	4	0.0367	0.0244	0.0238
2	0	0.677	0.02	0.0198
2	1	0.0145	0.0282	0.0267
2	2	0.0111	0.0275	0.0259
2	3	0.0842	0.029	0.0275
2	4	0.0717	0.0283	0.0267
4	0	1.2853	0.0241	0.025
4	1	-0.0279	0.0354	0.035
4	2	-0.0279	0.0345	0.0335
4	3	0.195	0.0355	0.0358
4	4	0.1683	0.0346	0.034
6	0	1.9689	0.0276	0.0279
6	1	-0.0035	0.041	0.0404
6	2	0.0016	0.0398	0.0393
6	3	0.379	0.0408	0.0397
6	4	0.3391	0.0395	0.0386

Table B.1: Simulation-Based Estimation of AMOR with Different Weight Models: Non-absorbing Time-varying Treatment, $T = 10$

^aEmpirically derived aggregated marginal odds ratio

^bBias for the estimated aggregated marginal odds ratio

^cAverage standard error of log AMOR

^dEmpirical standard error of log AMOR

AMOR ^a	Weight Model	Bias ^b	ASE ^c	ESE ^d
1	0	0.0762	0.0244	0.022
1	1	0.0193	0.0275	0.0256
1	2	-0.002	0.0261	0.024
2	0	0.2405	0.0268	0.023
2	1	0.1196	0.0318	0.0291
2	2	0.0688	0.0297	0.0274
4	0	0.6476	0.0309	0.0276
4	1	0.388	0.0383	0.0374
4	2	0.2713	0.0353	0.0344
6	0	1.0968	0.0342	0.0309
6	1	0.6951	0.0433	0.0405
6	2	0.5143	0.0395	0.0363

Table B.2: Simulation-Based Estimation of AMOR with Different Weight Models: Absorbing Treatment Case with $T = 10$

^aEmpirically derived aggregated marginal odds ratio

^bBias for the estimated aggregated marginal odds ratio

^cAverage standard error of log AMOR

^dEmpirical standard error of log AMOR

AMOR ^a	Weight Model	Bias ^b	ASE ^c	ESE ^d
1	0	0.362	0.0178	0.0179
1	1	0.0333	0.0241	0.0245
1	2	0.0276	0.0236	0.0232
1	3	0.0524	0.0249	0.0253
1	4	0.0436	0.0243	0.024
2	0	0.677	0.0199	0.0195
2	1	0.0186	0.0282	0.0278
2	2	0.0144	0.0276	0.0273
2	3	0.1065	0.0287	0.0281
2	4	0.0917	0.028	0.0276
4	0	1.2765	0.0238	0.0222
4	1	-0.0263	0.0351	0.0342
4	2	-0.0269	0.0341	0.0339
4	3	0.2399	0.035	0.0326
4	4	0.2102	0.0339	0.0324
6	0	1.9598	0.0274	0.0269
6	1	-0.01	0.0409	0.0394
6	2	-0.0017	0.0398	0.0381
6	3	0.4442	0.0403	0.0391
6	4	0.4036	0.0391	0.0376

Table B.3: Simulation-Based Estimation of AMOR with Different Weight Models: Non-absorbing Treatment with Comparative Setting, $T = 10$

^aEmpirically derived aggregated marginal odds ratio

^bBias for the estimated aggregated marginal odds ratio

^cAverage standard error of log AMOR

^dEmpirical standard error of log AMOR

B.4 Variable Information for the Real Data Example in Section 2.4

Please find the Table of related variables for NLSY97 in the next page.

Type	Variable Name	Description	Specifications	Reference Number (NLSY97)
Baseline ($t = 0$)	marijuana	respondent(R)'s marijuana usage	binary with levels 1: R has used marijuana at $t = 0$ 0: otherwise	R03589.00
	race	R's race or ethnicity	categorical with levels: 1: R is black 2: R is hispanic 3: R is non-hispanic mixed race 4: R is non-hispanic and non-black race	R14826.00
	gender	R's gender	binary with levels 1: R is female 0: R is male	R05363.00
	mother's highest degree	the highest degree R's residential mother(res mom) completed	categorical with levels: 1: res mom had education until grade school 2: res mom had education until college 3: res mom had education until graduate school	R13027.00
	father's highest degree	the highest degree R's biological father(bio dad) completed	categorical with levels: 1: bio dad had education until grade school 2: bio dad had education until college 3: bio dad had education until graduate school	R13024.00
	time-varying age	R's age as of the survey year (varies with time but recorded at baseline)	continuous ranging from 14 to 22	R05364.02
	mother's age at birth	the age of R's Biological Mom when R was born	continuous ranging from 12 to 52	R12002.00
	poverty ratio	the ratio of R's household income and poverty level survey at $t = 0$	continuous ranging from 0.00 to 16.27	R12049.00
	smoking status ($A(0)$)	R's previous smoking experience	binary with levels 1: R has smoked in his/her life at $t = 0$ 0: otherwise	R03579.00
	smoking status ($A(t)$)	R's smoking status since the date of the last interview (DLI) to survey date at time t	binary with levels 1: R has smoked since DLI 0: otherwise	R21894.00 ($t = 1$) R35085.00 ($t = 2$) R49066.00 ($t = 3$) R65341.00 ($t = 4$) S09216.00 ($t = 5$)
Time-Varying ($t = 1, \dots, 5$)	school enrollment status ($Y(t)$)	R's school enrollment or attainment status as of the survey date at time t	binary with levels 1: R is enrolled in a grade school or college, or has attained a college degree 0: otherwise	R12014.00($t = 0$) R25600.01 ($t = 1$) R38815.01 ($t = 2$) R54606.01 ($t = 3$) R72242.01 ($t = 4$) S15380.01 ($t = 5$)

Appendix C

Supplementary Information for Chapter 3

In this appendix, we show derivations for identification and the efficient influence function (EIF) for our defined causal estimand of interest. Then we will provide detailed estimation procedures for some of the estimators discussed in Chapter 3.

C.1 Identification

With a static point treatment a , we start by identifying the causal estimand of interest without censoring, denoted by $\psi_{I^0,k}$; then we incorporate censoring into our analysis, where the causal estimand is denoted as $\psi_{I,k}, k = 1, \dots, K + 1$.

C.1.1 Without Censoring

When censoring is absent, our relevant data can be expressed as

$$O = (L_0, A, Y_1, L_1, \dots, L_{K-1}, Y_K).$$

Figure C.1 shows the corresponding single world intervention graph (SWIG) that graphically illustrate the intervening process and variable relationships.

Since Y_k and L_k usually co-exist, denote $L_j^* = (Y_j, L_j), (j = 1, \dots, K + 1)$ for




Figure C.1: A single world intervention graph (SWIG) that identifies the average causal effect and illustrates the relationship between the measured baseline variables L_0 , the static treatment $A = a$, the time-varying variable $\bar{L}_k$ and the (intermediate) recurrent outcome $\bar{Y}_{k+1}$, assuming no unmeasured confounding.

abbreviation from this point on. Then, the relevant data can be abbreviated as

$$O = (L_0, A, L_1^*, \dots, L_{k-1}^*, Y_k).$$

In addition, we simplify the notation $p(b) = P(B = b)$, and $p(b|d) = P(B = b|D = d)$ for any random variables B and D .

Based on the causal assumptions discussed in Section 3.2 and probability rules, the causal quantity can be rewritten in the iterative conditional expectation (ICE) form as follows.

$$\begin{aligned}
E(Y_k^a) &= \sum_{\forall l_0} E(Y_k^a | l_0) p(l_0) \\
(\text{exchangeability}) &= \sum_{\forall l_0} E(Y_k^a | A = a, l_0) p(l_0) \\
(\text{probability rules}) &= \sum_{\forall \bar{l}_1^*} E(Y_k^a | \bar{l}_1^*, a) p(l_1 | a, l_0) p(l_0) \\
(\text{probability rules recursively}) &= \sum_{\forall \bar{l}_{k-1}^*} E(Y_k^a | \bar{l}_{k-1}^*, a) p(l_{k-1} | \bar{l}_{k-2}^*, a) \dots p(l_1 | l_0, a) p(l_0) \\
(\text{consistency}) &= E\{E[E(\dots E[E(Y_k | \bar{L}_{k-1}^*, A = a) | \bar{L}_{k-2}^*, A = a] \dots | \bar{L}_1^*, A = a) | L_0]\}
\end{aligned}$$

Therefore,

$$\begin{aligned}\psi_{I^0,k} &= E_P(R_k^a) = \sum_{j=1}^k E_P(Y_j^a) \\ &= \sum_{j=1}^k E\{E[E(\dots E[E(Y_j|\bar{L}_{j-1}^*, A=a)|\bar{L}_{j-2}^*, A=a] \dots |\bar{L}_1^*, A=a)|L_0]\}\end{aligned}$$

C.1.2 With Censoring

When censoring is taken into consideration, our relevant data becomes

$$O = (L_0, A, C_1, L_1^*, C_2, \dots, L_{k-1}^*, C_k, Y_k).$$

Then the ICE form of our causal estimand of interest, $\psi_{I,k} = E(Y_k^{a, \bar{c}_k=\bar{0}})$, can be derived as follows.

$$E(Y_k^{a, \bar{c}_k=\bar{0}}) = \sum_{\forall l_0} E(Y_k^{a, \bar{c}_k=\bar{0}}|l_0)p(l_0).$$

By the causal assumptions of exchangeability and consistency,

$$E(Y_k^{a, \bar{c}_k=\bar{0}}) = \sum_{\forall l_0} E(Y_k^{a, \bar{c}_k=\bar{0}}|A=a, l_0)p(l_0).$$

By exchangeability,

$$E(Y_k^{a, \bar{c}_k=\bar{0}}) = \sum_{\forall l_0} E(Y_k^{a, \bar{c}_k=\bar{0}}|C_1^a=0, A=a, l_0)p(l_0).$$

By consistency and probability rules,

$$E(Y_k^{a, \bar{c}_k=\bar{0}}) = \sum_{\forall \bar{l}_1^*} E(Y_k^{a, \bar{c}_k=\bar{0}}|C_1=0, \bar{l}_1^*, a)p(\bar{l}_1^*|C_1=0, a, l_0)p(l_0).$$

For $t \in \{1, \dots, k-1\}$, we have

$$\sum_{\forall \bar{l}_t^*} E(Y_k^{a, \bar{c}_k=\bar{0}}|\bar{C}_t=\bar{0}, \bar{l}_t^*, a)p(\bar{l}_t^*|\bar{C}_t=\bar{0}, \bar{l}_{t-1}^*, a) \dots p(\bar{l}_1^*|C_1=0, a, l_0)p(l_0)$$

(exchangeability)

$$\begin{aligned}
&= \sum_{\forall \bar{l}_t^*} E(Y_k^{a, \bar{c}_k = \bar{0}} | C_{t+1}^{a, \bar{c}_t = \bar{0}} = 0, \bar{C}_t = \bar{0}, \bar{l}_t^*, a) p(l_t^* | \bar{C}_t = \bar{0}, \bar{l}_{t-1}^*, a) \dots p(l_1^* | C_1 = 0, a, l_0) p(l_0) \\
&\text{(consistency)} \\
&= \sum_{\forall \bar{l}_t^*} E(Y_k^{a, \bar{c}_k = \bar{0}} | \bar{C}_{t+1} = \bar{0}, \bar{l}_t^*, a) p(l_t^* | \bar{C}_t = \bar{0}, \bar{l}_{t-1}^*, a) \dots p(l_1^* | C_1 = 0, a, l_0) p(l_0)
\end{aligned}$$

By induction, we argue that

$$\begin{aligned}
E(Y_k^{a, \bar{c}_k = \bar{0}}) &= \sum_{\forall \bar{l}_1^*} E(Y_k^{a, \bar{c}_k = \bar{0}} | C_1 = 0, a, l_0) p(l_0) \\
&\text{(recursively for } t = \{1, \dots, k-2\}) \\
&= \sum_{\forall \bar{l}_t^*} E(Y_k^{a, \bar{c}_k = \bar{0}} | \bar{C}_{t+1} = \bar{0}, \bar{l}_t^*, a) p(l_t^* | \bar{C}_t = \bar{0}, \bar{l}_{t-1}^*, a) \dots p(l_1^* | C_1 = 0, a, l_0) p(l_0) \\
&= \sum_{\forall \bar{l}_{k-1}^*} E(Y_k^{a, \bar{c}_k = \bar{0}} | \bar{C}_k = \bar{0}, \bar{l}_{k-1}^*, a) p(l_k^* | \bar{C}_{k-1} = \bar{0}, \bar{l}_{k-2}^*, a) \dots p(l_1^* | C_1 = 0, a, l_0) p(l_0) \\
&\text{(consistency)} \\
&= \sum_{\forall \bar{l}_{k-1}^*} E(Y_k | \bar{l}_{k-1}^*, a, \bar{C}_k = \bar{0}) p(l_{k-1}^* | \bar{l}_{k-2}^*, a, \bar{C}_{k-1} = \bar{0}) \dots p(l_1^* | l_0, a, C_1 = 0) p(l_0) \\
&= E\{E[E(\dots E[E(Y_k | \bar{l}_{k-1}^*, a, \bar{C}_k = \bar{0}) | \bar{l}_{k-2}^*, a, \bar{C}_{k-1} = \bar{0}] \dots | \bar{l}_1^*, a, C_1 = 0) | L_0]\}.
\end{aligned}$$

The corresponding estimand is:

$$\begin{aligned}
\psi_{I,k} &= E\{E[E(\dots E[E(Y_j | \bar{C}_j = \bar{0}, \bar{l}_{j-1}^*, A = a) | \bar{C}_{j-1} = \bar{0}, \bar{l}_{j-2}^*, A = a] \\
&\quad \dots | \bar{C}_2 = \bar{0}, \bar{l}_1^*, A = a) | L_0]\}.
\end{aligned}$$

Alternatively, $E(Y_k^{a, \bar{c}_k = \bar{0}})$ can be expressed by the inverse probability weighted form specified by Equation 3.4. The derivation process is demonstrated as follows.

$$\begin{aligned}
E(Y_k^{a, \bar{c}_k = \bar{0}}) &= \sum_{\forall l_0} E(Y_k^{a, \bar{c}_k = \bar{0}} | l_0) p(l_0) \\
&\text{(exchangeability)} \\
&= \sum_{\forall l_0} E(Y_k^{a, \bar{c}_k = \bar{0}} | C_1, A, l_0) p(l_0) \\
&\text{(probability rules)} \\
&= \frac{\mathbb{I}(A = a)(1 - C_1)}{\pi_A^a \bar{\pi}_{C_1}^0} \sum_{\forall l_0} E(Y_k^{a, \bar{c}_k = \bar{0}} | a, C_1 = 0, l_0) p(l_0) \\
&\text{(for } t = 1, \dots, k-1, \text{ recursively)} \\
&= \frac{\mathbb{I}(A = a) \prod_{h=1}^t (1 - C_h)}{\pi_A^a \bar{\pi}_{C_t}^0} \sum_{\forall \bar{l}_t^*} E(Y_k^{a, \bar{c}_k = \bar{0}} | a, \bar{C}_t = \bar{0}, \bar{l}_t^*) p(\bar{l}_t^*) \tag{C.1} \\
&\text{(exchangeability \& probability rules)} \\
&= \frac{\mathbb{I}(A = a) \prod_{h=1}^{t+1} (1 - C_h)}{\pi_A^a \bar{\pi}_{C_{t+1}}^0} \sum_{\forall \bar{l}_t^*} E(Y_k^{a, \bar{c}_k = \bar{0}} | a, \bar{C}_{t+1} = \bar{0}, \bar{l}_t^*) p(\bar{l}_t^*) \\
&\text{(by induction and } t = k - 1) \\
&= \frac{\mathbb{I}(A = a) \prod_{h=1}^k (1 - C_h)}{\pi_A^a \bar{\pi}_{C_k}^0} \sum_{\forall \bar{l}_t^*} E(Y_k^{a, \bar{c}_k = \bar{0}} | a, \bar{C}_k = \bar{0}, \bar{l}_{k-1}^*) p(\bar{l}_{k-1}^*) \\
&\text{(consistency)} \\
&= \frac{\mathbb{I}(A = a) * \prod_{h=1}^k (1 - C_h)}{\pi_A^a \bar{\pi}_{C_k}^0} E(Y_k | \bar{C}_k = \bar{0}, a, \bar{L}_{k-1}^*)
\end{aligned}$$

C.2 EIF Derivation for $\psi_{I,k}$

Next, we derive the efficient influence function (EIF) for the causal estimands based on their identification formulas given in Section C.1 of this Appendix. Let $p(O; e)$ be a regular parametric submodel of the nonparametric main model indexed by e , where its score function in $e = 0$ can be written as $S(O; e) = S(\bar{Y}_k, \bar{L}_{k-1}, A = a; e)$. Note that we suppress e in p for convenience of notation. The EIF for ψ in a nonparametric main model is defined as the unique mean zero, finite-variance random variable D^* such that $\frac{\partial}{\partial e}|_{e=0} \psi(e) = E(D^* S(O))$, which is the pathwise derivative of e . As before,

we will first consider the case where censoring is absent, and then the case where censoring is present.

C.2.1 Without Censoring

Based on Figure C.1, all causal paths from the static treatment a to the intermediate outcome Y_k are d-separated (blocked) by conditioning on L_0 [41]. This reduces the effect of the time-varying confounder L_k^* ($k > 0$) to L_0 . Then, we can rewrite the estimand expression as $\psi_{I^0,k} = \sum_{j=1}^k E(Y_j | L_0, A = a)$, where the corresponding EIF, denoted by $D_{I^0,k}^*(P) = D_{I^0,k}^*$, is simply the sum of the EIF expressions with a time-fixed confounder [51].

$$D_{I^0,k}^* = \sum_{j=1}^k \left[\frac{\mathbb{I}(A=a)}{\pi_A^a} (Y_j - Q_{j,0}) + Q_{j,0}^a - \psi_{I^0,k,j} \right],$$

where $\psi_{I^0,k,j} = E(Y_j | L_0, A = a)$, $Q_{j,h} = E(Q_{j,h+1} | \bar{Y}_h, \bar{L}_h, A)$, and $Q_{j,h}^a = E(Q_{j,h+1}^a | \bar{Y}_h, \bar{L}_h, A = a)$ for $0 \leq h \leq j \leq k$.

C.2.2 With Censoring

With $\psi_{I,k}$ represented by the ICE form, the corresponding EIF can be derived as:

$$\begin{aligned} D_{I,k}^* = \sum_{j=1}^k \left\{ \frac{\mathbb{I}(A=a) \prod_{i=1}^j (1 - C_i)}{\pi_A^a \bar{\pi}_{C_j}^0} [Y_j - Q_{j,j-1}] \right. \\ \left. + \sum_{m=2}^j \frac{\mathbb{I}(A=a) \prod_{h=1}^{m-1} (1 - C_h)}{\pi_A^a \bar{\pi}_{C_{m-1}}^0} [Q_{j,m-1}(C_m = 0) - Q_{j,m-2}] \right. \\ \left. + Q_{j,0}^a - \psi_{I,k,j} \right\}, \end{aligned}$$

where $Q_{k,k} = Y_k$ and for $0 \leq h < k$,

$$\begin{aligned} Q_{k,h} &= Q_{k,h}(\bar{C}_{h+1}, \bar{Y}_h, \bar{L}_h, A) = E(Q_{k,h+1} | \bar{C}_{h+1}, \bar{Y}_h, \bar{L}_h, A), \\ Q_{k,h}(C_{h+1} = 0) &= Q_{k,h}(\bar{C}_{h+1} = \bar{0}, \bar{Y}_h, \bar{L}_h, A) = E(Q_{k,h+1} | \bar{C}_{h+1} = \bar{0}, \bar{Y}_h, \bar{L}_h, A), \\ Q_{k,h}^a &= E(Q_{k,h+1} | \bar{C}_{h+1} = \bar{0}, \bar{Y}_h, \bar{L}_h, A = a), \text{ and} \\ \psi_{I,k,j} &= E\{E[\dots E(Y_k | \bar{C}_k = \bar{0}, \bar{L}_{k-1}, \bar{Y}_{k-1}, a) \dots | L_0]\}. \end{aligned}$$

Proof by induction: For the base case $k = 1$, we have the relevant observed data $O = (L_0, A = a, C_1 = 0, Y_1)$. Similarly to the result in [51], we have the corresponding EIF as follows.

$$D_{I,1}^* = \frac{\mathbb{I}(A = a)(1 - C_1)}{\pi_A^a \bar{\pi}_{C_1}^0} [Y_1 - Q_{1,0}] + Q_{1,0}^a - \psi_{I,1},$$

where

$$\psi_{I,1} = E[E(Y_1 | \bar{C}_1 = \bar{0}, \bar{L}_0^*, A = a)]$$

To illustrate the process of addition, we further derive the case for $k = 2$, where $O = (L_0, A = a, C_1 = 0, Y_1, L_1, C_2 = 0, Y_2)$; and the estimand of interest is written as

$$\begin{aligned} \psi_{I,2} &= \sum_{j=1}^2 E\{E[\dots E(Y_2 | \bar{C}_1 = \bar{0}, \bar{L}_1^*, A = a) \dots | C_1 = 0, L_0, A = a]\} \\ &= E[E(Y_1 | C_1 = 0, L_0, A = a)] + E\{E[E(Y_2 | \bar{C}_2 = \bar{0}, \bar{L}_1^*, A = a) | C_1 = 0, L_0, A = a]\} \\ &= \psi_{I,2,1} + \psi_{I,2,2}. \end{aligned}$$

Since $\psi_{I,2,1}$ is identical to the base case $\psi_{I,1}$, it is straightforward to write

$$D_{I,2,1}^* = \frac{\mathbb{I}(A = a)(1 - C_1)}{\pi_A^a \bar{\pi}_{C_1}^0} [Y_1 - Q_{1,0}] + Q_{1,0}^a - \psi_{I,2,1}.$$

Next, we focus on $\psi_{I,2,2}$.

$$\begin{aligned} \frac{\partial}{\partial e}|_{e=0} \psi_{I,2,2} &= \frac{\partial}{\partial e}|_{e=0} E\{E[E(Y_2 | \bar{C}_2 = \bar{0}, \bar{L}_1^*, A = a) | C_1 = 0, L_0, A = a]\} \\ &= E\{E\left[\sum_{\forall y_2} y_2 \frac{\partial}{\partial e} p(y_2 | \bar{C}_2 = \bar{0}, \bar{L}_1^*, a)|_{e=0} | \bar{C}_1 = \bar{0}, \bar{L}_0^*, a\right]\} \end{aligned} \quad (A)$$

$$+ E\left[\sum_{\forall l_1^*} E(Y_2 | \bar{C}_2 = \bar{0}, l_1^*, L_0, a) \frac{\partial}{\partial e} p(l_1^* | \bar{C}_1 = \bar{0}, L_0, a)|_{e=0}\right] \quad (B)$$

$$+ \sum_{\forall l_0} E[E(Y_2 | \bar{C}_2 = \bar{0}, \bar{L}_1^*, a) | C_1 = 0, l_0, a] \frac{\partial}{\partial e} p(l_0)|_{e=0} \quad (C)$$

To help with the derivation, we introduce two known identities based on [51]. For some observed data $(O_0, O_s, \dots, O_i, \dots, O_n)$, where $0 \leq s \leq i \leq n$.

$$\frac{\partial}{\partial e} p_{O_i, e}(o_i | \bar{o}_s) |_{e=0} = \{E[S(O) | \bar{O}_i = \bar{o}_i] - E[S(O) | \bar{O}_s = \bar{o}_s]\} p_{O_i}(o_i | \bar{o}_s). \quad (\text{C.2})$$

$$E[h(O_i) | O_s = o_s^*] = \frac{E[h(O_i) | \bar{O}_s] \mathbb{I}(O_s = o_s^*)}{P(O_s = o_s^* | \bar{O}_{s-1})} \quad (\text{C.3})$$

We discuss (A), (B) and (C) separately.

$$\begin{aligned} (A) &= E \left\{ E \left[\sum_{\forall y_2} y_2 \frac{\partial}{\partial e} p(y_2 | \bar{C}_2 = \bar{0}, \bar{L}_1^*, a) |_{e=0} | \bar{C}_1 = \bar{0}, \bar{L}_0^*, a \right] \right\} \\ \text{by (C.2)} \quad &= E \left\{ E \left[\sum_{\forall y_2} y_2 \{ E(S(O) | O = o) \right. \right. \\ &\quad \left. \left. - E(S(O) | \bar{C}_2, \bar{L}_1^*, A) \} p(y_2 | \bar{C}_2 = \bar{0}, \bar{L}_1^*, a) | \bar{C}_1 = \bar{0}, \bar{L}_0^*, a \right] \right\} \\ \text{by (C.3)} \quad &= \frac{\mathbb{I}(A = a)(1 - C_2)(1 - C_1)}{\pi_A^a \bar{\pi}_{C_2}^0} E \left\{ E \left[\sum_{\forall y_2} y_2 \{ S(o) \right. \right. \\ &\quad \left. \left. - E(S(O) | \bar{C}_2, \bar{L}_1^*, A) \} p(y_2 | \bar{C}_2, \bar{L}_1^*, A) | \bar{C}_1, \bar{L}_0^*, A \right] \right\} \\ &= \frac{\mathbb{I}(A = a)(1 - C_2)(1 - C_1)}{\pi_A^a \bar{\pi}_{C_2}^0} \sum_{\forall l_0, l_1^*, y_2} \{ y_2 S(o) \\ &\quad - \sum_{\forall y_2} y_2 S(o) p(y_2 | \bar{c}_2, \bar{l}_1^*, a) y_2 \} p(y_2 | \bar{c}_2, \bar{l}_1^*, a) p(l_1^* | \bar{c}_1, \bar{l}_0^*, a) p(l_0) \\ &= \frac{\mathbb{I}(A = a)(1 - C_2)(1 - C_1)}{\pi_A^a \bar{\pi}_{C_2}^0} \sum_{\forall o} [y_2 - E(y_2 | \bar{c}_2, \bar{l}_1^*, a)] S(o) p(o) \\ &= \frac{\mathbb{I}(A = a)(1 - C_2)(1 - C_1)}{\pi_A^a \bar{\pi}_{C_2}^0} \sum_{\forall o} [y_2 - Q_{2,1}] S(o) p(o) \\ (B) &= E \left[\sum_{\forall l_1^*} E(Y_2 | \bar{C}_2 = \bar{0}, l_1^*, L_0, a) \frac{\partial}{\partial e} p(l_1^* | \bar{C}_1 = \bar{0}, L_0, a) |_{e=0} \right] \end{aligned}$$

$$\begin{aligned}
\text{by (C.2)} \quad &= \sum_{\forall \bar{l}_1^*} E(Y_2 | \bar{C}_2 = \bar{0}, \bar{l}_1^*, a) [E(S(O) | \bar{C}_2, \bar{l}_1^*, A) \\
&\quad - E(S(O) | C_1, l_0, A)] p(l_1^* | C_1 = 0, l_0, a) p(l_0) \\
&= \sum_{\forall l_0, l_1^*, y_2} y_2 \left\{ \sum_{\forall y_2} S(O) p(y_2 | \bar{C}_2, \bar{l}_1^*, A) - \sum_{\forall y_2, l_1^*} S(O) p(y_2, l_1^* | C_1, \bar{l}_0^*, A) \right\} \\
&\quad p(y_2 | \bar{C}_2 = \bar{0}, \bar{l}_1^*, a) p(l_1^* | C_1 = 0, l_0, a) p(l_0) \\
&= \sum_{\forall l_0, l_1^*, y_2} [Q_{2,1}^a - Q_{2,0}^a] S(o) p(y_2, l_1^* | C_1 = 0, l_0, a) p(l_0) \\
\text{by (C.3)} \quad &= \sum_{\forall l_0, l_1^*, y_2} \frac{\mathbb{I}(A = a)(1 - C_1)}{\pi_A^a \bar{\pi}_{C_1}^0} [Q_{2,1}(\bar{C}_2 = \bar{0}, \bar{l}_1^*, A) - Q_{2,0}] S(o) p(y_2, l_1^* | C_1, l_0, a) p(l_0) \\
&= \sum_{\forall o} \frac{\mathbb{I}(A = a)(1 - c_1)}{\pi_A^a \bar{\pi}_{C_1}^0} [Q_{2,1}(c_2 = 0) - Q_{2,0}] S(o) p(o)
\end{aligned}$$

$$\begin{aligned}
(C) &= \sum_{\forall l_0} E \left[E(Y_2 | \bar{C}_2 = \bar{0}, \bar{l}_1^*, a) | C_1 = 0, l_0, a \right] \frac{\partial}{\partial e} p(l_0) |_{e=0} \\
\text{by (C.2)} \quad &= \sum_{\forall l_0, l_1^*, y_2} y_2 p(y_2 | \bar{C}_2 = \bar{0}, \bar{l}_1^*, a) p(l_1^* | C_1 = 0, l_0, a) \\
&\quad \left[E(S(O) | C_1 = 0, l_0, a) - E(S(O)) \right] p(l_0) \\
&= \sum_{\forall l_0, l_1^*, y_2} Q_{2,0}^a \sum_{\forall y_2, l_1^*} (S(o) p(y_2, l_1^* | C_1, l_0, A) p(l_0) - \sum_{\forall o} \psi_{I,2,2} S(o) p(o)) \\
&= \sum_{\forall o} [Q_{2,0}^a - \psi_{I,2,2}] S(o) p(o)
\end{aligned}$$

Combining the results, we derive the following.

$$\begin{aligned}
\frac{\partial}{\partial e} |_{e=0} \psi_{I,2,2} &= (A) + (B) + (C) \\
&= \left\{ \frac{\mathbb{I}(A = a)(1 - C_2)(1 - C_1)}{\pi_A^a \bar{\pi}_{C_2}^0} [Y_2 - Q_{2,1}] \right. \\
&\quad + \frac{\mathbb{I}(A = a)(1 - c_1)}{\pi_A^a \bar{\pi}_{C_1}^0} Q_{2,1}(C_2 = 0) - Q_{2,0} \\
&\quad \left. + [Q_{2,0}^a - \psi_{I,2,2}] S(o) \right\}
\end{aligned}$$

So, for the base case $k = 2$, we have the following as the end of our base case(s).

$$\begin{aligned}
D_{I,2}^* &= D_{I,2,1}^* + D_{I,2,2}^* \\
&= \frac{\mathbb{I}(A=a)(1-C_1)}{\pi_A^a \bar{\pi}_{C_1}^0} [Y_1 - Q_{1,0}] + Q_{1,0}^a - \psi_{I,2,1} \\
&\quad + \left\{ \frac{\mathbb{I}(A=a)(1-C_2)(1-C_1)}{\pi_A^a \bar{\pi}_{C_2}^0} [Y_2 - Q_{2,1}] \right. \\
&\quad + \frac{\mathbb{I}(A=a)(1-C_1)}{\pi_A^a \bar{\pi}_{C_1}^0} [Q_{2,1}(C_2=0) - Q_{2,0}] \\
&\quad \left. + [Q_{2,0}^a - \psi_{I,2,2}] \right\}
\end{aligned}$$

Next, we present the induction hypothesis. If the EIF derivation holds for $k = n > 2$, then we need to show that it also works for $k = n + 1$. i.e., suppose

$$\begin{aligned}
D_{I,n}^* &= \sum_{j=1}^n \left\{ \frac{\mathbb{I}(A=a) \prod_{h=1}^j (1-C_{h_1})}{\pi_A^a \bar{\pi}_{C_j}^0} [Y_j - Q_{j,j}] \right. \\
&\quad + \sum_{m=2}^j \frac{\mathbb{I}(A=a) \prod_{h=1}^{m-1} (1-C_h)}{\pi_A^a \bar{\pi}_{C_{m-1}}^0} [Q_{j,m}(C_{m+1}=0) - Q_{j,m-1}] \\
&\quad \left. + Q_{j,1}^a - \psi_{I,n,j} \right\},
\end{aligned}$$

is true, then we need to prove

$$\begin{aligned}
D_{I,n+1}^* &= \sum_{j=1}^{n+1} \left\{ \frac{\mathbb{I}(A=a) \prod_{h=1}^j (1-C_{h_1})}{\pi_A^a \bar{\pi}_{C_j}^0} [Y_j - Q_{j,j}] \right. \\
&\quad + \sum_{m=2}^j \frac{\mathbb{I}(A=a) \prod_{h=1}^{m-1} (1-C_h)}{\pi_A^a \bar{\pi}_{C_{m-1}}^0} [Q_{j,m}(C_{m+1}=0) - Q_{j,m-1}] \\
&\quad \left. + Q_{j,1}^a - \psi_{I,n+1,j} \right\}
\end{aligned}$$

holds.

proof:

$$\begin{aligned}
D_{I,n+1}^* &= D_{I,n+1,1}^* + D_{I,n+1,2}^* + \dots + D_{I,n+1,n+1}^* \\
&= \sum_{j=1}^{n+1} D_{I,n+1,j}^*
\end{aligned}$$

Each of $D_{I,n+1,j}^*$ ($j = 1, \dots, n+1$) can be derived by taking the pathwise derivative of $\psi_{I,n+1,j}$, where

$$\begin{aligned}
\psi_{I,n+1,j} &= E\{E[\dots E(Y_j|\bar{C}_j = \bar{0}, \bar{L}_{j-1}^*, A = a)|\bar{C}_{j-1} = \bar{0}, \bar{L}_{j-2}^*, A = a] \dots | C_1 = 0, L_0, A = a]\} \\
&= \sum_{\forall o} y_j p(y_j|\bar{c}_j = \bar{0}, \bar{l}_{j-1}^*, a) p(l_{j-1}^*|\bar{c}_{j-1} = \bar{0}, \bar{l}_j^*, a) \dots p(l_1^*|c_1 = 0, l_0, a) p(l_0)
\end{aligned}$$

$$\frac{\partial}{\partial e}|_{e=0} \psi_{I,n+1,j} = E\{E[\dots E(Y_j \frac{\partial}{\partial e} p(Y_j|\bar{C}_j = \bar{0}, \bar{L}_{j-1}^*, a)|_{e=0}|\bar{C}_j = \bar{0}, \bar{L}_{j-2}^*, a) \dots | C_1 = 0, L_0, a]\} \quad (E1)$$

$$+ E\{E[\dots E(Y_j|\bar{C}_j = \bar{0}, \bar{L}_{j-1}^*, a) \frac{\partial}{\partial e} p(L_{j-1}^*|\bar{C}_{j-1} = \bar{0}, \bar{L}_{j-2}^*, a)|_{e=0} \dots | C_1 = 0, L_0, a]\} \quad (E2)$$

.

.

.

$$+ E\{\sum_{\forall l_1^*} E[\dots E(Y_j|\bar{L}_{j-1}^*, a) \dots |\bar{l}_1^*] \frac{\partial}{\partial e} p(l_1^*|L_0)|_{e=0}\} \quad (Ej)$$

$$+ \sum_{\forall l_0} E[\dots E(Y_j|\bar{L}_{j-1}^*, a) \dots |l_0] \frac{\partial}{\partial e} p(l_0)|_{e=0} \quad (E(j+1))$$

Similarly, by the identities (C.2) and (C.3), it is easy to show the following.

$$(E1) = \sum_{\forall o} \frac{\mathbb{I}(A = a) \prod_{h_1=1}^j (1 - c_{h_1})}{\pi_A^a \bar{\pi}_{C_j}^0} (y_j - Q_{j,j-1}) S(o) p(o)$$

$$\begin{aligned}
(E2) &= \sum_{\forall o} \frac{\mathbb{I}(A = a) \prod_{h_1=1}^{j-1} (1 - c_{h_1})}{\pi_A^a \bar{\pi}_{C_j}^0} (Q_{j,j-1}(\bar{c}_j = 0) - Q_{j,j-2}) S(o) p(o) \\
&\quad \cdot \\
&\quad \cdot \\
&\quad \cdot \\
(Ej) &= \sum_{\forall o} \frac{\mathbb{I}(A = a) (1 - c_1)}{\pi_A^a \bar{\pi}_{C_1}^0} (Q_{j,1}(\bar{c}_2 = 0) - Q_{j,0}) S(o) p(o) \\
(E(j+1)) &= \sum_{\forall o} (Q_{j,0}^a - \psi_{I,n+1,j}) S(o) p(o)
\end{aligned}$$

Combining the results, we have the following.

$$\begin{aligned}
\frac{\partial}{\partial e} \Big|_{e=0} \psi_{I,n+1,j} &= (E1) + \dots + (E(j+1)) \\
&= \sum_{\forall o} \left\{ \frac{\mathbb{I}(A = a) \prod_{h_1=1}^j (1 - C_{h_1})}{\pi_A^a \bar{\pi}_{C_j}^0} [Y_j - Q_{j,j}] \right. \\
&\quad + \sum_{m=2}^j \frac{\mathbb{I}(A = a) \prod_{h=1}^{m-1} (1 - C_h)}{\pi_A^a \bar{\pi}_{C_{m-1}}^0} [Q_{j,m}(C_{m+1} = 0) - Q_{j,m-1}] \\
&\quad \left. + Q_{j,1}^a - \psi_{I,n+1,j} \right\} S(o) p(o)
\end{aligned}$$

For $j = 1, \dots, n+1$,

$$\begin{aligned}
D_{I,n+1,j}^* &= \frac{\mathbb{I}(A = a) \prod_{h_1=1}^j (1 - C_{h_1})}{\pi_A^a \bar{\pi}_{C_j}^0} [Y_j - Q_{j,j}] \\
&\quad + \sum_{m=2}^j \frac{\mathbb{I}(A = a) \prod_{h=1}^{m-1} (1 - C_h)}{\pi_A^a \bar{\pi}_{C_{m-1}}^0} [Q_{j,m}(C_{m+1} = 0) - Q_{j,m-1}] \\
&\quad + Q_{j,1}^a - \psi_{I,n+1,j}
\end{aligned}$$

Therefore,

$$D_{I,n+1}^* = \sum_{j=1}^{n+1} \left\{ \frac{\mathbb{I}(A = a) \prod_{h_1=1}^j (1 - C_{h_1})}{\pi_A^a \bar{\pi}_{C_j}^0} [Y_j - Q_{j,j}] \right.$$

$$\begin{aligned}
& + \sum_{m=2}^j \frac{\mathbb{I}(A=a) \prod_{h=1}^{m-1} (1-C_h)}{\pi_A^a \bar{\pi}_{C_{m-1}}^0} [Q_{j,m}(C_{m+1}=0) - Q_{j,m-1}] \\
& + Q_{j,1}^a - \psi_{I,n+1,j} \}
\end{aligned}$$

This completes the induction process to arrive at the conclusion that the EIF for $\psi_{I,k} = \sum_{j=1}^k E_P(Y_j^a)$ (with censoring) is

$$\begin{aligned}
D_{I,k}^* = \sum_{j=1}^k \Big\{ & \frac{\mathbb{I}(A=a) \prod_{h=1}^j (1-C_h)}{\pi_A^a \bar{\pi}_{C_j}^0} [Y_j - Q_{j,j}] \\
& + \sum_{m=2}^j \frac{\mathbb{I}(A=a) \prod_{h=1}^{m-1} (1-C_h)}{\pi_A^a \bar{\pi}_{C_{m-1}}^0} [Q_{j,m}(C_{m+1}=0) - Q_{j,m-1}] \\
& + Q_{j,1}^a - \psi_{I,k,j} \Big\},
\end{aligned}$$

where

$$\psi_{I,k,j} = E\{E[\dots E[E(Y_j|\bar{C}_j = \bar{0}, \bar{L}_{j-1}^*, A=a)|\bar{C}_{j-1} = \bar{0}, \bar{L}_{j-2}^*, A=a] \dots | C_1=0, L_0, A=a]\}.$$

end of proof

C.3 Generalized form for some of the estimators

C.3.1 Iterative conditional expectation (ICE) estimator:

$$\hat{\psi}_{\text{ICE}}$$

The ICE estimator can be naturally formulated by iteratively fitting a outcome model on its corresponding conditions. Define $Q_{k,k} = Y_k$, and iteratively for $0 \leq h < k$, let

$$Q_{k,h} = E(Q_{k,h+1} | \bar{C}_{h+1} = \bar{0}, \bar{Y}_h, \bar{L}_h, A=a),$$

and $\hat{Q}_{k,h}$ is the corresponding estimates for $Q_{k,h}$. Then, the ICE estimator for $\psi_{I,k}$, denoted by $\hat{\psi}_{\text{ICE}}$, can be acquired as follows.

1. Set index $i = 1$.

2. Estimate $Q_{k,k-i}$:
 - a. model $\hat{Q}_{k,k-i+1}$ on past measured variables using observations with treatment level a which are uncensored until time interval $k - i + 1$;
 - b. predict $\hat{Q}_{k,k-i}$ using the fitted model with observations with treatment level a and are uncensored until time interval $k - i$.
3. $i = i + 1$. If $i \leq k$, go back to the last step.
4. Calculate $\hat{\psi}_{\text{ICE}} = \mathbb{P}_n(\hat{Q}_{k,0})$ for all individuals of the study.

C.3.2 Augmented inverse probability weighted (AIPW) estimator: $\hat{\psi}_{\text{AIPW}}$

The explicit form of the EIF for $\psi_{I,k}$, denoted by $D_{I,k}^*$, is given as:

$$\begin{aligned}
 D_{I,k}^* = & \left\{ \frac{\mathbb{I}(A = a) \prod_{j=1}^k (1 - C_j)}{\pi_A^a \bar{\pi}_{C_k}^0} [Y_k - Q_{k,k-1}] \right. \\
 & + \sum_{m=1}^{k-1} \frac{\mathbb{I}(A = a) \prod_{h=1}^{m+1} (1 - C_h)}{\pi_A^a \bar{\pi}_{C_{m+1}}^0} [Q_{k,m}(C_{m+1} = 0) - Q_{k,m-1}] \\
 & \left. + Q_{k,0}^a - \psi_{I,k} \right\}, \tag{C.4}
 \end{aligned}$$

where for $0 \leq h < k$ and with a slight abuse of notation,

$$Q_{k,h}^a = E(Y_k | \bar{C}_{h+1} = \bar{0}, \bar{Y}_h, \bar{L}_h, A = a).$$

$$Q_{k,h} = Q_{k,h}(\bar{C}_{h+1}, \bar{Y}_h, \bar{L}_h, A) = E(Y_k | \bar{C}_{h+1}, \bar{Y}_h, \bar{L}_h, A),$$

and

$$Q_{k,h}(C_{h+1} = 0) = Q_{k,h}(\bar{C}_{h+1} = \bar{0}, \bar{Y}_h, \bar{L}_h, A) = E(Y_k | \bar{C}_{h+1} = \bar{0}, \bar{Y}_h, \bar{L}_h, A).$$

The detailed derivation of $D_{I,k}^*$ is provided in Section C.2. Then, the AIPW estimator can be written as:

$$\begin{aligned}
\hat{\psi}_{\text{AIPW}} = & \mathbb{P}_n \left\{ \frac{\mathbb{I}(A = a) \mathbb{I}(C_k = 0)}{\hat{\pi}_A^a \bar{\pi}_{C_k}^0} [Y_k - \hat{Q}_{k,k-1}] \right. \\
& + \sum_{m=1}^{k-1} \frac{\mathbb{I}(A = a) \mathbb{I}(C_{m+1} = 0)}{\hat{\pi}_A^a \bar{\pi}_{C_{m+1}}^0} [\hat{Q}_{k,m} - \hat{Q}_{k,m-1}] \\
& \left. + \hat{Q}_{k,0}^a \right\}, \tag{C.5}
\end{aligned}$$

where $\hat{\pi}_A^a$, $\bar{\pi}_{C_{h+1}}^0$, and $\hat{Q}_{k,h}$ ($0 \leq h < k$) are derived the same way as in Section 3.2.3 of the IPW estimator and in Section 3.2.5, the ICE estimators, respectively, with suitable working models.

C.3.3 Weighted iterative conditional expectation (WICE) estimator: $\hat{\psi}_{\text{WICE}}$

Given the EIF expression of Equation C.4, the WICE estimation process follows the steps below.

1. Obtain $\bar{\pi}_{C_j}^0$ ($j \in \{1, \dots, k\}$) and $\hat{\pi}_A^a$ as in Section 3.2.3.
2. Set index $i = 1$.
3. Estimate $Q_{k,k-i}$ as in Section 3.2.5 except with weights as follows:
 - a. model $\hat{Q}_{k,k-i+1}$ on past measured variables, where each observation is weighted by inverse of the estimated cumulative probability of treatment and censoring, or $[\bar{\pi}_{C_{k-i}}^0 \hat{\pi}_A^a]^{-1}$;
 - b. predict $\hat{Q}_{k,k-i}$ with derived model on observations with treatment level $A = a$ and are uncensored until time interval $k - i$.
4. $i = i + 1$. If $i \leq k$, go back to the last step.
5. Calculate $\hat{\psi}_{\text{WICE}} = \mathbb{P}_n(\hat{Q}_{k,0})$ for all individuals of the study.

C.3.4 Targeted maximum likelihood (TML) estimator:

$$\hat{\psi}_{\text{TML}}$$

The detailed estimation steps for TML estimator are given as follows. Define $\hat{Q}_{k,k}^* = Y_k$ for ease of notation.

1. Obtain $\bar{\pi}_{C_j}^0 (j \in \{1, \dots, k\})$ and $\hat{\pi}_A^a$ as in Section 3.2.3.
2. Set index $i = 1$.
3. Estimate $Q_{k,k-i}^*$:
 - i. select individuals that are uncensored until time interval $k - i + 1$ with treatment level a ;
 - ii. model $\hat{Q}_{k,k-i+1}^*$ on individuals selected in step i to obtain the initial estimate $\hat{Q}_{k,k-i}$ (as in Section 3.2.7); denote the model as F_0 ;
 - iii. regress $\hat{Q}_{k,k-i+1}^*$ on 1 with an offset (fixed intercept) of $\hat{Q}_{k,k-i}$ weighted by $[\bar{\pi}_{C_{k-i}}^0 \hat{\pi}_A^a]^{-1}$; denote the new model as F_1 ;
 - iv. select individuals that are uncensored until time interval $k - i$ with treatment level a ;
 - v. obtain a new offset term with the fitted model F_0 on newly selected data in step iv, denoted Q^* ;
 - vi. obtain the updated estimate $\hat{Q}_{k,k-i}^{a,*}$ by applying F_1 on new offset Q^* and the selected data in step iv.
4. $i = i + 1$ for $i \leq k$; go back to Step 3.
5. Calculate $\hat{\psi}_{\text{TML}} = \mathbb{P}_n(\hat{Q}_{k,0})$ for all individuals in the study.

Appendix D

Supplementary Information for Chapter 4

In this appendix, we follow a similar structure as in Appendix C, to show derivations for identification and the efficient influence function (EIF) for our defined causal estimand of interest in Chapter 4. Then we will provide generalized forms for the WICE, AIPW, and TML estimators.

D.1 Identification

Given a time-varying treatment strategy, our relevant data is

$$O = (L_0, A_0^*, C_1, Y_1, L_1^*, A_1^*, \dots, C_K, Y_K, L_K^*, A_K^*, C_{K+1}, Y_{K+1}),$$

or

$$O = (L_0, A_0, L_1, A_1, \dots, L_K, A_K, Y_{K+1}).$$

The intermediate causal estimand of interest, $E_P(Y_k^{\bar{a}_{k-1}})$, the expected count of events (0 or 1) occurred in interval $k-1$ with a treatment path $\bar{a}_{k-1}$, can be identified through causal assumptions and written into an iteratively conditional expectation (ICE) as follows.

$$\begin{aligned}
\psi_{II,k}(P) &= E_P(Y_k^{\bar{A}_{k-1}=\bar{a}_{k-1}}) = \sum_{\forall l_0} E(Y_k^{\bar{a}_{k-1}}|l_0)p(l_0) \\
&\text{(exchangeability)} \\
&= \sum_{\forall l_0} E(Y_k^{\bar{a}_{k-1}}|a_0, l_0)p(l_0) \\
&\text{(consistency \& exchangeability)} \\
&= \sum_{\forall \bar{l}_1} E(Y_k^{\bar{a}_{k-1}}|\bar{l}_1, \bar{a}_1)p(l_1|a_0, l_0)p(l_0) \\
&\text{(for } t=1, \dots, k-2, \text{ recursively:)} \\
&= \sum_{\forall \bar{l}_t} E(Y_k^{\bar{a}_{k-1}}|\bar{l}_t, \bar{a}_t)p(l_t|\bar{l}_{t-1}, \bar{a}_{t-1}) \dots p(l_1|a_0, l_0)p(l_0) \\
&\text{(exchangeability \& consistency)} \\
&= \sum_{\forall \bar{l}_{t+1}} E(Y_k^{\bar{a}_{k-1}}|\bar{l}_{t+1}, \bar{a}_{t+1})p(l_{t+1}|\bar{l}_t, \bar{a}_t) \dots p(l_1|a_0, l_0)p(l_0) \\
&\text{(by induction and for } t = k-2) \\
&= \sum_{\forall \bar{l}_{k-1}} E(Y_k^{\bar{a}_{k-1}}|\bar{l}_{k-1}, \bar{a}_{k-1})p(l_{k-1}|\bar{l}_{k-2}, \bar{a}_{k-2}) \dots p(l_1|a_0, l_0)p(l_0) \\
&\text{(consistency)} \\
&= \sum_{\forall \bar{l}_{k-1}} E(Y_k|\bar{l}_{k-1}, \bar{a}_{k-1})p(l_{k-1}|\bar{l}_{k-2}, \bar{a}_{k-2}) \dots p(l_1|a_0, l_0)p(l_0)
\end{aligned}$$

The corresponding estimand is:

$$\psi_{II,k} = E\{E[E(\dots E[E(Y_k|\bar{L}_{k-1}, a_{k-1})|\bar{L}_{k-2}, a_{k-2}] \dots |\bar{L}_1, a_1)|L_0]\}. \quad (\text{D.1})$$

Alternatively, $E(Y_k^{\bar{a}_{k-1}})$ can be expressed in the inverse probability weighted form. The derivation process is demonstrated as follows.

$$\begin{aligned}
E(Y_k^{\bar{a}_{k-1}}) &= \sum_{\forall l_0} E(Y_k^{\bar{a}_{k-1}} | l_0) p(l_0) \\
&\text{(exchangeability)} \\
&= \sum_{\forall l_0} E(Y_k^{\bar{a}_{k-1}} | A_0, l_0) p(l_0) \\
&\text{(probability rules)} \\
&= \sum_{\forall l_0} \frac{\mathbb{I}(A_0 = a_0)}{\pi_{A_0}^{a_0}} E(Y_k^{\bar{a}_{k-1}} | a_0, l_0) p(l_0) \\
&\text{(for } t = 0, \dots, k-2, \text{ recursively)} \\
&= \frac{\prod_{h=1}^{t-1} \mathbb{I}(A_h = a_h)}{\bar{\pi}_{A_{t-1}}^{a_{t-1}}} \sum_{\forall \bar{l}_t} E(Y_k^{\bar{a}_{k-1}} | \bar{a}_{t-1}, \bar{l}_t) p(\bar{l}_t) \\
&\text{(exchangeability \& probability rules)} \tag{D.2} \\
&= \frac{\prod_{h=1}^t \mathbb{I}(A_h = a_h)}{\bar{\pi}_{A_t}^{a_t}} \sum_{\forall \bar{l}_{k-1}} E(Y_k^{\bar{a}_{k-1}} | \bar{a}_t, \bar{l}_t) p(\bar{l}_t) \\
&\text{(by induction and } t = k-2) \\
&= \frac{\prod_{h=1}^{k-1} \mathbb{I}(A_h = a_h)}{\bar{\pi}_{A_{k-1}}^{a_{k-1}}} \sum_{\forall \bar{l}_{k-1}} E(Y_k^{\bar{a}_{k-1}} | \bar{a}_{k-1}, \bar{l}_{k-1}) p(\bar{l}_{k-1}) \\
&\text{(consistency)} \\
&= \frac{\prod_{h=1}^{k-1} \mathbb{I}(A_h = a_h)}{\bar{\pi}_{A_{k-1}}^{a_{k-1}}} \sum_{\forall \bar{l}_{k-1}} E(Y_k | \bar{a}_{k-1}, \bar{l}_{k-1}) p(\bar{l}_{k-1}) \\
&= \frac{\prod_{h=1}^{k-1} \mathbb{I}(A_h = a_h)}{\bar{\pi}_{A_{k-1}}^{a_{k-1}}} E(Y_k | \bar{a}_{k-1}, \bar{L}_{k-1})
\end{aligned}$$

D.2 EIF for $\psi_{II,k}$

Similarly as in Appendix C.2, we denote the EIF for $\psi_{II,k} = E_P(Y_k^{\bar{a}_k})$ as $D_{II,k}^*$, which can be derived as:

$$\begin{aligned} D_{II,k}^* = & \sum_{j=1}^k \left\{ \frac{\mathbb{I}(\bar{A}_{j-1} = \bar{a}_{j-1})}{\bar{\pi}_{A_{j-1}}^{a_{j-1}}} [Y_j - Q_{j,j-1}] \right. \\ & + \sum_{m=2}^j \frac{\mathbb{I}(\bar{A}_{m-2} = \bar{a}_{m-2})}{\bar{\pi}_{A_{m-2}}^{a_{m-2}}} [Q_{j,j-1} - Q_{j,j-2}] \\ & \left. + Q_{j,0}^{a_0} - \psi_{II,k,j} \right\}, \end{aligned}$$

where $Q_{k,k} = Q_{k,k}^{\bar{a}_k} = Y_k$, iteratively for $0 \leq h < k$

$$Q_{k,h} = E(Q_{k,h+1} | \bar{L}_h, \bar{A}_h), Q_{k,h}^{\bar{a}_h} = E(Q_{k,h+1}^{\bar{a}_{h+1}} | \bar{a}_h, \bar{L}_h),$$

and

$$\psi_{II,k,j} = E\{E[E(\dots E[E(Y_j | \bar{L}_{j-1}, \bar{a}_{j-1}) | \bar{L}_{j-2}, \bar{a}_{j-2}] \dots | \bar{L}_1, \bar{a}_1) | L_0]\}.$$

Proof by induction: For the base case $k = 1$, we have the relevant observed data $O = (L_0, A_0, Y_1)$. Again, according to previous work (Section 4.2 of [51]), we have the following.

$$D_{II,1}^* = \left[\frac{\mathbb{I}(A_0 = a_0)}{P(A_0 = a_0 | L_0)} (Y_1 - Q_{1,0}) + Q_{1,0}^{a_0} - \psi_{II,1} \right].$$

Next, the induction hypothesis is: If the EIF derivation holds for $k = n > 1$, then we need to show that it also works for $k = n + 1$. i.e., suppose that

$$\begin{aligned} D_{II,n}^* = & \sum_{j=1}^n \left\{ \frac{\mathbb{I}(\bar{A}_{j-1} = \bar{a}_{j-1})}{\bar{\pi}_{A_{j-1}}^{a_{j-1}}} [Y_j - Q_{j,j-1}] \right. \\ & + \sum_{m=2}^j \frac{\mathbb{I}(\bar{A}_{m-2} = \bar{a}_{m-2})}{\bar{\pi}_{A_{m-2}}^{a_{m-2}}} [Q_{j,j-1} - Q_{j,j-2}] \\ & \left. + Q_{j,0}^{a_0} - \psi_{II,n,j} \right\}, \end{aligned}$$

is true, then we are required to prove

$$D_{II,n+1}^* = \sum_{j=1}^{n+1} \left\{ \frac{\mathbb{I}(\bar{A}_{j-1} = \bar{a}_{j-1})}{\bar{\pi}_{\bar{A}_{j-1}}^{a_{j-1}}} [Y_j - Q_{j,j-1}] \right. \\ \left. + \sum_{m=2}^j \frac{\mathbb{I}(\bar{A}_{m-2} = \bar{a}_{m-2})}{\bar{\pi}_{\bar{A}_{m-2}}^{a_{m-2}}} [Q_{j,j-1} - Q_{j,j-2}] \right. \\ \left. + Q_{j,0}^{a_0} - \psi_{II,n+1,j} \right\}$$

holds.

proof:

$$D_{II,n+1}^* = D_{II,n+1,1}^* + D_{II,n+1,2}^* + \dots + D_{II,n+1,n+1}^* \\ = \sum_{j=1}^{n+1} D_{II,n+1,j}^*$$

Each of $D_{II,n+1,j}^*$ ($j = 1, \dots, n+1$) can be derived by taking the pathwise derivative of $\psi_{II,n+1,j}$, where

$$\psi_{II,n+1,j} = E\{E[E(\dots E[E(Y_j|\bar{a}_{j-1}, \bar{l}_{j-1})|\bar{a}_{j-2}, \bar{l}_{j-2}] \dots |\bar{a}_1, \bar{l}_1)|L_0]\} \\ = \sum_{\forall o} y_j p(y_j|\bar{a}_{j-1}, \bar{l}_{j-1}) p(l_{j-1}|\bar{a}_{j-1}, \bar{l}_j) \dots p(l_1|a_0, l_0) p(l_0).$$

Then,

$$\frac{\partial}{\partial e}|_{e=0} \psi_{II,n+1,j} = E\{E[\dots E(\sum_{\forall y_j} y_j \frac{\partial}{\partial e} p(y_j|\bar{a}_{j-1}, \bar{l}_{j-1})|_{e=0} |\bar{a}_{j-2}, \bar{l}_{j-2}] \dots |L_0]\} \quad (E1)$$

$$+ E\{E[\dots \sum_{\forall l_{j-1}} E(Y_j|\bar{a}_{j-1}, l_{j-1}, \bar{l}_{j-2}) \cdot \\ \frac{\partial}{\partial e} p(l_{j-1}|\bar{a}_{j-2}, \bar{l}_{j-2})|_{e=0} \dots |L_0]\} \quad (E2)$$

.

.

.

$$\begin{aligned}
& + E\left\{\sum_{\forall l_1} E[\dots E(Y_j|\bar{L}_{j-1}, a_{j-1})\dots|\bar{a}_1, l_1, L_0] \frac{\partial}{\partial e} p(l_1|a_0, L_0)|_{e=0}\right\} & (Ej) \\
& + \sum_{\forall l_0} E[\dots E(Y_j|\bar{a}_{j-1}, \bar{L}_{j-1})\dots|a_0, l_0] \frac{\partial}{\partial e} p(l_0)|_{e=0} & (E(j+1))
\end{aligned}$$

Similarly, by the identities (C.2) and (C.3), it is easy to show the following.

$$\begin{aligned}
(E1) &= \sum_{\forall o} \frac{\mathbb{I}(\bar{A}_{j-1} = \bar{a}_{j-1})}{\bar{\pi}_{A_{j-1}}^{a_{j-1}}} (y_j - Q_{j,j-1}) S(o) p(o) \\
(E2) &= \sum_{\forall o} \frac{\mathbb{I}(\bar{A}_{j-2} = \bar{a}_{j-2})}{\bar{\pi}_{A_{j-2}}^{a_{j-2}}} (Q_{j,j-1} - Q_{j,j-2}) S(o) p(o) \\
&\quad \cdot \\
&\quad \cdot \\
&\quad \cdot \\
(Ej) &= \sum_{\forall o} \frac{\mathbb{I}(\bar{A}_1 = \bar{a}_1)}{\bar{\pi}_{A_1}^{a_1}} (Q_{j,1} - Q_{j,0}) S(o) p(o) \\
(E(j+1)) &= \sum_{\forall o} (Q_{j,0}^{a_0} - \psi_{II,n+1,j}) S(o) p(o)
\end{aligned}$$

Combining the results, we have the following.

$$\begin{aligned}
\frac{\partial}{\partial e} \big|_{e=0} \psi_{II,n+1,j} &= (E1) + \dots + (E(j+1)) \\
&= \sum_{\forall o} \left\{ \frac{\mathbb{I}(\bar{A}_{j-1} = \bar{a}_{j-1})}{\bar{\pi}_{A_{j-1}}^{a_{j-1}}} [y_j - Q_{j,j-1}] \right. \\
&\quad + \sum_{m=2}^j \frac{\mathbb{I}(\bar{A}_{m-2} = \bar{a}_{m-2})}{\bar{\pi}_{A_{m-2}}^{a_{m-2}}} [Q_{j,j-1} - Q_{j,j-2}] \\
&\quad \left. + Q_{j,0}^{a_0} - \psi_{II,n+1,j} \right\} S(o) p(o)
\end{aligned}$$

This leads us to

$$D_{II,n+1,j}^* = \frac{\mathbb{I}(\bar{A}_{j-1} = \bar{a}_{j-1})}{\bar{\pi}_{A_{j-1}}^{a_{j-1}}} [Y_j - Q_{j,j-1}]$$

$$\begin{aligned}
& + \sum_{m=2}^j \frac{\mathbb{I}(\bar{A}_{m-2} = \bar{a}_{m-2})}{\bar{\pi}_{A_{m-2}}^{a_{m-2}}} [Q_{j,j-1} - Q_{j,j-2}] \\
& + Q_{j,0}^{a_0} - \psi_{II,n+1,j}
\end{aligned}$$

Therefore,

$$\begin{aligned}
D_{II,n+1}^* = \sum_{j=1}^{n+1} \Big\{ & \frac{\mathbb{I}(\bar{A}_{j-1} = \bar{a}_{j-1})}{\bar{\pi}_{A_{j-1}}^{a_{j-1}}} [Y_j - Q_{j,j-1}] \\
& + \sum_{m=2}^j \frac{\mathbb{I}(\bar{A}_{m-2} = \bar{a}_{m-2})}{\bar{\pi}_{A_{m-2}}^{a_{m-2}}} [Q_{j,j-1} - Q_{j,j-2}] \\
& + Q_{j,0}^{a_0} - \psi_{II,n+1,j} \Big\}.
\end{aligned}$$

end of proof

D.3 Generalized procedures for some estimators

D.3.1 Iterative conditional expectation (ICE) estimator:

$\hat{\psi}_{\text{ICE}}$

Define $Q_{k,k} = \hat{Q}_{k,k} = Y_k$, and iteratively for $j \in \{0, 1, \dots, k-1\}$, let

$$Q_{k,j} = E(Q_{k,j+1} | \bar{a}_j^*, \bar{C}_{j+1} = 0, \bar{L}_j),$$

and $\hat{Q}_{k,j}$ is the corresponding estimate for $Q_{k,j}$. Then, we describe an ICE estimator for $\psi_{II,k}$, denoted by $\hat{\psi}_{\text{ICE}}$, to be acquired as follows.

1. Set index $i = 1$.
2. Estimate $Q_{k,k-i}$ as follows:
 - i. select individuals that are uncensored until interval $k-i+1$ with treatment path a_{k-i}^* ;

- ii. model $\hat{Q}_{k,k-i+1}$ on past measured variables $(\bar{a}_{k-i}^*, \bar{Y}_{k-i}, \bar{L}_{k-i}^*)$ using data selected in the , denote the model as F_a ;
 - iii. select individuals that are uncensored until interval $k - i$ with treatment path a_{k-i-1}^* ;
 - iv. acquire the fitted values using model F_a on data from the last step, denote the predicted value as $\hat{Q}_{k,k-i}$.
3. $i = i + 1$. If $i \leq k$, go back to the last step.
4. Calculate $\hat{\psi}_{\text{ICE}} = \mathbb{P}_n(\hat{Q}_{k,0})$ for all observations of the study.

D.3.2 Augmented inverse probability weighted (AIPW) estimator: $\hat{\psi}_{\text{AIPW}}$

The EIF for $\psi_{II,k}$, denoted by $D_{II,k}^*$, is given as:

$$\begin{aligned}
 D_{II,k}^* = & \left\{ \frac{\mathbb{I}(\bar{A}_{k-1}^* = \bar{a}_{k-1}^*, \bar{C}_k = \bar{0})}{\bar{\pi}_{A_{k-1}}^{a_{k-1}}} [Y_k - Q_{k,k-1}] \right. \\
 & + \sum_{j=0}^{k-1} \frac{\mathbb{I}(\bar{A}_j^* = \bar{a}_j^*, \bar{C}_{j+1} = \bar{0})}{\bar{\pi}_{A_j}^{a_j}} [Q_{k,j+1} - Q_{k,j}] \\
 & \left. + Q_{k,0} - \psi_{II,k} \right\}, \tag{D.3}
 \end{aligned}$$

where $Q_{k,k} = Y_k$ and for $j \in \{0, 1, \dots, k-1\}$, $Q_{k,j} = E(Q_{k,j+1} | \bar{a}_j^*, \bar{C}_{j+1} = 0, \bar{L}_j)$. Note that the detailed derivation process of $D_{II,k}^*$ is given in Section D.2.

The AIPW estimator, can be constructed by solving the estimating equation of the EIF's empirical form. Specifically,

$$\begin{aligned}
 \hat{\psi}_{\text{AIPW}} = & \mathbb{P}_n \left\{ \frac{\mathbb{I}(\bar{A}_{k-1}^* = \bar{a}_{k-1}^*, \bar{C}_k = \bar{0})}{\hat{\pi}_{A_{k-1}}^{a_{k-1}}} [Y_k - \hat{Q}_{k,k-1}] \right. \\
 & + \sum_{j=0}^{k-1} \frac{\mathbb{I}(\bar{A}_j^* = \bar{a}_j^*, \bar{C}_{j+1} = \bar{0})}{\hat{\pi}_{A_j}^{a_j}} [\hat{Q}_{k,j+1} - \hat{Q}_{k,j}] \\
 & \left. + \hat{Q}_{k,0} \right\}, \tag{D.4}
 \end{aligned}$$

where $\bar{\pi}_{A_j}^{a_j}$ and $\hat{Q}_{k,j}$ ($j \in \{0, 1, \dots, k-1\}$) are derived as for the IPW estimator and the ICE estimator, respectively, with suitable working models.

D.3.3 Weighted iterative conditional expectation (WICE) estimator: $\hat{\psi}_{\text{WICE}}$

The WICE estimator can be derived following the steps below.

1. Obtain $\bar{\pi}_{A_j}^{a_j}$ ($j \in \{0, \dots, k-1\}$) as described in Section 4.2.2.
2. Set index $i = 1$.
3. Estimate $Q_{k,k-i}$ similar to the ICE estimator but with weights as follows:
 - i. select subjects that are uncensored until interval $k-i+1$ with treatment path a_{k-i}^* ;
 - ii. model $\hat{Q}_{k,k-i+1}$ on past measured variables $(\bar{a}_{k-i}^*, \bar{Y}_{k-i}, \bar{L}_{k-i}^*)$ using selected data from step i, while each observation is weighted by $[\bar{\pi}_{A_{k-i}}^{a_{k-i}}]^{-1}$, denote the model as F_a ;
 - iii. select subjects that are uncensored until interval $k-i$ with treatment path a_{k-i-1}^* ;
 - iv. acquire the fitted values using model F_a on the selected data from step iii, denote the predicted value as $\hat{Q}_{k,k-i}$.
4. $i = i + 1$. If $i \leq k$, go back to the last step.
5. Calculate $\hat{\psi}_{\text{WICE}} = \mathbb{P}_n(\hat{Q}_{k,0})$ for all observations of the study.

D.3.4 Targeted maximum likelihood (TML) estimator: $\hat{\psi}_{\text{TML}}$

The TML estimator is obtained as follows.

1. Obtain $\bar{\pi}_{A_j}^{a_j}$ ($j \in \{0, \dots, k-1\}$) as described for the IPW estimator.
2. Set index $i = 1$.

3. Estimate $Q_{k,k-i}$ similar to the WICE estimator but with an additional update:
 - i. select observations that are uncensored until time interval $k - i + 1$ with treatment path a_{k-i}^* ;
 - ii. model $\hat{Q}_{k,k-i+1}^*$ on past measured variables with data selected from step i and then weighted by $[\bar{\pi}_{A_{k-i}}^{a_{k-i}}]^{-1}$, denote the model as F_a ;
 - iii. obtain (the initial estimate of) $\hat{Q}_{k,k-i}$ by fitting model F_a on the data selected from step i;
 - iv. model $\hat{Q}_{k,k-i+1}^*$ on 1 with an offset of $\hat{Q}_{k,k-i}$ on the data selected from step i and then weighted by $[\bar{\pi}_{A_{k-i}}^{a_{k-i}}]^{-1}$; denote the new model as F_a^* ;
 - v. select observations that are uncensored until time interval $k - i$ with treatment path a_{k-i-1}^* ;
 - vi. obtain the fitted values, denoted as Q^* , using model F_a on selected data from the last step;
 - vii. obtain the updated estimate $\hat{Q}_{k,k-i}^*$ by fitting F_a^* using the new offset Q^* with the newly selected data.
4. $i = i + 1$ for $i \leq k$; go back to Step 3.
5. Calculate $\hat{\psi}_{\text{TML}} = \mathbb{P}_n(\hat{Q}_{k,0})$ for all observations in the study.

D.4 Supplementary Application Results with the SPRINT dataset

By implementing the targeted maximum likelihood estimation with static or time-varying treatments on the SPRINT dataset, we can answer a variety of related research questions as follows.

D.4.1 Total Intention-to-treat Effect of Intensive Blood Pressure Therapy

With the assumptions 1–4 listed in Section 3.2, our causal contrast of $\Psi_{I,TE}(P) = E_P(R_2^{z=1} - R_2^{z=0})$ can also be identified from the observed data of SPRINT, by treating D_k as a time-varying confounder (part of L_k) with a deterministic relationship with




Figure D.1: A single world intervention graph (SWIG) that identifies the causal treatment effect and illustrates the relationship between the measured baseline variables L_0 , the initial treatment assignment (Z), the competing event of death $\bar{D}_2$, the confounders $\bar{L}_1$, and the (intermediate) recurrent outcome $\bar{Y}_2$, assuming there is no unmeasured confounding.

the recurrent outcome. Figure D.1 demonstrates the causal relationship among key variables. We applied the proposed TML estimator in Section 3.2.8 and found an estimated average total effect of 0.019 with a 95% confidence interval of [0.011,0.027] during the first four years of the SPRINT study. Hence, all individuals assigned the intensive blood pressure lowering target are expected to have more acute kidney injury events than if they were assigned the standard target, regardless of their survival status in the first four years of the SPRINT study.

It is noteworthy that the estimated total effect is smaller than the previously estimated controlled direct effect. This observation suggests that the competing risk of death operates in the opposite direction to the direct treatment effect of blood pressure therapies. By controlling for death, we may inadvertently select for a healthier subgroup of the study population, thereby influencing the observed treatment effect.

D.4.2 Population Intervention Effect of Adherence

In considering the rarity of complete non-adherence within this population, our emphasis could instead be shifted to a population intervention effect (PIE) by comparing the average outcomes when the entire population is treated versus the actual outcomes observed in the real-world population. We assume that PIE would be different under the standard and intensive blood pressure therapy assignment. Therefore, the PIE will be defined conditioning on the initial intensive or standard blood pressure lowering therapy assignment Z , so that we could consider the PIE within either the intensive or the standard the blood pressure lowering therapy group. In this case, we are estimating the expected potential number of recurrent event outcomes under full adherence and the observed outcomes as it is, conditional on participants' assigned therapy group. Hence, our causal estimand of interests can be expressed as $E_P(R_{K+1}^{\bar{a}_{K+1}^1} \mid Z = z) - E_P(R_{K+1} \mid Z = z), z \in \{0, 1\}$, which captures both the direct effect and the indirect effect (through death) of full adherence on recurrent acute kidney injuries. We assume the key variables follow the chronological order as in Figure D.2.

Result

Our causal contrast of $E_P(R_2^{\bar{a}_1} \mid Z = z) - E_P(R_2 \mid Z = z)$ can be identified from the observed data of SPRINT, where $Z = z, z \in \{0, 1\}$ indicates the original standard or intensive blood pressure lowering therapy assignment, respectively. Note that we can safely disregard the censoring process since death is considered as a competing event to recurrent outcomes and because no participant experience lost to follow-up during our study period. Therefore, the PIE represents the difference in the expected number of acute kidney injury events if participants fully adhere to their prescribed blood pressure lowering medications versus the observed expected recurrent events in their assigned blood pressure lowering therapy arm. Applying the TMLE algorithms, we found that in the intensive blood pressure lowering therapy arm the estimated PIE is 0 with 95% CI [-0.004, 0.003], while for standard therapy arm the PIE is estimated to be 0.002 with 95% CI [-0.001, 0.005]. The standard error for the the confidence interval construction was obtained using 1,000 nonparametric bootstrap samples.

Overall, in the intensive blood pressure lowering therapy arm, the full adherence to prescribed blood pressure lowering medications has no impact on number of recurrent acute kidney injuries in the first four years of the study regardless of survival status; while full adherence in standard blood pressure lowering therapy arm the PIE effect is opposite but also insignificant. Since we are estimating the total effect which also includes the indirect treatment effect through mortality, the ACE




Figure D.2: A single world intervention graph (SWIG) that identifies the (conditional) population intervention effect and illustrates the relationship between the measured baseline variables L_0 , time-varying intervention of participants' adherence $\bar{A}_1$, the competing event of death $\bar{D}_2$, the confounders $\bar{L}_1$, and the (intermediate) recurrent outcome $\bar{Y}_2$, assuming there is no unmeasured confounding.

of medication adherence on acute kidney injury recurrences could be masked by the competing effect of death, especially in the standard blood pressure lowering therapy arm. This explanation is consistent with the observation mentioned in Chapter 1 of lower rates of acute kidney injury but higher incidence of all-cause mortality found in the standard blood pressure lowering therapy arm than in the intensive arm.

To further address the competing effect of death on the causal relationship between medication adherence and the recurrences of acute kidney injuries, one could estimate the separable effect of adherence on the acute kidney recurrences only, while controlling for the separable effect of adherence on death [48, 94]. The details of this method are beyond the scope of the current topic, and we will discuss it in Chapter 5 as our future work.

D.4.3 Summery Table for Results

Estimation results of the research questions discussed above are summarized in Table [D.1](#).

Section	Estimand	Effect Type	Estimation Results
3.4.4	average causal effect of intensive vs standard therapy	ITT, CDE	0.021 [0.012, 0.030]
D.4.1	average causal effect of intensive vs standard therapy	ITT, TE	0.019 [0.011, 0.027]
4.4.3	average causal effect of intensive vs standard therapy	PP, TE	0.017 [0.008, 0.026]
D.4.2	population intervention effect of adherence in standard therapy	TE	0.002 [-0.001, 0.005]
D.4.2	population intervention effect of adherence in intensive therapy	TE	0 [-0.004, 0.003]

ITT: intention-to-treat

PP: per-protocol

ACE: average causal effect

PIE: population intervention effect

CDE: controlled direct effect

TE: total effect

Table D.1: Table of estimated causal effects with the SPRINT