

UC San Diego

UC San Diego Electronic Theses and Dissertations

Title

Estimating Treatment Effects with Time-to-Event Outcomes: Methods and Applications in Dementia and Inflammatory Bowel Disease

Permalink

<https://escholarship.org/uc/item/8sm7w6b5>

ISBN

9798190926422

Author

Qi, Yuchen

Publication Date

2026-09-11

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA SAN DIEGO

Estimating Treatment Effects with Time-to-Event Outcomes: Methods and Applications in
Dementia and Inflammatory Bowel Disease

A dissertation submitted in partial satisfaction of the
requirements for the degree Doctor of Philosophy

in

Biostatistics

by

Yuchen Qi

Committee in charge:

Professor Ronghui Xu, Chair
Professor Sonia Jain
Professor Xin Tu
Professor Daren Wang

Copyright

Yuchen Qi, 2026

All rights reserved.

The Dissertation of Yuchen Qi is approved, and it is acceptable in quality and form for publication on microfilm and electronically.

University of California San Diego

2026

TABLE OF CONTENTS

Dissertation Approval Page	iii
Table of Contents	iv
List of Figures	vi
List of Tables	vii
Acknowledgements	viii
Vita	ix
Abstract of the Dissertation	x
Chapter 1 Introduction and Data Resources	1
1.1 Introduction	1
1.2 Honolulu–Asia Aging Study	2
1.3 OptumLabs® Data	4
1.4 Aims and Organization of this Dissertation	5
Chapter 2 Martingale R-learner: Estimating Time-varying Heterogeneous Treatment Effects for Time-to-event Outcomes	7
2.1 Introduction	7
2.2 Model and Method	9
2.2.1 Notation, Model, and Assumptions	9
2.2.2 R-Learner Martingale Residual and Orthogonal Functional Score	11
2.3 Estimation Procedure	14
2.4 Oracle Property of the Functional Score	16
2.5 Numerical Experiments	18
2.6 Discussion	25
Chapter 3 Applying the Martingale R-learner to the Honolulu–Asia Aging Study	28
3.1 Introduction	28
3.2 Study Population and Variables	29
3.3 Statistical Analysis	30
3.4 Estimated Heterogeneous Treatment Effects	32
3.5 HTE-based Stratification	34
3.6 Discussion	34
Chapter 4 Prior Exposure to Advanced Therapy and Risk of Serious Infections in Inflammatory Bowel Disease	37
4.1 Introduction	37
4.2 Data Source and Study Population	39

4.3	Exposure, Outcome, and Covariates	39
4.4	Statistical Analysis	40
4.5	Cohort Characteristics	42
4.6	Prior Advanced Therapy Exposure and Serious Infection Risk	43
4.7	Timing of Prior Advanced Therapy Discontinuation	46
4.8	Short-term Infection Risk and Corticosteroid Use	47
4.9	Discussion	48
Appendix A Proofs for Chapter 2		51
A.1	Derivation of the Risk-set Propensity Score and Marginal Hazard	51
A.2	Proof of Lemma 1	52
A.3	Proof of Lemma 2	53
A.4	Proof of Theorem 1	56
Bibliography		58

LIST OF FIGURES

Figure 2.1.	Distribution of root MSE under scenarios A - E.	24
Figure 2.2.	Ratio of MSE between Martingale R-learners with estimated and oracle nuisance models under scenarios A - E.	26
Figure 3.1.	Estimated heterogeneous treatment effect surfaces for midlife heavy drinking in the HAAS application.	33
Figure 3.2.	Kaplan–Meier curves for cognitive-impairment-free survival stratified by tertiles of the estimated 10-year cumulative HTE. Participants were divided into low, middle, and high groups based on $\widehat{\mathcal{T}}(10, \mathbf{Z}_i)$. The Kaplan–Meier curves are descriptive and are not adjusted for measured confounding. ...	35
Figure 4.1.	Standardized mean differences of baseline covariates before and after inverse probability of treatment weighting for prior advanced therapy exposure. Dashed vertical lines indicate absolute standardized mean difference of 0.1.	45

LIST OF TABLES

Table 3.1.	Characteristics of HAAS participants included in the Martingale R-learner analysis. Baseline was HAAS Exam 4. The event was defined as CASI score below 74 or mortality, whichever occurred first. Columns 0 and 1 denote non-heavy drinkers and midlife heavy drinkers, respectively.	30
Table 4.1.	Baseline characteristics of patients with IBD initiating a new advanced therapy, by prior advanced therapy exposure within the preceding 12 months.	43
Table 4.2.	Association between prior advanced therapy exposure and serious infection risk after initiation of a new advanced therapy.	45
Table 4.3.	Association between timing of prior advanced therapy discontinuation and serious infection risk among patients with prior advanced therapy exposure.	47
Table 4.4.	Short-term serious infection risk and corticosteroid use among patients with prior advanced therapy exposure.	48

ACKNOWLEDGEMENTS

I would like to first thank my chair, Professor Ronghui (Lily) Xu, for her guidance, support, and mentorship throughout my doctoral training. I am deeply grateful to Marquis for his guidance and collaboration on the methodological work that became an important part of this dissertation. I would also like to thank Dr. Siddharth (Sid) Singh for his support, mentorship, and the opportunity to work together on clinical research projects that greatly enriched my doctoral experience.

I would also like to thank my mom, Chuntao Chen, and my dad, Yonglin Qi, for their unconditional love, encouragement, and support throughout my education and all the years that led to the completion of this dissertation.

Chapters 2 and 3, in part, contain material that is currently submitted for publication. Jue Hou; Yuchen Qi; Ronghui Xu. The dissertation author was a coauthor of this material.

Chapter 4, in part, is a reprint of the material as it appears in *Journal of Crohn's and Colitis*. Virginia Solitano; Dhruv Ahuja; Yuchen Qi; Sagar B. Patel; Marianne Hupé; Sudheer K. Vuyyuru; Anna L. Silverman; Ronghui Xu; Shane W. Goodwin; Ola Olén; Ashwin N. Ananthakrishnan; Vipul Jairath; Siddharth Singh, 2026. The dissertation author was a coauthor of this paper.

VITA

- 2014–2018 Bachelor of Science in Mathematics, The Chinese University of Hong Kong, Shatin, Hong Kong
- 2019–2021 Master of Science in Biostatistics, Columbia University in the City of New York, New York, NY
- 2021–2026 Doctor of Philosophy in Biostatistics, University of California San Diego, La Jolla, CA

ABSTRACT OF THE DISSERTATION

Estimating Treatment Effects with Time-to-Event Outcomes: Methods and Applications in
Dementia and Inflammatory Bowel Disease

by

Yuchen Qi

Doctor of Philosophy in Biostatistics

University of California San Diego, 2026

Professor Ronghui Xu, Chair

Treatment effects in biomedical studies are often evaluated using time-to-event outcomes, such as disease onset, treatment failure, infection, or mortality. In many clinical and epidemiologic settings, these effects may vary across patient characteristics, and treatment decisions must also account for risks observed in real-world data. This dissertation develops and applies causal survival analysis methods for estimating treatment effects in precision medicine and clinical decision-making.

The first part introduces a martingale R-learner for estimating heterogeneous treatment effects in survival data. We develop a functional score framework that extends estimating equations

for survival outcomes to nonparametric heterogeneous treatment effects and generalizes Neyman orthogonality to this setting. The proposed approach decomposes conditional martingale residuals into components involving the risk-set propensity score and marginal martingale, reducing sensitivity to nuisance estimation. This framework allows flexible machine learning estimation of nuisance functions while preserving favorable theoretical and empirical performance. Numerical studies demonstrate that the proposed method performs well relative to existing approaches.

The second part applies this framework to the Honolulu–Asia Aging Study to examine whether the effect of midlife alcohol exposure on dementia risk varies across baseline characteristics, including age, education, APOE genotype, and physical function. This application illustrates how flexible heterogeneous treatment effect estimation can provide insight into individualized risk and potential effect modification in aging research.

The third part studies a real-world clinical question in inflammatory bowel disease: whether prior exposure to advanced therapies and the timing of therapy switching are associated with serious infection risk after initiation of a new advanced therapy. Using a large claims-based cohort, we apply survival models and inverse probability of treatment weighting to estimate treatment-related infection risk while addressing measured confounding. The analysis suggests that prior advanced therapy exposure is associated with a modestly higher risk of serious infection, whereas the timing of discontinuation is not strongly associated with infection risk. Together, these three parts show how causal survival analysis can support both methodological development and practical evidence generation in biomedical research.

Chapter 1

Introduction and Data Resources

1.1 Introduction

Time-to-event outcomes are central to biomedical research. In clinical trials, observational cohorts, and real-world evidence studies, investigators are often interested not only in whether an outcome occurs, but also in when it occurs. Examples include onset of dementia, disease progression, treatment failure, serious infection, hospitalization, and mortality. These outcomes are naturally analyzed using survival methods because they are subject to censoring, delayed occurrence, and changing risk over follow-up. At the same time, many scientific questions in medicine are causal: researchers seek to understand the effect of an exposure or treatment on future risk, rather than merely describe associations.

Causal inference for survival outcomes presents several methodological challenges. First, treatment or exposure assignment is often not randomized in observational data, so adjustment for measured confounding is required. Second, censoring and competing clinical processes may affect which individuals remain under observation and at risk. Third, treatment effects may vary over time and across patient characteristics. A single average treatment effect may be useful for population-level summaries, but it may be insufficient for precision medicine, where treatment decisions are made for individuals with different genetic, demographic, behavioral, and clinical profiles.

This dissertation focuses on causal survival analysis for estimating treatment effects in

biomedical studies. The central methodological contribution is the development of a martingale R-learner for estimating time-varying heterogeneous treatment effects for time-to-event outcomes. The proposed framework extends ideas from causal machine learning to survival data by using martingale residuals, risk-set propensity scores, and Neyman orthogonal functional scores. This approach is designed to reduce sensitivity to nuisance function estimation while allowing flexible modeling of treatment effect heterogeneity.

The dissertation also includes two applied components that motivate and illustrate the methodological issues. The first application uses data from the Honolulu–Asia Aging Study to study the effect of alcohol exposure on dementia risk and to evaluate whether this effect varies across baseline characteristics. This application emphasizes heterogeneous treatment effect estimation in a longitudinal aging cohort. The second application uses OptumLabs® Data to study serious infection risk among patients with inflammatory bowel disease initiating advanced therapies. This application emphasizes real-world evidence, treatment sequencing, and inverse probability weighting for causal survival analysis in administrative claims data.

Together, these projects address a common theme: how to estimate and interpret treatment or exposure effects when the outcome is time-to-event, the data are observational, and clinical decisions require either individualized or population-level causal evidence.

1.2 Honolulu–Asia Aging Study

The Honolulu–Asia Aging Study (HAAS) is a longitudinal epidemiologic study of cognitive impairment, dementia, and aging-related outcomes among Japanese-American men. HAAS was developed as an extension of the Honolulu Heart Program, a prospective cohort study that began in 1965 with Japanese-American men living on the island of Oahu, Hawaii (Syme et al., 1975; National Heart, Lung, and Blood Institute, 2024). The original Honolulu Heart Program was designed to study cardiovascular disease and stroke, but its long follow-up and extensive collection of midlife and late-life risk factors later made it a valuable resource for

studying aging and dementia.

The first HAAS examination was conducted in 1991–1993 among surviving members of the original Honolulu Heart Program cohort. Participants were elderly Japanese-American men, approximately 71 to 93 years of age at the start of the HAAS dementia examinations (White et al., 1996; Gelber et al., 2012). HAAS collected information on cognitive function, clinical diagnoses of dementia, mortality, genetic factors such as apolipoprotein E genotype, demographic and socioeconomic factors, behavioral exposures, and physical function measures. Cognitive function was assessed longitudinally, and dementia outcomes were adjudicated using standardized clinical procedures in prior HAAS studies (White et al., 1996; Gelber et al., 2012).

In this dissertation, HAAS is used to evaluate the effect of alcohol exposure on the risk of dementia. Alcohol exposure is a scientifically interesting and methodologically challenging exposure because its relationship with dementia may depend on age, genetic susceptibility, education, baseline health, and physical function. In addition, dementia is a time-to-event outcome with long latency, and mortality is common in elderly cohorts. These features make HAAS a natural setting for studying time-varying and heterogeneous treatment effects in survival analysis.

The HAAS application focuses on estimating how the effect of alcohol exposure varies across baseline characteristics, including age, education, APOE genotype, and grip strength. These variables are clinically meaningful because they may reflect demographic, genetic, socioeconomic, and functional dimensions of dementia risk. From a methodological perspective, the application illustrates how flexible heterogeneous treatment effect methods can move beyond a single average effect and instead characterize how the exposure-outcome relationship changes across subgroups and covariate profiles.

The analysis also highlights practical issues that commonly arise in cohort studies of aging. Follow-up may be incomplete due to censoring or death, event times may be observed over long time horizons, and covariates may play different roles as confounders, prognostic factors, or effect modifiers. The martingale R-learner developed in this dissertation is motivated

by precisely this type of setting: observational survival data in which treatment effects may vary over time and across baseline characteristics.

1.3 OptumLabs[®] Data

The second data resource used in this dissertation is OptumLabs[®] Data. OptumLabs[®] Data Warehouse is a large real-world data source containing de-identified longitudinal health information, including administrative claims for commercially insured and Medicare Advantage enrollees (OptumLabs, 2019; Li et al., 2020). The claims component includes medical and pharmacy claims, such as physician, facility, and prescription drug claims, submitted during routine health care delivery (Li et al., 2020). These data are well suited for studying treatment utilization, treatment switching, safety outcomes, and health care events that are captured in claims, such as hospitalizations, procedures, and medication dispensations.

Administrative claims data have several advantages for real-world evidence research. They can provide large sample sizes, broad geographic coverage, longitudinal enrollment information, and detailed records of reimbursed health care encounters. These features are particularly useful for studying uncommon safety outcomes and treatment patterns that may be difficult to evaluate in smaller clinical cohorts. However, claims data also have important limitations. They are collected for billing rather than research, may lack detailed clinical measures, and are subject to outcome and covariate misclassification. Therefore, causal analyses using claims data require careful cohort construction, exposure definition, outcome ascertainment, and adjustment for measured confounding.

In this dissertation, OptumLabs[®] Data are used to study serious infection risk among patients with inflammatory bowel disease (IBD) initiating advanced therapies. Advanced therapies, including biologic and small-molecule therapies, are increasingly used for IBD management. In clinical practice, patients may switch from one advanced therapy to another due to nonresponse, loss of response, adverse events, or other clinical considerations. The timing of switching is

clinically important because short washout intervals may raise concerns about overlapping immunosuppression, whereas longer intervals may risk worsening disease activity or corticosteroid exposure.

The applied question is whether prior exposure to advanced therapies, and the timing of discontinuation of the prior therapy before initiation of a new therapy, are associated with serious infection risk after initiation of the new therapy. The analysis uses a large cohort of treatment episodes among patients with IBD initiating advanced therapies. Survival models are used to evaluate time-to-serious infection after treatment initiation. In addition, inverse probability of treatment weighting is used to address measured confounding when estimating the effect of prior advanced therapy exposure on infection risk.

This application complements the HAAS analysis in two ways. First, it focuses on a different type of real-world survival outcome: a treatment-related safety event rather than dementia onset. Second, it emphasizes population-level causal effect estimation using weighting rather than individualized heterogeneous effect estimation. Together, the two applications show how causal survival analysis can be used both for precision medicine questions and for real-world evidence questions about treatment safety and sequencing.

1.4 Aims and Organization of this Dissertation

The overall goal of this dissertation is to develop and apply causal survival analysis methods for treatment effect estimation in biomedical studies. The dissertation is organized around three specific aims.

Aim 1 is to develop a martingale R-learner for estimating time-varying heterogeneous treatment effects in survival data. This aim addresses the methodological challenge of estimating flexible treatment effect functions when outcomes are subject to censoring and the nuisance functions, such as survival probabilities and risk-set treatment mechanisms, must be estimated from data. The proposed method uses a functional score framework and Neyman orthogonality

to reduce sensitivity to nuisance estimation and to allow flexible machine learning methods to be used in survival treatment effect estimation.

Aim 2 is to apply the martingale R-learner to the Honolulu–Asia Aging Study to investigate heterogeneity in the effect of alcohol exposure on dementia risk. This aim demonstrates how the proposed method can be used to examine effect modification over clinically meaningful baseline characteristics. The HAAS analysis provides an example of precision medicine and aging research in which individualized risk patterns and treatment effect heterogeneity are of primary scientific interest.

Aim 3 is to evaluate serious infection risk associated with prior advanced therapy exposure and treatment-switching timing among patients with inflammatory bowel disease using OptumLabs® Data. This aim applies causal survival analysis and inverse probability weighting to a real-world evidence question in treatment safety. The analysis estimates whether prior advanced therapy exposure is associated with subsequent serious infection risk and whether the timing of discontinuation of the prior therapy modifies that risk.

The remainder of the dissertation is organized as follows. Chapter 2 presents the martingale R-learner methodology, including the estimand, functional score construction, nuisance functions, estimation procedure, theoretical properties, and numerical experiments. Chapter 3 presents the HAAS application to alcohol exposure and dementia risk, with emphasis on time-varying heterogeneous treatment effects across demographic, genetic, and functional characteristics. Chapter 4 presents the OptumLabs® Data application on advanced therapy exposure, treatment switching, and serious infection risk in inflammatory bowel disease.

Chapter 2

Martingale R-learner: Estimating Time-varying Heterogeneous Treatment Effects for Time-to-event Outcomes

2.1 Introduction

A critical question in epidemiology and biomedical research is understanding how an individual's response to treatment or exposure may vary according to personal characteristics. For example, when evaluating the effect of midlife heavy drinking on time to cognitive impairment in the Honolulu–Asia Aging Study, the effect may vary by age, education, APOE genotype, and physical function. This variability in treatment effects, known as heterogeneous treatment effects (HTE), is central to clinical decision-making at the individual level.

In this chapter, we focus on causal effects for time-to-event outcomes. For such outcomes, HTE can be formalized as contrasts in survival functions or hazard functions conditional on baseline covariates under alternative treatment assignments. Let $T^{(d)}$ denote the potential event time under treatment d , where $d = 0, 1$, and let $\mathbf{Z}$ denote baseline covariates. We consider a time-varying HTE that can be interpreted through the difference in cumulative hazards or, equivalently, through an additive hazard representation.

Structural nested models provide a natural framework for modeling such causal contrasts. Robins (1986, 1989); Robins et al. (1992) developed structural nested models, including structural

nested mean models for mean contrasts (Robins, 1994, 2000). For survival outcomes, structural nested cumulative failure time models have been used to model contrasts in cumulative failure probabilities or survival-related quantities (Robins and Hernan, 2008; Picciotto et al., 2012). Marginal structural models are another common class of causal models, often used to estimate average treatment effects through inverse probability weighting (Robins, 2000). Although marginal structural models can examine heterogeneity indirectly by including interactions between treatment and effect modifiers, they typically model marginal effects averaged over other covariates. By contrast, structural nested models target the causal contrast directly through a blip or treatment-effect function, making them particularly suitable for heterogeneous treatment effect estimation (Vansteelandt and Joffe, 2014).

Considerable effort has been devoted to developing methods for HTE estimation, including meta-learners such as S-, T-, X-, M-, and R-learners (Athey and Imbens, 2016; Künzel et al., 2019; Nie and Wager, 2021), as well as machine learning and nonparametric approaches such as lasso (Imai and Ratkovic, 2013), random forests (Wager and Athey, 2018), boosting (Powers et al., 2018), neural networks (Shalit et al., 2017), and Bayesian additive regression trees (Hill, 2011; Hahn et al., 2020). Theoretical guarantees for HTE estimation have been studied for uncensored outcomes (Nie and Wager, 2021; Foster and Syrgkanis, 2023; Kennedy, 2023). However, these methods and guarantees do not immediately generalize to censored survival outcomes.

In healthcare applications, survival data with right censoring are often of primary interest. Recent work has extended meta-learners and causal machine learning methods to time-to-event outcomes (Xu et al., 2023; Jaroszewicz and Rzepakowski, 2014; Zhang et al., 2017; Henderson et al., 2020; Tabib and Larocque, 2020; Zhu and Gallego, 2020; Cui et al., 2023; Curth et al., 2021; Hu et al., 2021; Chapfuwa et al., 2021; Schrod et al., 2022; Abraich et al., 2022). However, relatively few methods provide theoretical guarantees for HTE estimation in censored survival settings. Existing approaches include censoring unbiased transformations for adapting HTE learners to time-to-event outcomes (Xu et al., 2024) and causal survival forests (Cui et al., 2023).

Moreover, methods for estimating time-varying HTEs remain less developed.

In this chapter, we propose a martingale R-learner for estimating time-varying heterogeneous treatment effects in survival data. The method is analogous in spirit to the R-learner of Nie and Wager (2021), but is built around counting-process martingales and risk sets. The proposed learner incorporates two nuisance components: a marginal cumulative hazard function for the factual survival outcome and a time-varying propensity score within risk sets. The resulting functional score is constructed to be Neyman orthogonal, reducing first-order sensitivity to nuisance estimation. Computationally, the score can be converted into a weighted least squares problem using pseudo-outcomes and pseudo-covariates, allowing flexible basis expansion and penalized regression.

The remainder of this chapter is organized as follows. Section 2.2 introduces the notation, causal estimand, and identifying assumptions. Then it constructs the R-learner martingale residual and the orthogonal functional score. Section 2.3 presents the estimation procedure. Section 2.4 gives a general oracle-property result for functional scores. Section 2.5 presents numerical experiments, and Section 2.6 concludes.

2.2 Model and Method

2.2.1 Notation, Model, and Assumptions

Let $T^{(d)}$ denote the potential event time that would be observed if a subject received treatment d , where $d = 0, 1$. Let $\mathbf{Z}$ denote a p -dimensional vector of baseline covariates, and let $D \in \{0, 1\}$ denote the observed treatment assignment at baseline.

We consider the following structural nested cumulative failure time model:

$$\frac{\mathbb{P}\{T^{(d)} \geq t \mid \mathbf{Z}\}}{\mathbb{P}\{T^{(0)} \geq t \mid \mathbf{Z}\}} = \exp\{-d \mathcal{T}(t, \mathbf{Z})\}, \quad (2.1)$$

where $\mathcal{T}(t, \mathbf{Z})$ is the cumulative time-varying heterogeneous treatment effect. Using the rela-

tionship

$$\mathbb{P}\{T^{(d)} \geq t \mid \mathbf{Z}\} = \exp\{-\Lambda(t; \mathbf{Z}, d)\},$$

where $\Lambda(t; \mathbf{Z}, d)$ is the cumulative hazard under treatment d , model (2.1) can be written as

$$\Lambda(t; \mathbf{Z}, d) = \Lambda(t; \mathbf{Z}, 0) + d \mathcal{T}(t, \mathbf{Z}). \quad (2.2)$$

This connects directly to the additive hazard model (Aalen, 1989),

$$\lambda(t; \mathbf{Z}, d) = \lambda(t; \mathbf{Z}, 0) + d \tau(t, \mathbf{Z}), \quad (2.3)$$

where

$$\mathcal{T}(t, \mathbf{Z}) = \int_0^t \tau(u, \mathbf{Z}) du.$$

Thus, $\tau(t, \mathbf{Z})$ is the instantaneous time-varying HTE on the hazard-difference scale, and $\mathcal{T}(t, \mathbf{Z})$ is the corresponding cumulative hazard difference.

Let $C^{(d)}$ denote the potential censoring time under treatment d . Let T and C denote the observed event and censoring times, respectively. Define

$$X = \min(T, C), \quad \delta = I(T \leq C).$$

The corresponding counting process and at-risk process are

$$N(t) = I(X \leq t, \delta = 1), \quad Y(t) = I(X \geq t).$$

Throughout this chapter, a subscript 0 denotes the true value of a parameter or function.

To identify $\tau(t, \mathbf{Z})$, we use the following assumptions.

Assumption A. SUTVA. There are no multiple versions of treatment and no interference between units (Rubin, 1980).

Assumption B. Unconfoundedness.

$$\{T^{(0)}, T^{(1)}\} \perp D \mid \mathbf{Z}.$$

Assumption C. Consistency.

$$T = T^{(D)}.$$

Assumption D. Risk-set positivity.

$$0 < \rho(t, \mathbf{Z}) \equiv \mathbb{P}(D = 1 \mid T \geq t, \mathbf{Z}) < 1$$

almost surely for all t and $\mathbf{Z}$ in the support of the data.

Assumption E. Non-informative censoring.

$$C^{(d)} \perp T^{(d)} \mid \mathbf{Z}, \quad d = 0, 1.$$

These assumptions are standard in causal survival analysis under the potential outcome framework. SUTVA and consistency link observed outcomes to well-defined counterfactual outcomes. Unconfoundedness requires that all baseline predictors of both treatment and the event time are included in $\mathbf{Z}$. Non-informative censoring ensures that censoring does not introduce bias in estimating treatment effects. The risk-set positivity assumption is specific to the proposed construction because the nuisance treatment mechanism is defined within the risk set at each time point, rather than only at baseline.

2.2.2 R-Learner Martingale Residual and Orthogonal Functional Score

When estimating a target treatment effect in the presence of nuisance functions, regularization bias from nuisance estimation can affect the target estimator. Neyman orthogonality is a key device for reducing this sensitivity. For a functional score $\psi(\tau; g)$, where τ is the target

function and g denotes nuisance functions, we say the score is Neyman orthogonal if

$$D_g \mathbb{E}\{\psi(\tau_0; g_0)\}[g - g_0] = 0, \quad (2.4)$$

where D_g denotes the directional derivative with respect to the nuisance functions. Equation (2.4) means that the score is locally insensitive to first-order perturbations of the nuisance functions.

For survival outcomes, achieving such orthogonality is more challenging than in uncensored regression settings because the at-risk process is time-dependent and related to treatment assignment. Inspired by the R-learner for uncensored outcomes (Nie and Wager, 2021), we construct a martingale residual decomposition that residualizes both the outcome process and treatment assignment within the risk set.

Define the marginal hazard of the factual event time conditional on $\mathbf{Z}$:

$$\lambda_m(t; \mathbf{Z}) = \lim_{\Delta t \rightarrow 0} \frac{\mathbb{P}(T \leq t + \Delta t \mid T \geq t, \mathbf{Z})}{\Delta t}.$$

Under the additive hazard model and conditional ignorability,

$$\lambda_m(t; \mathbf{Z}) = \lambda(t; \mathbf{Z}, 0) + \tau(t, \mathbf{Z})\rho(t, \mathbf{Z}), \quad (2.5)$$

where

$$\rho(t, \mathbf{Z}) = \mathbb{P}(D = 1 \mid T \geq t, \mathbf{Z})$$

is the risk-set propensity score (Kang et al., 2018). Let

$$\pi(\mathbf{Z}) = \mathbb{P}(D = 1 \mid \mathbf{Z})$$

denote the baseline propensity score. Under model (2.3), the risk-set propensity score can be

written as

$$\rho(t, \mathbf{Z}) = \frac{\pi(\mathbf{Z}) \exp \left\{ - \int_0^t \tau(u, \mathbf{Z}) du \right\}}{1 + \left[\exp \left\{ - \int_0^t \tau(u, \mathbf{Z}) du \right\} - 1 \right] \pi(\mathbf{Z})}. \quad (2.6)$$

The following lemma gives the martingale residual decomposition underlying the proposed learner.

Lemma 1. *Under equations (2.3), (2.5), and (2.6),*

$$dN(t) - Y(t)d\Lambda_m(t; \mathbf{Z}) = \tau(t, \mathbf{Z})Y(t)\{D - \rho(t, \mathbf{Z})\}dt + dM(t), \quad (2.7)$$

where

$$dM(t) = dN(t) - Y(t)\lambda(t; \mathbf{Z}, D)dt$$

is a martingale increment with respect to the natural filtration.

We refer to (2.7) as the *R-learner martingale residual*. The left-hand side is the residual from the marginal survival model, the term $\{D - \rho(t, \mathbf{Z})\}$ is the residualized treatment assignment within the risk set, and $dM(t)$ is a mean-zero martingale error term.

Based on this decomposition, define the functional score for τ by

$$\begin{aligned} \psi(\tau; \Lambda_m, \rho)[\phi] &= \int_0^{t_{\max}} \phi(t, \mathbf{Z}) \{D - \rho(t, \mathbf{Z})\} [dN(t) - Y(t)d\Lambda_m(t; \mathbf{Z})] \\ &\quad - \int_0^{t_{\max}} \phi(t, \mathbf{Z}) Y(t) \{D - \rho(t, \mathbf{Z})\}^2 \tau(t, \mathbf{Z}) dt, \end{aligned} \quad (2.8)$$

where $\phi(t, \mathbf{Z})$ is an arbitrary test function and $t_{\max}$ is the maximum follow-up time.

Lemma 2. *The score function (2.8) satisfies*

$$\mathbb{E}\{\psi(\tau_0; \Lambda_{m,0}, \rho_0)[\phi]\} = 0$$

for all suitable test functions ϕ . Moreover, the score is Neyman orthogonal with respect to the nuisance functions (Λ_m, ρ) .

Lemma 2 shows that the proposed score is unbiased at the true nuisance functions and locally insensitive to nuisance perturbations. This property motivates the use of flexible nuisance estimators in the martingale R-learner. In the next section, we describe how the functional score can be converted into a computable estimation procedure using basis expansion, cross-fitting, and penalized least squares.

2.3 Estimation Procedure

Suppose we observe independent and identically distributed data

$$\{(\mathbf{Z}_i, D_i, X_i, \delta_i) : i = 1, \dots, n\}$$

satisfying Assumptions A–E. We approximate the treatment effect function using a basis expansion

$$\tau(t, \mathbf{Z}) = \boldsymbol{\theta}^\top \boldsymbol{\phi}(t, \mathbf{Z}),$$

where $\boldsymbol{\phi}(t, \mathbf{Z})$ is a vector of pre-specified basis functions of time and baseline covariates, and $\boldsymbol{\theta}$ is the coefficient vector.

After estimating the nuisance functions, the estimating equation for $\boldsymbol{\theta}$ is

$$\begin{aligned} 0 = & \frac{1}{n} \sum_{i=1}^n \int_0^{t_{\max}} \boldsymbol{\phi}(t, \mathbf{Z}_i) \{D_i - \widehat{\rho}(t, \mathbf{Z}_i)\} \left[dN_i(t) - Y_i(t) d\widehat{\Lambda}_m(t; \mathbf{Z}_i) \right] \\ & - \frac{1}{n} \sum_{i=1}^n \int_0^{t_{\max}} \boldsymbol{\phi}(t, \mathbf{Z}_i) Y_i(t) \{D_i - \widehat{\rho}(t, \mathbf{Z}_i)\}^2 \boldsymbol{\theta}^\top \boldsymbol{\phi}(t, \mathbf{Z}_i) dt. \end{aligned} \quad (2.9)$$

Solving (2.9) is equivalent to minimizing the quadratic loss

$$\begin{aligned} L(\boldsymbol{\theta}) = & \frac{1}{2} \boldsymbol{\theta}^\top \left[\frac{1}{n} \sum_{i=1}^n \int_0^{t_{\max}} Y_i(t) \{D_i - \widehat{\rho}(t, \mathbf{Z}_i)\}^2 \boldsymbol{\phi}(t, \mathbf{Z}_i) \boldsymbol{\phi}(t, \mathbf{Z}_i)^\top dt \right] \boldsymbol{\theta} \\ & - \boldsymbol{\theta}^\top \left[\frac{1}{n} \sum_{i=1}^n \int_0^{t_{\max}} \boldsymbol{\phi}(t, \mathbf{Z}_i) \{D_i - \widehat{\rho}(t, \mathbf{Z}_i)\} \left[dN_i(t) - Y_i(t) d\widehat{\Lambda}_m(t; \mathbf{Z}_i) \right] \right]. \end{aligned} \quad (2.10)$$

Let $0 = t_0 < t_1 < \dots < t_J = t_{\max}$ denote the ordered unique event times, and define

$$\Delta f(t_j) = f(t_j) - f(t_{j-1}).$$

Approximating the integrals over this grid gives a pseudo-data representation. For $i = 1, \dots, n$ and j such that $t_j \leq X_i$, define

$$\tilde{Y}_{ij} = \frac{\Delta N_i(t_j) - Y_i(t_j) \Delta \hat{\Lambda}_m(t_j; \mathbf{Z}_i)}{\Delta t_j},$$

$$\tilde{\mathbf{X}}_{ij} = \phi(t_j, \mathbf{Z}_i) \{D_i - \hat{\rho}(t_j, \mathbf{Z}_i)\},$$

and

$$\tilde{w}_{ij} = \frac{1}{2} Y_i(t_j) \Delta t_j.$$

Then the target estimation step can be written as the weighted least squares problem

$$\hat{\boldsymbol{\theta}} = \arg \min_{\boldsymbol{\theta}} \sum_{i=1}^n \sum_{j: t_j \leq X_i} \tilde{w}_{ij} \left(\boldsymbol{\theta}^\top \tilde{\mathbf{X}}_{ij} - \tilde{Y}_{ij} \right)^2. \quad (2.11)$$

We use cross-fitting to reduce overfitting bias from nuisance estimation. The data are divided into Q folds. For each fold, nuisance functions are estimated using the data outside the fold and then evaluated for subjects in the held-out fold. The final treatment effect estimator is

$$\hat{\tau}(t, \mathbf{Z}) = \hat{\boldsymbol{\theta}}^\top \phi(t, \mathbf{Z}).$$

Algorithm 1. Martingale R-learner estimation procedure

- 1: Divide the data into Q approximately equal folds $\mathcal{I}_1, \dots, \mathcal{I}_Q$.
- 2: **for** $q = 1, \dots, Q$ **do**
- 3: Using observations not in $\mathcal{I}_q$, estimate $\widehat{\Lambda}_m^{(-q)}(t; \mathbf{Z})$, $\widehat{\pi}^{(-q)}(\mathbf{Z})$, and $\widehat{S}^{(-q)}(t; \mathbf{Z}, D)$.
- 4: For $i \in \mathcal{I}_q$, compute

$$\widehat{\rho}^{(-q)}(t, \mathbf{Z}_i) = \frac{\widehat{\pi}^{(-q)}(\mathbf{Z}_i) \widehat{S}^{(-q)}(t; \mathbf{Z}_i, 1) / \widehat{S}^{(-q)}(t; \mathbf{Z}_i, 0)}{1 + \left[\widehat{S}^{(-q)}(t; \mathbf{Z}_i, 1) / \widehat{S}^{(-q)}(t; \mathbf{Z}_i, 0) - 1 \right] \widehat{\pi}^{(-q)}(\mathbf{Z}_i)}.$$

- 5: **end for**
 - 6: Construct pseudo-outcomes $\widetilde{Y}_{ij}$, pseudo-covariates $\widetilde{\mathbf{X}}_{ij}$, and weights $\widetilde{w}_{ij}$ using cross-fitted nuisance estimates.
 - 7: Estimate θ by minimizing the weighted ridge-penalized version of (2.11).
 - 8: Return $\widehat{\tau}(t, \mathbf{Z}) = \widehat{\theta}^\top \phi(t, \mathbf{Z})$.
-

2.4 Oracle Property of the Functional Score

This section presents a general oracle-property result for functional scores. The result is stated at the level of an abstract functional score and is not tied to a specific convergence-rate calculation for the martingale R-learner.

Let $\mathcal{H}$ denote the target function class and $\mathcal{G}$ denote the nuisance function space. Let $\tau_0 \in \mathcal{H}$ and $g_0 \in \mathcal{G}$ be the true target and nuisance functions. Let $\psi_i(\tau, g)[\phi]$ denote a functional score indexed by a test function ϕ . Within $\mathcal{H}$, define the inner product and induced norm

$$\langle \phi_1, \phi_2 \rangle_{\mathcal{H}} = \int_0^{t_{\max}} \mathbb{E}\{\phi_1(t, \mathbf{Z}) \phi_2(t, \mathbf{Z})\} dt, \quad \|\phi\|_{\mathcal{H}} = \sqrt{\langle \phi, \phi \rangle_{\mathcal{H}}}. \quad (2.12)$$

For estimation of τ_0 , define the pseudo-regret

$$\mathcal{R}(\tau, g) = \mathbb{E}\{\psi_i(\tau, g)[\tau - \tau_0]\} - \mathbb{E}\{\psi_i(\tau_0, g)[\tau - \tau_0]\}. \quad (2.13)$$

This quantity is analogous to the regret associated with a loss function, except that it is defined through the functional score rather than through a finite-dimensional likelihood score or loss gradient.

We use the following assumptions.

Assumption F. Neyman orthogonality. For all $\tau \in \mathcal{H}$,

$$D_g \mathbb{E}\{\psi_i(\tau_0, g_0)[\tau - \tau_0]\}[g - g_0] = 0.$$

Assumption G. Identification.

$$\mathbb{E}\{\psi_i(\tau_0, g_0)[\tau - \tau_0]\} = 0$$

for all τ in the relevant star hull of the target class around τ_0 .

Assumption H. Smoothness. There exist constants $\beta_1, \beta_2 > 0$ and $r \in [0, 1)$ such that the functional score is first-order smooth in τ and second-order smooth in g :

$$\mathbb{E}\{D_\tau \psi_i(\bar{\tau}, g_0)[\tau - \tau_0, \tau - \tau_0]\} \leq \beta_1 \|\tau - \tau_0\|_{\mathcal{H}}^2,$$

and

$$\left| \mathbb{E}\{D_g^2 \psi_i(\tau_0, \bar{g})[\tau - \tau_0, g - g_0, g - g_0]\} \right| \leq \beta_2 \|\tau - \tau_0\|_{\mathcal{H}}^{1-r} \|g - g_0\|_{\mathcal{G}}^2.$$

Assumption I. Local convexity. There exist constants $\lambda, \kappa > 0$ such that

$$\mathbb{E}\{D_\tau \psi_i(\bar{\tau}, g)[\tau - \tau_0, \tau - \tau_0]\} \geq \lambda \|\tau - \tau_0\|_{\mathcal{H}}^2 - \kappa \|g - g_0\|_{\mathcal{G}}^{4/(1+r)}.$$

Let $\text{Rate}(\mathcal{G}, S, \delta)$ denote the nuisance estimation rate on sample S with probability at least $1 - \delta$. Let

$$\text{Rate}(\mathcal{H}, S, \delta; \hat{\tau}, \hat{g})$$

denote the corresponding target estimation rate when the target algorithm is applied with nuisance function $\widehat{g}$.

Theorem 1. *Suppose Assumptions F–I hold. With sample splitting, the estimator $\widehat{\tau}$ satisfies, with probability at least $1 - \delta$,*

$$\|\widehat{\tau} - \tau_0\|_{\mathcal{H}} \lesssim \text{Rate}(\mathcal{H}, S_2, \delta/2; \widehat{\tau}, \widehat{g}) + \text{Rate}(\mathcal{G}, S_1, \delta/2)^{2/(1+r)}. \quad (2.14)$$

Theorem 1 states that, under orthogonality and smoothness conditions, nuisance estimation affects the target estimator only through a higher-order term. When $r = 0$, the nuisance contribution is quadratic in the nuisance estimation rate. Therefore, if the nuisance functions are estimated at a rate faster than $n^{-1/4}$ and the target estimation rate is of order $n^{-1/2}$ in the appropriate norm, the nuisance estimation error does not dominate the target estimation error. This is the oracle property: the estimator behaves, to first order, as if the nuisance functions were known.

2.5 Numerical Experiments

We assessed the finite-sample performance of the Martingale R-learner in simulation studies. Based on simulation designs from related literature (Nie and Wager, 2021; Foster and Syrgkanis, 2023), we specified scenarios by the dimension of covariates p , baseline propensity score $\pi(\mathbf{Z}_i)$, heterogeneous treatment effect $\tau(t, \mathbf{Z}_i)$, and untreated hazard $\lambda(t; \mathbf{Z}_i, 0)$.

- A. $p = 1$, $\pi(Z_i) = 0.5$, $\tau(t, Z_i) = Z_i + Z_i^2 + t + t^2$, and $\lambda(t; Z_i, 0) = 0.5$.
- B. $p = 3$, $\pi(\mathbf{Z}_i) = \text{expit}(Z_{i,1} + Z_{i,2} + Z_{i,3})$, $\tau(t, \mathbf{Z}_i) = \sin(Z_{i,1} + Z_{i,2} + Z_{i,3}) + \sin(t) + 4$, and $\lambda(t; \mathbf{Z}_i, 0) = 4 + Z_{i,1}$.
- C. $p = 3$, $\pi(\mathbf{Z}_i) = \text{expit}(Z_{i,1})$, $\tau(t, \mathbf{Z}_i) = 4 + \log\{1 + \exp(Z_{i,2} + Z_{i,3})\} + t$, and $\lambda(t; \mathbf{Z}_i, 0) = 3 + Z_{i,1}$.

D. $p = 3$, $\pi(\mathbf{Z}_i) = \text{expit}(-Z_{i,1} - Z_{i,2} - Z_{i,3})$, $\tau(t, \mathbf{Z}_i) = 6$, and $\lambda(t; \mathbf{Z}_i, 0) = \log\{1 + \exp(Z_{i,1} + Z_{i,2})\} + t^2$.

E. $p = 5$, $\pi(\mathbf{Z}_i) = \text{expit}(-Z_{i,1} - Z_{i,2} - Z_{i,3})$, $\tau(t, \mathbf{Z}_i) = \log\{6 + Z_{i,3} + Z_{i,4}\} + \sin(Z_{i,5}) + \sin(t)$, and $\lambda(t; \mathbf{Z}_i, 0) = 3 + Z_{i,1}$.

Here $\text{expit}(x) = 1/(1 + e^{-x})$. Across all scenarios, covariates were generated independently from $\text{Unif}(-2, 2)$, and censoring times were generated from $\text{Unif}(0, c_{\max})$, with $c_{\max}$ chosen to produce approximately 15% censoring. Scenario A represents randomized treatment assignment with a polynomial HTE and constant untreated hazard. Scenarios B–E include observational treatment assignment and more complex HTE structures involving trigonometric, logarithmic, or constant treatment effect functions. Scenario D includes a nonlinear time-varying untreated hazard.

For each scenario, we considered sample sizes $n = 500, 750, 1000$ to evaluate the trajectory of estimation performance with increasing sample size. For each simulation replicate, we generated one training data set from the specified data-generating mechanism using a predefined seed. For each simulation setting and sample size, we also generated one independent test data set of the same sample size and kept it fixed across all replicates. Each learner was fit on the training data set and used to estimate the treatment effect function at the test data.

In all simulation settings, we implemented the Martingale R-learner in R using the packages `polyspline`, `splines`, `gam`, and `caret`. We estimated the nuisance models under a cross-fitting scheme with $Q = 5$ folds. For each fold, nuisance models were estimated using the out-of-fold data, and the resulting fitted nuisance functions were used for subjects in the held-out fold when constructing pseudo-outcomes and performing target estimation.

Specifically, the cumulative hazard of the marginal survival model, $\widehat{\Lambda}_m^{(-q)}(t; \mathbf{Z}_i)$, was estimated using the `hare` function in the `polyspline` package, applied to the observed right-censored outcome (X_i, δ_i) with baseline covariates $\mathbf{Z}_i$ as predictors. Unless otherwise stated, `hare` was used with its default parameter settings. To estimate the risk-set propensity score $\rho(t, \mathbf{Z}_i)$, we

first estimated the baseline propensity score $\pi(\mathbf{Z}_i)$ using a generalized additive logistic regression model with binomial family and basis dimension equal to 3. We then estimated the conditional survival functions $S(t; \mathbf{Z}_i, 0)$ and $S(t; \mathbf{Z}_i, 1)$ using an additional hane model fit to (X_i, δ_i) with predictors $\mathbf{Z}_i$ and D_i , evaluated at $D_i = 0$ and $D_i = 1$, respectively. The risk-set propensity score was then computed as

$$\hat{\rho}^{(-q)}(t, \mathbf{Z}_i) = \frac{\hat{\pi}^{(-q)}(\mathbf{Z}_i) \hat{S}^{(-q)}(t; \mathbf{Z}_i, 1) / \hat{S}^{(-q)}(t; \mathbf{Z}_i, 0)}{1 + \left[\hat{S}^{(-q)}(t; \mathbf{Z}_i, 1) / \hat{S}^{(-q)}(t; \mathbf{Z}_i, 0) - 1 \right] \hat{\pi}^{(-q)}(\mathbf{Z}_i)}. \quad (2.15)$$

Using the cross-fitted nuisance estimates, we implemented the target-estimation step through the pseudo-data representation described in Section 2.3. We set the time grid $0 = t_0 < t_1 < \dots < t_J$ to be the ordered unique observed event times. The main computational cost of the Martingale R-learner arises from the pseudo-data expansion over ordered event times. When the number of distinct event times grows proportionally with sample size, the pseudo-data size is of order $O(n^2)$. Consequently, the computational cost of estimation increases approximately quadratically with sample size.

To approximate $\tau(t, \mathbf{Z}_i)$, we used natural spline basis expansions generated by the `ns` function in the `splines` package. In the simulation studies, we used the same degree of freedom for time and for each baseline covariate, and selected this degree of freedom from the candidate grid $\{5, 6, 7, 8, 9\}$. For Scenario A, with univariate Z_i , we approximated the HTE using pairwise interactions of a natural spline basis expansion for Z_i and time. For the remaining scenarios, we approximated the HTE using pairwise interactions of natural spline basis expansions for each baseline covariate and time. For each candidate basis specification, the design matrix included main spline basis terms and pairwise interaction terms among these basis expansions. An intercept term was included explicitly in the design matrix.

The coefficient vector θ was estimated by weighted ridge regression. Instead of fitting the expanded pseudo-data directly, we accumulated fold-specific components of the normal

equations. For the v th validation fold, define

$$G_v = \sum_{i \in \mathcal{I}_v} \sum_{j: t_j \leq X_i} \Delta t_j \tilde{\mathbf{X}}_{ij} \tilde{\mathbf{X}}_{ij}^\top,$$

and

$$h_v = \sum_{i \in \mathcal{I}_v} \sum_{j: t_j \leq X_i} \Delta t_j \tilde{\mathbf{X}}_{ij} \tilde{Y}_{ij}.$$

Computationally, these fold-specific sums were evaluated without explicitly storing the full expanded pseudo-data set. Within each validation fold $\mathcal{I}_v$, subjects were divided into smaller batches $\mathcal{B}_{v1}, \dots, \mathcal{B}_{vM_v}$. For each batch $\mathcal{B}_{vm}$, we computed

$$G_{vm} = \sum_{i \in \mathcal{B}_{vm}} \sum_{j: t_j \leq X_i} \Delta t_j \tilde{\mathbf{X}}_{ij} \tilde{\mathbf{X}}_{ij}^\top, \quad h_{vm} = \sum_{i \in \mathcal{B}_{vm}} \sum_{j: t_j \leq X_i} \Delta t_j \tilde{\mathbf{X}}_{ij} \tilde{Y}_{ij}.$$

The fold-specific components were obtained by summing the batch-level contributions:

$$G_v = \sum_{m=1}^{M_v} G_{vm}, \quad h_v = \sum_{m=1}^{M_v} h_{vm}.$$

Thus, batching gives the same G_v and h_v as the full pseudo-data formulation, but only requires storing the current batch-level quantities and the accumulated $p_\tau \times p_\tau$ matrix and p_τ -dimensional vector.

Let

$$G = \sum_{v=1}^5 G_v, \quad h = \sum_{v=1}^5 h_v.$$

The ridge penalty matrix was defined as

$$P = \text{diag}(0, 1, \dots, 1),$$

so that the intercept term was left unpenalized and all remaining coefficients were penalized equally. For each candidate basis specification and penalty parameter λ , we performed 5-fold

cross-validation using the fold-specific estimating-equation components. Let

$$G_{-v} = G - G_v, \quad h_{-v} = h - h_v.$$

The coefficient vector fitted without the v th validation fold was

$$\hat{\theta}_{-v}(\lambda) = \{G_{-v} + \lambda P\}^{-1} h_{-v}.$$

The validation criterion was

$$\text{CV}_v(\lambda) = \hat{\theta}_{-v}(\lambda)^\top G_v \hat{\theta}_{-v}(\lambda) - 2h_v^\top \hat{\theta}_{-v}(\lambda),$$

which corresponds to the validation estimating loss up to constants that do not depend on θ . The basis specification and penalty parameter were selected by minimizing

$$\sum_{v=1}^5 \text{CV}_v(\lambda)$$

over all candidate basis specifications and penalty parameters. The final estimator was computed using all observations:

$$\hat{\theta} = \{G + \hat{\lambda} P\}^{-1} h.$$

Predicted HTEs were obtained by evaluating the selected spline basis at the target time and covariate values and multiplying the resulting design matrix by $\hat{\theta}$.

As benchmark meta-learners, we implemented S-learner and T-learner using hazard regression through the `hazreg` function in the `polyspline` package. In both learners, the outcome (X_i, δ_i) was modeled directly, and the predicted HTE was obtained by taking the difference between predicted hazards under treatment and control. For the S-learner, we fit a single hazard regression to the full training data set with predictors $(\mathbf{Z}_i, D_i)$. Let $\hat{\lambda}_S(t; \mathbf{Z}_i, D_i)$ denote the fitted

hazard function from this model. For test data $(X_{*,i}, \mathbf{Z}_{*,i})$, we evaluated the fitted hazard twice, once with treatment fixed at $D_i = 1$ and once with treatment fixed at $D_i = 0$, and defined

$$\widehat{\tau}_S(X_{*,i}, \mathbf{Z}_{*,i}) = \widehat{\lambda}_S(X_{*,i}; \mathbf{Z}_{*,i}, 1) - \widehat{\lambda}_S(X_{*,i}; \mathbf{Z}_{*,i}, 0).$$

For the T-learner, we fit two separate hazard models, one in the treated subsample and one in the control subsample, each using only baseline covariates $\mathbf{Z}_i$ as predictors. Let $\widehat{\lambda}_{T,1}(t; \mathbf{Z}_i)$ and $\widehat{\lambda}_{T,0}(t; \mathbf{Z}_i)$ denote the fitted hazards from the treated and control models, respectively. We defined

$$\widehat{\tau}_T(X_{*,i}, \mathbf{Z}_{*,i}) = \widehat{\lambda}_{T,1}(X_{*,i}; \mathbf{Z}_{*,i}) - \widehat{\lambda}_{T,0}(X_{*,i}; \mathbf{Z}_{*,i}).$$

To assess oracle behavior, we also implemented the Martingale R-learner with oracle nuisance functions. To evaluate performance in a region with sufficient data support, we restricted evaluation to the subset of test observations with observed time less than the 85th percentile of the event time, denoted by $t_{0.85}$. Denote the test data by

$$\{(X_{*,i}, \mathbf{Z}_{*,i}) : i = 1, \dots, n_*\}.$$

The restricted MSE was

$$\text{MSE}(\widehat{\tau}, \tau_0) = \frac{\sum_{i=1}^{n_*} I(X_{*,i} \leq t_{0.85}) [\widehat{\tau}(X_{*,i}, \mathbf{Z}_{*,i}) - \tau_0(X_{*,i}, \mathbf{Z}_{*,i})]^2}{\sum_{i=1}^{n_*} I(X_{*,i} \leq t_{0.85})}.$$

Figure 2.1 reports the simulation results comparing the Martingale R-learner with estimated nuisance functions, the oracle Martingale R-learner, and the benchmark S- and T-learners. Across the five simulation scenarios and all sample sizes considered, the Martingale R-learner with estimated nuisance functions achieved lower median root MSE than the S- and T-learners. The improvement was particularly pronounced in Scenarios A and D, where the benchmark learners exhibited substantially larger RMSE and greater variability, whereas the Martingale

R-learner remained comparatively stable. In Scenarios B, C, and E, the differences were more moderate, but the same overall ranking was maintained.

As expected, the oracle Martingale R-learner generally achieved the lowest RMSE. The gap between the estimated- nuisance and oracle- nuisance learners tended to decrease as sample size increased, indicating that the effect of nuisance estimation error diminished with increasing sample size. This provides empirical support for the oracle behavior of the proposed learner. RMSE also generally decreased with sample size for all methods, while the relative advantage of the Martingale R-learner remained evident.




Figure 2.1. Distribution of root MSE under scenarios A - E.

Figure 2.2 shows ratios of MSE using estimated nuisance functions to MSE using oracle nuisance functions. These results provide an additional empirical check on the oracle

behavior of the Martingale R-learner. As sample size increases, the estimated-nuisance learner generally moves closer to the oracle-nuisance learner, indicating diminishing impact from nuisance estimation error.

2.6 Discussion

In this chapter, we proposed a martingale R-learner for estimating time-varying heterogeneous treatment effects for time-to-event outcomes. The method targets the treatment effect function directly under an additive hazard representation and uses an R-learner-type decomposition of martingale residuals. The proposed functional score incorporates two nuisance components: the marginal cumulative hazard and the risk-set propensity score. By residualizing both the outcome process and treatment assignment within risk sets, the method provides a flexible framework for estimating treatment effect heterogeneity over both time and baseline covariates.

The estimation procedure converts the functional score into a weighted least squares problem through discretization over event times and construction of pseudo-outcomes, pseudo-covariates, and risk-set weights. This formulation allows standard penalized regression or other machine learning tools to be used for estimating the treatment effect function. Cross-fitting is used so that nuisance functions and target treatment effects are estimated on separate folds, reducing overfitting bias from nuisance estimation.

The general oracle-property result shows that, under Neyman orthogonality and smoothness conditions, nuisance estimation contributes only higher-order error to the target estimator. This result motivates the use of flexible nuisance estimation methods while retaining stable target estimation. The chapter intentionally focuses on the general oracle-property framework and does not include the specific rate calculation for the martingale R-learner.

The simulation studies showed that the martingale R-learner performed favorably across settings with different propensity mechanisms, untreated hazards, and heterogeneous treatment




Figure 2.2. Ratio of MSE between Martingale R-learners with estimated and oracle nuisance models under scenarios A - E.

effect functions. Compared with S- and T-learners based on hazard regression, the martingale R-learner generally achieved lower root mean squared error. The comparison between estimated and oracle nuisance functions further suggested that the impact of nuisance estimation decreased with increasing sample size, providing empirical support for the stability of the proposed residualization strategy.

The next chapter applies the martingale R-learner to the Honolulu–Asia Aging Study to investigate heterogeneity in the effect of midlife heavy drinking on cognitive-impairment-free survival.

Acknowledgements

Chapters 2 and 3, in part, contain material that is currently submitted for publication. Jue Hou; Yuchen Qi; Ronghui Xu. The dissertation author was a coauthor of this material.

Chapter 3

Applying the Martingale R-learner to the Honolulu–Asia Aging Study

3.1 Introduction

The Honolulu–Asia Aging Study (HAAS) is a longitudinal study of dementia, cognitive decline, and aging-related outcomes among Japanese-American men (Peila et al., 2001). The study provides a useful setting for applying the Martingale R-learner because participants were followed over a long period, cognitive function was assessed repeatedly, and several baseline characteristics may modify the effect of midlife exposures on late-life cognitive outcomes.

In this chapter, we apply the Martingale R-learner developed in Chapter 2 to evaluate the effect of midlife heavy drinking on cognitive-impairment-free survival in HAAS. The exposure of interest is heavy alcohol consumption during midlife, defined as consuming at least 14 drinks per week at either Honolulu Heart Program Exam 1 or Exam 3. The outcome is time from HAAS Exam 4 to the first examination at which the participant had a Cognitive Abilities Screening Instrument (CASI) score below 74, or mortality, whichever occurred first.

This application is motivated by two features of the HAAS data. First, treatment effect heterogeneity is scientifically important. The effect of midlife heavy drinking on late-life cognitive outcomes may differ by baseline characteristics such as age, education, APOE genotype, and grip strength. Second, the timing of the outcome is important. HAAS participants were followed for many years, so the effect of heavy drinking may vary over follow-up time. These features

make HAAS a natural application for the Martingale R-learner, which estimates treatment effect heterogeneity as a function of both time and baseline covariates.

3.2 Study Population and Variables

The baseline for this analysis was HAAS Exam 4. The analysis included participants with available information on midlife heavy-drinking status, outcome and follow-up time, and the baseline covariates used for treatment effect estimation. After excluding observations with missing midlife heavy-drinking status, missing outcome or follow-up information, or missing covariates, the analytic data set contained 1874 participants.

The exposure was midlife heavy drinking, defined as at least 14 drinks per week at either HHP Exam 1 or Exam 3. Participants who did not meet this criterion were classified as non-heavy drinkers. Among the 1874 participants in the analytic cohort, 427 participants, or 22.8%, were classified as heavy drinkers during midlife.

The outcome was cognitive-impairment-free survival, defined as the time from Exam 4 to the first HAAS examination at which the participant had a CASI score below 74, or mortality, whichever occurred first. The CASI is a 0–100 cognitive assessment scale covering domains including attention, concentration, orientation, memory, language, visual construction, fluency, abstraction, and judgment. Death was ascertained through study surveillance and follow-up contacts, with death dates obtained from death certificates. Participants who remained alive with all observed CASI scores above 74 were censored at their last HAAS examination with an available CASI assessment.

The baseline covariates used in the analysis were age at Exam 4, education, APOE genotype, and grip strength. Education was dichotomized as ≤ 12 years versus > 12 years. APOE genotype was coded as APOE E4 positive versus APOE E4 negative. Grip strength was measured in kilograms of force and defined as the maximum over six grip measurements at Exam 4, consisting of three dominant-hand and three opposite-hand measurements.

Table 3.1 summarizes the baseline characteristics and outcome information for the analytic cohort, overall and by midlife heavy-drinking status. Participants were between 71.3 and 91.0 years old at baseline, with an average age of 77.1 years. The mean observed follow-up time was 3240 days, approximately 8.9 years. Overall, 47.1% of participants experienced the composite event of CASI score below 74 or death during follow-up.

Table 3.1. Characteristics of HAAS participants included in the Martingale R-learner analysis. Baseline was HAAS Exam 4. The event was defined as CASI score below 74 or mortality, whichever occurred first. Columns 0 and 1 denote non-heavy drinkers and midlife heavy drinkers, respectively.

	0 (N = 1447)	1 (N = 427)	Overall (N = 1874)
Observed time, days			
Mean (SD)	3360 (1860)	2830 (1620)	3240 (1820)
Median [Min, Max]	2840 [781, 7480]	2440 [888, 7410]	2770 [781, 7480]
Observed event			
Mean (SD)	0.461 (0.499)	0.506 (0.501)	0.471 (0.499)
Median [Min, Max]	0 [0, 1.00]	1.00 [0, 1.00]	0 [0, 1.00]
Age at Exam 4			
Mean (SD)	77.0 (3.79)	77.2 (3.78)	77.1 (3.79)
Median [Min, Max]	76.1 [71.4, 91.0]	76.4 [71.3, 90.9]	76.2 [71.3, 91.0]
Education, years			
Mean (SD)	11.2 (3.10)	10.4 (2.93)	11.0 (3.08)
Median [Min, Max]	12.0 [2.00, 24.0]	10.0 [3.00, 24.0]	12.0 [2.00, 24.0]
APOE E4 positive			
0	1187 (82.0%)	334 (78.2%)	1521 (81.2%)
1	260 (18.0%)	93 (21.8%)	353 (18.8%)
Grip strength			
Mean (SD)	32.0 (5.73)	31.2 (6.77)	31.8 (5.99)
Median [Min, Max]	32.0 [12.0, 55.0]	31.0 [0, 49.0]	32.0 [0, 55.0]

3.3 Statistical Analysis

We applied the Martingale R-learner described in Chapter 2 to estimate the time-varying heterogeneous treatment effect of midlife heavy drinking on cognitive-impairment-free survival.

Let $D = 1$ denote midlife heavy drinking and $D = 0$ denote non-heavy drinking. Let $\mathbf{Z}$ denote the baseline covariates included in the analysis: age at Exam 4, education, APOE E4 status, and grip strength.

The target estimand was the time-varying hazard difference

$$\tau(t, \mathbf{Z}) = \lambda(t; \mathbf{Z}, 1) - \lambda(t; \mathbf{Z}, 0),$$

where $\lambda(t; \mathbf{Z}, d)$ is the hazard of cognitive impairment or mortality under drinking status d . We also considered the cumulative treatment effect

$$\mathcal{T}(t, \mathbf{Z}) = \int_0^t \tau(u, \mathbf{Z}) du,$$

which can be interpreted as the cumulative hazard difference comparing midlife heavy drinkers with non-heavy drinkers, conditional on baseline covariates.

The Martingale R-learner requires estimation of two nuisance components: the marginal cumulative hazard function and the risk-set propensity score. The marginal cumulative hazard was estimated from the observed survival outcome and baseline covariates. The risk-set propensity score was estimated using the baseline propensity score together with conditional survival models, as described in Chapter 2. Cross-fitting was used to reduce overfitting bias from nuisance estimation.

The treatment effect function was estimated using basis expansions of time and baseline covariates. To visualize the estimated HTE, we focused on two continuous covariates, age and grip strength, and two categorical covariates, education and APOE E4 status. We plotted the estimated HTE surfaces over follow-up time and age while fixing grip strength at its mean, and over follow-up time and grip strength while fixing age at its mean. Each surface was displayed separately for the four strata defined by education and APOE E4 status.

Causal interpretation of this analysis relies on the assumptions described in Chapter 2,

including consistency, no unmeasured confounding conditional on the measured covariates, risk-set positivity, and non-informative censoring. Because HAAS is an observational study, these assumptions cannot be verified from the observed data alone. Therefore, the results should be interpreted as conditional causal effect estimates under the measured-covariate adjustment set used in this analysis.

3.4 Estimated Heterogeneous Treatment Effects

Figure 3.1 shows the estimated HTE surfaces varying by follow-up time and baseline age or grip strength for the four education and APOE E4 strata. The top row shows the instantaneous treatment effect $\tau(t, \mathbf{Z})$ on the hazard-difference scale. The bottom row shows the cumulative treatment effect $\mathcal{T}(t, \mathbf{Z})$ on the cumulative hazard-difference scale. Positive values indicate that midlife heavy drinking is associated with a higher hazard of cognitive impairment or mortality, corresponding to shorter cognitive-impairment-free survival.

The estimated HTE surfaces displayed nonlinear patterns over follow-up time and baseline covariates. Across most covariate values, the estimated effects were positive, suggesting that midlife heavy drinking was associated with increased hazard of cognitive impairment or mortality. The magnitude of the estimated harmful effect varied across baseline characteristics. In particular, the estimated harm of midlife heavy drinking appeared smaller among participants with more than 12 years of education. This pattern is consistent with prior evidence that higher education is associated with better cognitive outcomes, although the present analysis estimates the effect of midlife heavy drinking rather than the direct effect of education.

The age-varying surfaces suggest that the estimated effect of midlife heavy drinking is not constant across baseline age. Similarly, the grip-strength surfaces suggest that physical function may modify the estimated effect. These patterns support the use of a time-varying HTE method rather than a model that assumes a constant treatment effect over time or across individuals.




Figure 3.1. Estimated heterogeneous treatment effect surfaces for midlife heavy drinking in the HAAS application. The top row shows the instantaneous treatment effect $\tau(t, \mathbf{Z})$ on the hazard-difference scale, and the bottom row shows the cumulative treatment effect $\mathcal{T}(t, \mathbf{Z})$ on the cumulative hazard-difference scale. The left column varies baseline age and follow-up time, with grip strength fixed at its mean. The right column varies grip strength and follow-up time, with baseline age fixed at its mean. Surfaces are shown separately by education and APOE E4 status. Positive values indicate higher hazard of cognitive impairment or mortality among midlife heavy drinkers compared with non-heavy drinkers.

3.5 HTE-based Stratification

To further summarize the estimated treatment effect heterogeneity, we stratified participants according to their estimated 10-year cumulative HTE. For each participant, we first evaluated the estimated instantaneous treatment effect $\hat{\tau}(t, \mathbf{Z}_i)$ over a grid of 300 equally spaced time points from baseline to 10 years. We then approximated the 10-year cumulative HTE using the trapezoidal rule,

$$\widehat{\mathcal{T}}(10, \mathbf{Z}_i) = \int_0^{10 \text{ years}} \hat{\tau}(t, \mathbf{Z}_i) dt.$$

This participant-specific cumulative HTE summarizes the estimated effect of midlife heavy drinking over the first 10 years of follow-up.

Participants were then divided into three groups based on tertiles of the estimated 10-year cumulative HTE: low, middle, and high. These groups represent participants with relatively smaller, intermediate, and larger estimated cumulative treatment effects. We then estimated cognitive-impairment-free survival within each HTE tertile using the Kaplan–Meier estimator, separately for midlife heavy drinkers and non-heavy drinkers.

Figure 3.2 shows the resulting Kaplan–Meier curves. These curves provide a descriptive summary of observed survival patterns within groups defined by the estimated cumulative HTE. Because the Kaplan–Meier curves themselves are not adjusted for measured confounding, they should not be interpreted as causal treatment effect estimates. Instead, they are used to visualize whether participants with different estimated HTE levels also show different observed survival patterns and crude contrasts between heavy drinkers and non-heavy drinkers.

3.6 Discussion

This chapter applied the Martingale R-learner to the Honolulu–Asia Aging Study to investigate heterogeneity in the effect of midlife heavy drinking on cognitive-impairment-free survival. The analysis focused on a composite endpoint defined by CASI score below 74 or




Figure 3.2. Kaplan–Meier curves for cognitive-impairment-free survival stratified by tertiles of the estimated 10-year cumulative HTE. Participants were divided into low, middle, and high groups based on $\widehat{\mathcal{T}}(10, \mathbf{Z}_i)$. The Kaplan–Meier curves are descriptive and are not adjusted for measured confounding.

mortality, whichever occurred first. This endpoint reflects loss of cognitive-impairment-free survival in an elderly cohort.

The results suggest that the association between midlife heavy drinking and cognitive-impairment-free survival varies across baseline characteristics and over follow-up time. The estimated HTE surfaces were generally positive, indicating that midlife heavy drinking was associated with a higher hazard of cognitive impairment or mortality for most covariate profiles. The magnitude of this harmful association appeared smaller among participants with more than 12 years of education, and the estimated effect also varied by age, APOE E4 status, and grip strength.

The HTE-based stratification further summarized these patterns by dividing participants into tertiles of the estimated 10-year cumulative HTE. The resulting Kaplan–Meier curves showed observed cognitive-impairment-free survival patterns within groups defined by the estimated treatment effect. These curves were used descriptively and should not be interpreted

as confounding-adjusted causal estimates.

This application illustrates the scientific value of estimating time-varying heterogeneous treatment effects. A single average effect would summarize the overall association between midlife heavy drinking and cognitive outcomes, but it would not show how the estimated effect changes across time or across clinically meaningful baseline characteristics. The Martingale R-learner provides a flexible way to visualize and summarize these patterns.

Several limitations should be noted. First, HAAS is an observational study, so causal interpretation depends on adjustment for measured confounders and the absence of unmeasured confounding. Second, the analytic cohort is restricted to participants who had required data available at Exam 4; therefore, the estimand applies to this late-life analytic cohort. Third, some subgroups, particularly participants with education greater than 12 years and APOE E4 positive status, had relatively small sample sizes, which may lead to greater uncertainty in the estimated HTE surfaces.

Despite these limitations, the HAAS application demonstrates how the Martingale R-learner can be used to study treatment or exposure effect heterogeneity in long-term cohort studies with time-to-event outcomes. The method complements traditional survival analysis by estimating how the effect of an exposure varies jointly with time and participant characteristics.

Acknowledgements

Chapters 2 and 3, in part, contain material that is currently submitted for publication. Jue Hou; Yuchen Qi; Ronghui Xu. The dissertation author was a coauthor of this material.

Chapter 4

Prior Exposure to Advanced Therapy and Risk of Serious Infections in Inflammatory Bowel Disease

4.1 Introduction

Advanced therapies, including biologic and small-molecule agents, are central to the management of moderate-to-severe inflammatory bowel disease (IBD). Despite their effectiveness, a substantial proportion of patients discontinue or switch advanced therapies because of primary nonresponse, loss of response, intolerance, adverse events, or other clinical considerations. As treatment options expand, patients increasingly move from one advanced therapy to another, raising important questions about the safety of treatment sequencing.

One clinically relevant question is whether patients should have a washout period between discontinuation of a prior advanced therapy and initiation of a new advanced therapy. Shorter intervals may raise concerns about overlapping immunosuppressive effects and increased infection risk. Conversely, longer intervals may delay effective treatment, worsen disease activity, and increase corticosteroid exposure, which itself is an important risk factor for infection. Regulatory guidance has recommended washout periods in IBD trials based on pharmacokinetic considerations, such as approximately five drug half-lives or undetectable serum drug concentrations (U.S. Food and Drug Administration, 2022a,b). However, the empirical evidence supporting specific

washout durations remains limited, and washout requirements vary substantially across trials (Wils et al., 2024).

Prior observational studies have evaluated serious infection risk associated with biologic and small-molecule therapies in IBD (Kirchgesner et al., 2018; Singh et al., 2021; Khan et al., 2026; Axelrad et al., 2026; Ahuja et al., 2026). More recent work has also examined treatment switching, advanced combination therapy, and washout intervals in immune-mediated inflammatory diseases (Kruger et al., 2025; Vuyyuru et al., 2023; Feagan et al., 2023; Solitano et al., 2024, 2026). Nevertheless, it remains uncertain whether prior exposure to an advanced therapy, and the timing of discontinuation of that prior therapy, independently influence serious infection risk after initiation of a new advanced therapy.

In this chapter, we evaluate serious infection risk among patients with IBD initiating a new advanced therapy using OptumLabs[®] Data. The analysis addresses two related questions. First, among patients initiating an advanced therapy, is prior exposure to a different advanced therapy within the preceding 12 months associated with serious infection risk after initiation of the new therapy? Second, among patients with prior advanced therapy exposure, is the timing of discontinuation of the prior therapy associated with subsequent serious infection risk?

This chapter emphasizes inverse probability of treatment weighting (IPTW) as a causal inference tool for the primary exposure comparison. Although the corresponding manuscript focused primarily on multivariable regression, IPTW is important for this dissertation because it connects the real-world evidence application to the broader theme of causal treatment effect estimation with time-to-event outcomes. The IPTW analysis creates a weighted pseudo-population in which measured baseline covariates are better balanced between patients with and without prior advanced therapy exposure, allowing estimation of the association between prior exposure and serious infection risk under a causal weighting framework.

4.2 Data Source and Study Population

We conducted a retrospective cohort study using medical and pharmacy administrative claims from OptumLabs® Data Warehouse. The database includes longitudinal de-identified health information on commercially insured and Medicare Advantage enrollees across the United States. The claims data contain information on demographics, enrollment, inpatient and outpatient encounters, procedures, and prescription medications.

The study population included patients with IBD who initiated an advanced therapy between January 1, 2018 and December 31, 2023. Advanced therapies included tumor necrosis factor (TNF) antagonists, vedolizumab, ustekinumab, risankizumab, and Janus kinase (JAK) inhibitors. The index date was defined as the date of initiation of the new advanced therapy, based on prescription fill date or infusion date.

Eligible patients were adults with at least one diagnosis code for Crohn's disease or ulcerative colitis before the index date, continuous health plan enrollment with pharmacy benefits, and no prescription or infusion for the respective index advanced therapy in the 12 months before the index date. This new-user design ensured that each treatment episode represented initiation of the index advanced therapy after a baseline period available for covariate assessment.

Patients were followed from the index date until the earliest occurrence of serious infection, health plan disenrollment, discontinuation of the index advanced therapy, three years after index therapy initiation, July 1, 2024, or end of available data. The final cohort included 21,564 advanced therapy initiators. The mean age was 47 years, 52% were female, and 57% had Crohn's disease. A total of 7,490 patients, or 34.7%, had prior advanced therapy exposure within the 12 months before initiation of the new therapy.

4.3 Exposure, Outcome, and Covariates

The primary exposure was prior advanced therapy exposure within the 12 months before the index date. Patients were classified as exposed if they filled a prescription or received an

infusion for an alternative advanced therapy, of a different class than the index advanced therapy, in the 12 months before initiation of the new therapy. The comparator group consisted of patients without exposure to an alternative advanced therapy in the preceding 12 months. Therefore, the comparison should be interpreted as prior exposure within the preceding 12 months versus no prior exposure within that window.

Among patients with prior advanced therapy exposure, timing of discontinuation of the prior therapy was classified according to the last prescription fill date or infusion date before the index date. Timing was categorized as discontinuation within ≤ 1 month, $1 < 3$ months, or 3–12 months before initiation of the new therapy. In secondary analyses, timing was also modeled continuously in weeks.

The primary outcome was time to first serious infection after initiation of the index advanced therapy. Serious infection was defined as hospitalization with a primary diagnosis code for serious or opportunistic infection. In a short-term analysis, we also evaluated serious infection within 6 months after initiation of the new therapy.

Baseline covariates were measured before or at the index date. Demographic covariates included age, sex, race or ethnicity, and calendar year. Disease-related covariates included IBD subtype, categorized as Crohn's disease or ulcerative colitis. Comorbidities included obesity, hyperlipidemia, smoking, diabetes mellitus, coronary artery disease, prior hip or knee surgery, and Charlson comorbidity index. Healthcare utilization variables included all-cause hospitalization or emergency department visits, lower endoscopy, abdominal imaging, IBD-related surgery, and prior serious infection. Medication variables included prior and concomitant immunomodulator use, oral corticosteroid use, and opioid use.

4.4 Statistical Analysis

Baseline characteristics were summarized using descriptive statistics. We first compared patients with prior advanced therapy exposure within the preceding 12 months with patients

without prior exposure in that window. Among patients with prior exposure, we also summarized baseline characteristics by timing of discontinuation of the prior advanced therapy.

The primary thesis analysis used inverse probability of treatment weighting to estimate the effect of prior advanced therapy exposure on serious infection risk. Let $D_i = 1$ denote prior advanced therapy exposure within the 12 months before the index date, and let $D_i = 0$ denote no prior exposure within that window. Let $\mathbf{Z}_i$ denote measured baseline covariates. We estimated the propensity score

$$e(\mathbf{Z}_i) = \mathbb{P}(D_i = 1 \mid \mathbf{Z}_i),$$

using a logistic regression model including baseline covariates considered potential confounders of the relationship between prior advanced therapy exposure and serious infection risk.

For the average treatment effect in the study population, inverse probability weights can be defined as

$$w_i = \frac{D_i}{\widehat{e}(\mathbf{Z}_i)} + \frac{1 - D_i}{1 - \widehat{e}(\mathbf{Z}_i)}.$$

To improve stability, we used stabilized weights,

$$sw_i = D_i \frac{\widehat{p}}{\widehat{e}(\mathbf{Z}_i)} + (1 - D_i) \frac{1 - \widehat{p}}{1 - \widehat{e}(\mathbf{Z}_i)},$$

where $\widehat{p} = n^{-1} \sum_{i=1}^n D_i$ is the marginal probability of prior advanced therapy exposure. Covariate balance before and after weighting was evaluated using standardized mean differences. The distribution of propensity scores and stabilized weights was also examined to assess positivity and the influence of extreme weights.

After constructing the weighted pseudo-population, we estimated the association between prior advanced therapy exposure and time to serious infection using a weighted Cox proportional hazards model. Robust standard errors were used for inference. This IPTW-weighted Cox model was the primary causal analysis for the prior advanced therapy exposure question in this dissertation.

As a regression-adjusted comparison, we also fit multivariable Cox proportional hazards models in the unweighted cohort. These models adjusted for prespecified covariates, including index advanced therapy class, age, IBD subtype, Charlson comorbidity index, prior serious infection, concomitant immunomodulator use, concomitant oral corticosteroid use, and concomitant opioid use, as well as additional confounders selected through regression screening.

Among patients with prior advanced therapy exposure, we evaluated whether timing of discontinuation of the prior therapy was associated with serious infection risk. Cox proportional hazards models compared patients who discontinued the prior therapy within ≤ 1 month or $1 < 3$ months before index date with those who discontinued 3–12 months before index date. Timing was also modeled as a continuous variable measured in weeks. We further evaluated effect modification by prior therapy class, categorized as prior TNF antagonist versus prior non-TNF advanced therapy, using likelihood ratio tests for interaction.

Because overlapping immunosuppressive exposure would be expected to affect short-term infection risk most strongly, we also evaluated serious infection within 6 months after initiation of the new advanced therapy using multivariable logistic regression. Finally, because shorter intervals between therapies may reflect greater underlying disease activity, we compared corticosteroid use within the first 6 months after index date across timing groups.

4.5 Cohort Characteristics

The cohort included 21,564 patients with IBD initiating a new advanced therapy. Of these, 7,490 patients had prior advanced therapy exposure within the preceding 12 months, and 14,074 had no prior exposure within that window. Among patients with prior exposure, 2,353 discontinued the prior therapy within ≤ 1 month before the index date, 3,379 discontinued $1 < 3$ months before index date, and 1,758 discontinued 3–12 months before index date.

Patients with prior advanced therapy exposure differed from those without prior exposure in several baseline characteristics. Patients with prior exposure were more likely to initiate

vedolizumab, IL-12/23 or IL-23 antagonists, or JAK inhibitors as the index therapy, whereas patients without prior exposure were more likely to initiate TNF antagonists. Prior advanced therapy exposure was also associated with greater prior medication use and healthcare utilization.

Among patients with prior advanced therapy exposure, shorter intervals between discontinuation of the prior therapy and initiation of the new therapy were associated with different baseline profiles. Patients with a ≤ 1 month interval were younger on average and had more frequent concomitant corticosteroid use than those with longer intervals. Patients with a 3–12 month interval had higher comorbidity burden and more prior IBD-related surgery. These differences support the need for careful confounding adjustment when evaluating the relationship between treatment timing and infection risk.

Table 4.1. Baseline characteristics of patients with IBD initiating a new advanced therapy, by prior advanced therapy exposure within the preceding 12 months.

Characteristic	Prior AT exposure (N = 7,490)	No prior AT exposure (N = 14,074)
Age, mean (SD)	46 (17)	47 (18)
Female sex	3,954 (53%)	7,270 (52%)
Crohn's disease	4,436 (59%)	7,809 (55%)
Ulcerative colitis	3,054 (41%)	6,265 (45%)
Charlson comorbidity index ≥ 2	1,727 (23%)	3,321 (24%)
Prior serious infection	901 (12%)	1,526 (11%)
Concomitant corticosteroids	2,996 (40%)	5,460 (39%)
Steroid use in prior year	4,763 (64%)	7,781 (55%)
Emergency department visit or hospitalization	4,520 (60%)	7,558 (54%)
Prior IBD-related surgery	350 (4.7%)	428 (3.0%)

4.6 Prior Advanced Therapy Exposure and Serious Infection Risk

Over a mean follow-up of 1.6 years, 1,955 patients, or 9.1%, experienced a serious infection. Serious infections occurred in 725 of 7,490 patients with prior advanced therapy exposure and in 1,230 of 14,074 patients without prior advanced therapy exposure in the

preceding 12 months.

In the multivariable Cox proportional hazards model, prior advanced therapy exposure was associated with a higher risk of serious infection compared with no prior exposure within the preceding 12 months:

$$\text{HR} = 1.19, \quad 95\% \text{ CI} : 1.07\text{--}1.33.$$

Other factors associated with higher serious infection risk included older age, higher comorbidity burden, diabetes mellitus, prior serious infection, concomitant corticosteroid use, concomitant opioid use, prior corticosteroid use, Hispanic race or ethnicity, and emergency department visit or hospitalization in the prior year.

In the IPTW analysis, prior advanced therapy exposure was evaluated after balancing measured baseline covariates between exposure groups. The estimated propensity scores ranged from 0.022 to 0.881, with a median of 0.358. Stabilized IPTW weights had mean 1.009 and ranged from 0.396 to 11.319. Figure 4.1 shows the standardized mean differences before and after weighting. Covariate balance improved substantially after weighting, with most adjusted standardized mean differences close to zero and within the commonly used threshold of 0.1 in absolute value.

The IPTW-weighted Cox model estimated a hazard ratio of

$$\text{HR}_{\text{IPTW}} = 1.22, \quad 95\% \text{ CI} : 1.07\text{--}1.38, \quad p = 0.003.$$

This estimate was similar to the multivariable Cox model estimate, supporting the finding that prior advanced therapy exposure within the preceding 12 months was associated with a modestly higher risk of serious infection after initiation of a new advanced therapy.




Figure 4.1. Standardized mean differences of baseline covariates before and after inverse probability of treatment weighting for prior advanced therapy exposure. Dashed vertical lines indicate absolute standardized mean difference of 0.1.

Table 4.2. Association between prior advanced therapy exposure and serious infection risk after initiation of a new advanced therapy.

Analysis	Hazard ratio	95% CI	Interpretation
Multivariable Cox model	1.19	1.07–1.33	Regression-adjusted association
IPTW-weighted Cox model	1.22	1.07–1.38	Weighted causal contrast

4.7 Timing of Prior Advanced Therapy Discontinuation

Among patients with prior advanced therapy exposure, timing of discontinuation of the prior therapy was not associated with overall serious infection risk in the primary timing analysis. Using discontinuation 3–12 months before index date as the reference group, the adjusted hazard ratio was

$$\text{HR} = 1.17, \quad 95\% \text{ CI} : 0.95\text{--}1.43$$

for discontinuation within ≤ 1 month, and

$$\text{HR} = 1.01, \quad 95\% \text{ CI} : 0.83\text{--}1.22$$

for discontinuation 1–< 3 months before index date. When timing was modeled continuously in weeks, there was no evidence of a dose-response relationship:

$$\text{HR per week} = 1.00, \quad 95\% \text{ CI} : 0.99\text{--}1.01.$$

Effect modification by prior therapy class was observed in subgroup analyses. The association between timing of discontinuation and infection risk differed according to whether the prior therapy was a TNF antagonist or a non-TNF advanced therapy. Among patients whose prior therapy was a TNF antagonist, discontinuation within ≤ 1 month was associated with higher serious infection risk compared with discontinuation 3–12 months before the index date:

$$\text{HR} = 1.37, \quad 95\% \text{ CI} : 1.06\text{--}1.77.$$

In contrast, no such association was observed among patients whose prior therapy was a non-TNF advanced therapy. This subgroup result should be interpreted cautiously because it may reflect residual confounding by disease activity, treatment complexity, and clinical reasons for rapid

switching.

Table 4.3. Association between timing of prior advanced therapy discontinuation and serious infection risk among patients with prior advanced therapy exposure.

Timing of discontinuation	Adjusted HR	95% CI
3–12 months before index date	Reference	Reference
≤ 1 month before index date	1.17	0.95–1.43
1–< 3 months before index date	1.01	0.83–1.22
Continuous timing, per week	1.00	0.99–1.01

4.8 Short-term Infection Risk and Corticosteroid Use

Among patients with prior advanced therapy exposure, 387 patients experienced a serious infection within 6 months after initiation of the new advanced therapy. In multivariable logistic regression models, timing of discontinuation of the prior therapy was not associated with short-term serious infection risk. Compared with discontinuation 3–12 months before the index date, the adjusted odds ratio was

$$\text{OR} = 1.27, \quad 95\% \text{ CI} : 0.94\text{--}1.71$$

for discontinuation within ≤ 1 month, and

$$\text{OR} = 1.18, \quad 95\% \text{ CI} : 0.89\text{--}1.56$$

for discontinuation 1–< 3 months before index date. Modeling timing continuously yielded similar results:

$$\text{OR per week} = 0.995, \quad 95\% \text{ CI} : 0.984\text{--}1.007.$$

Patients with shorter intervals between discontinuation of the prior therapy and initiation of the new therapy were more likely to receive corticosteroids within the first 6 months after the index date. The proportion with at least one corticosteroid prescription was 33% among

patients with a ≤ 1 month interval, 32% among patients with a $1-< 3$ month interval, and 27% among patients with a 3–12 month interval. The corresponding proportions with at least two corticosteroid prescriptions were 20%, 19%, and 15%, respectively. These findings suggest that shorter intervals between therapies may occur more often among patients with greater disease activity or more refractory disease.

Table 4.4. Short-term serious infection risk and corticosteroid use among patients with prior advanced therapy exposure.

Analysis	Effect estimate	95% CI
Serious infection within 6 months: ≤ 1 month interval	OR 1.27	0.94–1.71
Serious infection within 6 months: $1-< 3$ month interval	OR 1.18	0.89–1.56
Continuous timing, per week	OR 0.995	0.984–1.007

4.9 Discussion

In this large real-world cohort of patients with IBD initiating a new advanced therapy, prior advanced therapy exposure within the preceding 12 months was associated with a modestly higher risk of serious infection. This association was observed in the multivariable Cox model and was also supported by the IPTW-weighted Cox model. The consistency between regression adjustment and IPTW strengthens the interpretation that prior advanced therapy exposure is associated with increased infection risk after initiation of a new advanced therapy, under measured-confounding assumptions.

Among patients with prior advanced therapy exposure, the timing of discontinuation of the prior therapy was not associated with overall serious infection risk. Patients who discontinued the prior therapy within ≤ 1 month or $1-< 3$ months before initiation of the new therapy did not have statistically significantly higher overall infection risk compared with those who discontinued 3–12 months before initiation. Modeling timing continuously also did not show a dose-response relationship. These findings suggest that a shorter interval between therapies may not independently increase serious infection risk in the overall population of patients switching

advanced therapies.

The subgroup analysis by prior therapy class suggested that short intervals may be associated with higher infection risk among patients switching from prior TNF antagonists, but not among those switching from non-TNF advanced therapies. This result should be interpreted as hypothesis-generating. Patients requiring rapid treatment transitions may differ in disease severity, corticosteroid exposure, comorbidity, and other unmeasured clinical characteristics. Therefore, the apparent class-specific timing effect may partly reflect residual confounding.

Corticosteroid exposure appeared to be an important marker of risk during treatment transitions. Patients with shorter intervals between therapies were more likely to receive corticosteroids after initiation of the index therapy, suggesting higher underlying disease activity or more refractory disease. Because corticosteroids are strongly associated with infection risk, these findings support the interpretation that infection risk during treatment transitions may be driven more by disease activity and corticosteroid exposure than by the length of the washout interval alone.

This chapter has several strengths. The analysis used a large contemporary claims-based cohort with detailed longitudinal information on therapy initiation, prior medication exposure, healthcare utilization, and serious infection outcomes. The study design focused on new initiation of advanced therapy and used both long-term time-to-event outcomes and short-term infection outcomes. The dissertation analysis additionally incorporated IPTW, allowing the primary exposure comparison to be interpreted through a causal inference framework under measured confounding assumptions.

Several limitations should also be noted. First, as with all observational claims-based studies, causal interpretation depends on the absence of unmeasured confounding. Claims data do not include detailed clinical disease activity measures, endoscopic findings, biomarkers, or physician rationale for switching therapies. Second, corticosteroid use may function both as a confounder and as a marker of disease severity, making full adjustment for disease activity difficult. Third, treatment discontinuation was defined using prescription fill or infusion dates,

which may not perfectly capture actual pharmacologic exposure, adherence, or drug-specific clearance. Fourth, the outcome definition captured serious infections requiring hospitalization and may not generalize to milder infections treated in outpatient settings.

Overall, the results suggest that prior advanced therapy exposure is associated with modestly higher serious infection risk after initiation of a new advanced therapy, whereas timing of discontinuation of the prior therapy is not strongly associated with infection risk in the overall cohort. These findings support a risk-adapted view of therapy transitions in IBD, emphasizing disease control, minimization of corticosteroid exposure, and individualized safety assessment rather than fixed prolonged washout intervals.

Acknowledgements

Chapter 4, in part, is a reprint of the material as it appears in *Journal of Crohn's and Colitis*. Virginia Solitano; Dhruv Ahuja; Yuchen Qi; Sagar B. Patel; Marianne Hupé; Sudheer K. Vuyyuru; Anna L. Silverman; Ronghui Xu; Shane W. Goodwin; Ola Olén; Ashwin N. Ananthakrishnan; Vipul Jairath; Siddharth Singh, 2026. The dissertation author was a coauthor of this paper.

Appendix A

Proofs for Chapter 2

This appendix provides derivations and proofs for the main statements in Chapter 2. Throughout, expectations are taken under the true data-generating distribution. We assume the required regularity conditions for interchanging differentiation, integration, and expectation, and all stochastic integrals are assumed to be well-defined.

A.1 Derivation of the Risk-set Propensity Score and Marginal Hazard

Recall that

$$\pi(\mathbf{Z}) = \mathbb{P}(D = 1 \mid \mathbf{Z}), \quad \rho(t, \mathbf{Z}) = \mathbb{P}(D = 1 \mid T \geq t, \mathbf{Z}).$$

By Bayes' rule,

$$\begin{aligned} \rho(t, \mathbf{Z}) &= \frac{\mathbb{P}(T \geq t \mid D = 1, \mathbf{Z})\mathbb{P}(D = 1 \mid \mathbf{Z})}{\sum_{d=0}^1 \mathbb{P}(T \geq t \mid D = d, \mathbf{Z})\mathbb{P}(D = d \mid \mathbf{Z})} \\ &= \frac{S(t; \mathbf{Z}, 1)\pi(\mathbf{Z})}{S(t; \mathbf{Z}, 1)\pi(\mathbf{Z}) + S(t; \mathbf{Z}, 0)\{1 - \pi(\mathbf{Z})\}}. \end{aligned}$$

Under the additive hazard model,

$$\lambda(t; \mathbf{Z}, 1) = \lambda(t; \mathbf{Z}, 0) + \tau(t, \mathbf{Z}),$$

and therefore

$$\frac{S(t; \mathbf{Z}, 1)}{S(t; \mathbf{Z}, 0)} = \exp \left\{ - \int_0^t \tau(u, \mathbf{Z}) du \right\}.$$

Substituting this ratio into the expression for $\rho(t, \mathbf{Z})$ gives

$$\rho(t, \mathbf{Z}) = \frac{\pi(\mathbf{Z}) \exp \left\{ - \int_0^t \tau(u, \mathbf{Z}) du \right\}}{1 + \left[\exp \left\{ - \int_0^t \tau(u, \mathbf{Z}) du \right\} - 1 \right] \pi(\mathbf{Z})}, \quad (\text{A.1})$$

which is equation (2.6) in Chapter 2.

Next, the observed or marginal hazard conditional on $\mathbf{Z}$ can be written as

$$\begin{aligned} \lambda_m(t; \mathbf{Z}) &= \lim_{\Delta t \rightarrow 0} \frac{\mathbb{P}(T \leq t + \Delta t \mid T \geq t, \mathbf{Z})}{\Delta t} \\ &= \sum_{d=0}^1 \lambda(t; \mathbf{Z}, d) \mathbb{P}(D = d \mid T \geq t, \mathbf{Z}) \\ &= \lambda(t; \mathbf{Z}, 0) \{1 - \rho(t, \mathbf{Z})\} + \{\lambda(t; \mathbf{Z}, 0) + \tau(t, \mathbf{Z})\} \rho(t, \mathbf{Z}) \\ &= \lambda(t; \mathbf{Z}, 0) + \tau(t, \mathbf{Z}) \rho(t, \mathbf{Z}). \end{aligned}$$

Thus,

$$\lambda_m(t; \mathbf{Z}) = \lambda(t; \mathbf{Z}, 0) + \tau(t, \mathbf{Z}) \rho(t, \mathbf{Z}),$$

which is equation (2.5).

A.2 Proof of Lemma 1

Under the additive hazard model,

$$\lambda(t; \mathbf{Z}, D) = \lambda(t; \mathbf{Z}, 0) + D\tau(t, \mathbf{Z}).$$

By the counting-process representation,

$$dN(t) = Y(t) \lambda(t; \mathbf{Z}, D) dt + dM(t),$$

where $M(t)$ is a martingale with respect to the natural filtration.

Using the marginal hazard derived above,

$$d\Lambda_m(t; \mathbf{Z}) = \lambda_m(t; \mathbf{Z}) dt = \{\lambda(t; \mathbf{Z}, 0) + \tau(t, \mathbf{Z})\rho(t, \mathbf{Z})\} dt.$$

Therefore,

$$\begin{aligned} dN(t) - Y(t)d\Lambda_m(t; \mathbf{Z}) &= Y(t)\lambda(t; \mathbf{Z}, D) dt + dM(t) - Y(t)\lambda_m(t; \mathbf{Z}) dt \\ &= Y(t) [\lambda(t; \mathbf{Z}, 0) + D\tau(t, \mathbf{Z}) - \lambda(t; \mathbf{Z}, 0) - \rho(t, \mathbf{Z})\tau(t, \mathbf{Z})] dt + dM(t) \\ &= Y(t)\{D - \rho(t, \mathbf{Z})\}\tau(t, \mathbf{Z}) dt + dM(t). \end{aligned}$$

This proves the martingale residual decomposition

$$dN(t) - Y(t)d\Lambda_m(t; \mathbf{Z}) = Y(t)\{D - \rho(t, \mathbf{Z})\}\tau(t, \mathbf{Z}) dt + dM(t),$$

which is equation (2.7).

A.3 Proof of Lemma 2

Recall the functional score

$$\begin{aligned} \psi(\tau; \Lambda_m, \rho)[\phi] &= \int_0^{t_{\max}} \phi(t, \mathbf{Z})\{D - \rho(t, \mathbf{Z})\} [dN(t) - Y(t)d\Lambda_m(t; \mathbf{Z})] \\ &\quad - \int_0^{t_{\max}} \phi(t, \mathbf{Z})Y(t)\{D - \rho(t, \mathbf{Z})\}^2 \tau(t, \mathbf{Z}) dt. \end{aligned}$$

First, we show that the score has mean zero at the truth. Evaluating at $(\tau_0, \Lambda_{m,0}, \rho_0)$ and using Lemma 1,

$$dN(t) - Y(t)d\Lambda_{m,0}(t; \mathbf{Z}) = Y(t)\{D - \rho_0(t, \mathbf{Z})\}\tau_0(t, \mathbf{Z}) dt + dM(t).$$

Substituting this into the score gives

$$\begin{aligned}\psi(\tau_0; \Lambda_{m,0}, \rho_0)[\phi] &= \int_0^{t_{\max}} \phi(t, \mathbf{Z}) \{D - \rho_0(t, \mathbf{Z})\} [Y(t) \{D - \rho_0(t, \mathbf{Z})\} \tau_0(t, \mathbf{Z}) dt + dM(t)] \\ &\quad - \int_0^{t_{\max}} \phi(t, \mathbf{Z}) Y(t) \{D - \rho_0(t, \mathbf{Z})\}^2 \tau_0(t, \mathbf{Z}) dt \\ &= \int_0^{t_{\max}} \phi(t, \mathbf{Z}) \{D - \rho_0(t, \mathbf{Z})\} dM(t).\end{aligned}$$

The integrand is predictable, and under the regularity conditions for martingale integration,

$$\mathbb{E} \left[\int_0^{t_{\max}} \phi(t, \mathbf{Z}) \{D - \rho_0(t, \mathbf{Z})\} dM(t) \right] = 0.$$

Thus,

$$\mathbb{E}\{\psi(\tau_0; \Lambda_{m,0}, \rho_0)[\phi]\} = 0.$$

Next, we show Neyman orthogonality with respect to the nuisance functions. Consider perturbations

$$\Lambda_r(t; \mathbf{Z}) = \Lambda_{m,0}(t; \mathbf{Z}) + r\Delta_\Lambda(t; \mathbf{Z}), \quad \rho_r(t, \mathbf{Z}) = \rho_0(t, \mathbf{Z}) + r\Delta_\rho(t, \mathbf{Z}),$$

where Δ_Λ and Δ_ρ are arbitrary nuisance directions. Define

$$\Psi(r) = \mathbb{E}\{\psi(\tau_0; \Lambda_r, \rho_r)[\phi]\}.$$

Differentiating $\Psi(r)$ at $r = 0$ gives

$$\begin{aligned}\Psi'(0) &= \mathbb{E} \int_0^{t_{\max}} \phi(t, \mathbf{Z}) [-\Delta_\rho(t, \mathbf{Z}) \{dN(t) - Y(t) d\Lambda_{m,0}(t; \mathbf{Z})\} - Y(t) \{D - \rho_0(t, \mathbf{Z})\} d\Delta_\Lambda(t; \mathbf{Z})] \\ &\quad + \mathbb{E} \int_0^{t_{\max}} 2\phi(t, \mathbf{Z}) Y(t) \{D - \rho_0(t, \mathbf{Z})\} \Delta_\rho(t, \mathbf{Z}) \tau_0(t, \mathbf{Z}) dt.\end{aligned}$$

Using the residual decomposition again,

$$dN(t) - Y(t)d\Lambda_{m,0}(t; \mathbf{Z}) = Y(t)\{D - \rho_0(t, \mathbf{Z})\}\tau_0(t, \mathbf{Z})dt + dM(t),$$

we obtain

$$\begin{aligned} \Psi'(0) = & \mathbb{E} \int_0^{t_{\max}} \phi(t, \mathbf{Z}) Y(t) \{D - \rho_0(t, \mathbf{Z})\} [\Delta_\rho(t, \mathbf{Z}) \tau_0(t, \mathbf{Z}) dt - d\Delta_\Lambda(t; \mathbf{Z})] \\ & - \mathbb{E} \int_0^{t_{\max}} \phi(t, \mathbf{Z}) \Delta_\rho(t, \mathbf{Z}) dM(t). \end{aligned}$$

The second term is zero because it is the expectation of a martingale integral. For the first term, note that $\Delta_\rho(t, \mathbf{Z})$, $\tau_0(t, \mathbf{Z})$, and $d\Delta_\Lambda(t; \mathbf{Z})$ are functions of $(t, \mathbf{Z})$ and do not depend on D after conditioning on $(Y(t), \mathbf{Z})$. Moreover,

$$\mathbb{E}\{Y(t)[D - \rho_0(t, \mathbf{Z})] \mid \mathbf{Z}\} = \mathbb{P}(Y(t) = 1 \mid \mathbf{Z}) [\mathbb{E}\{D \mid Y(t) = 1, \mathbf{Z}\} - \rho_0(t, \mathbf{Z})] = 0,$$

because

$$\rho_0(t, \mathbf{Z}) = \mathbb{P}(D = 1 \mid Y(t) = 1, \mathbf{Z}).$$

Therefore, the first term is also zero, and hence

$$\Psi'(0) = 0.$$

This proves Neyman orthogonality:

$$D_g \mathbb{E}\{\psi(\tau_0; g_0)[\phi]\}[g - g_0] = 0.$$

A.4 Proof of Theorem 1

Let

$$R_{\mathcal{H}} = \text{Rate}(\mathcal{H}, S_2, \delta/2; \hat{\tau}, \hat{g}), \quad R_{\mathcal{G}} = \text{Rate}(\mathcal{G}, S_1, \delta/2).$$

By sample splitting, we may condition on the sample used to estimate the nuisance function $\hat{g}$ when analyzing the target estimator on the second sample. On the event with probability at least $1 - \delta$, we have

$$\|\hat{g} - g_0\|_{\mathcal{G}} \leq R_{\mathcal{G}},$$

and the target estimator satisfies the rate bound associated with $R_{\mathcal{H}}$.

Let

$$a = \|\hat{\tau} - \tau_0\|_{\mathcal{H}}.$$

By the definition of the target estimation rate and the pseudo-regret in equation (2.13),

$$\mathcal{R}(\hat{\tau}, \hat{g}) - \mathcal{R}(\tau_0, \hat{g}) \lesssim aR_{\mathcal{H}} - \mathbb{E}\{\psi_i(\tau_0, \hat{g})[\hat{\tau} - \tau_0]\}.$$

By the local convexity condition in Assumption I,

$$\mathcal{R}(\hat{\tau}, \hat{g}) - \mathcal{R}(\tau_0, \hat{g}) \gtrsim a^2 - R_{\mathcal{G}}^{4/(1+r)}.$$

Combining these two inequalities yields

$$a^2 \lesssim aR_{\mathcal{H}} + R_{\mathcal{G}}^{4/(1+r)} - \mathbb{E}\{\psi_i(\tau_0, \hat{g})[\hat{\tau} - \tau_0]\}. \quad (\text{A.2})$$

It remains to control the nuisance-induced term. By a Taylor expansion of the score in

the nuisance direction around g_0 ,

$$\begin{aligned}\mathbb{E}\{\psi_i(\tau_0, \widehat{g})[\widehat{\tau} - \tau_0]\} &= \mathbb{E}\{\psi_i(\tau_0, g_0)[\widehat{\tau} - \tau_0]\} \\ &\quad + D_g \mathbb{E}\{\psi_i(\tau_0, g_0)[\widehat{\tau} - \tau_0]\}[\widehat{g} - g_0] \\ &\quad + \frac{1}{2} D_g^2 \mathbb{E}\{\psi_i(\tau_0, \bar{g})[\widehat{\tau} - \tau_0]\}[\widehat{g} - g_0, \widehat{g} - g_0],\end{aligned}$$

for some $\bar{g}$ on the line segment between g_0 and $\widehat{g}$. The first term is zero by identification, and the first-order derivative term is zero by Neyman orthogonality. Therefore, only the second-order term remains. By Assumption H,

$$|\mathbb{E}\{\psi_i(\tau_0, \widehat{g})[\widehat{\tau} - \tau_0]\}| \lesssim a^{1-r} \|\widehat{g} - g_0\|_{\mathcal{G}}^2 \lesssim a^{1-r} R_{\mathcal{G}}^2.$$

By Young's inequality, for any $\eta > 0$,

$$a^{1-r} R_{\mathcal{G}}^2 \leq \eta a^2 + C_\eta R_{\mathcal{G}}^{4/(1+r)},$$

where C_η is a constant depending on η . Choosing η sufficiently small and substituting into (A.2), we obtain

$$a^2 \lesssim a R_{\mathcal{H}} + R_{\mathcal{G}}^{4/(1+r)}.$$

Solving this quadratic inequality gives

$$a \lesssim R_{\mathcal{H}} + R_{\mathcal{G}}^{2/(1+r)}.$$

Substituting the definitions of $R_{\mathcal{H}}$ and $R_{\mathcal{G}}$ yields

$$\|\widehat{\tau} - \tau_0\|_{\mathcal{H}} \lesssim \text{Rate}(\mathcal{H}, S_2, \delta/2; \widehat{\tau}, \widehat{g}) + \text{Rate}(\mathcal{G}, S_1, \delta/2)^{2/(1+r)}.$$

This proves Theorem 1.

Bibliography

- Aalen, O. O. (1989). A linear regression model for the analysis of life times. *Statistics in Medicine*, 8(8):907–925.
- Abraich, A., Guilloux, A., and Hanczar, B. (2022). Survcaus: Representation balancing for survival causal inference. arXiv:2203.15672.
- Ahuja, D., Dziegielewski, C., Yeh, K.-H., Patel, S. B., Goodwin, S. W., Ma, C., Ananthakrishnan, A. N., Singh, N., Jairath, V., Xu, R., and Singh, S. (2026). Comparative safety of advanced therapies in patients with ulcerative colitis: An administrative claims-based study. *American Journal of Gastroenterology*, 121:1192–1201.
- Athey, S. and Imbens, G. (2016). Recursive partitioning for heterogeneous causal effects. *Proceedings of the National Academy of Sciences*, 113(27):7353–7360.
- Axelrad, J., Forss, A., Söderling, J., Mårild, K., Halfvarson, J., Naucler, P., Ludvigsson, J. F., Olén, O., and SWIBREG Study Group (2026). Risk of serious infections in patients with inflammatory bowel disease treated with biologic and small molecule therapies: A nationwide cohort study. *Journal of Crohn’s and Colitis*, 20(2):jjaf229.
- Chapfuwa, P., Assaad, S., Zeng, S., Pencina, M. J., Carin, L., and Henao, R. (2021). Enabling counterfactual survival analysis with balanced representations. page 133–145.
- Cui, Y., Kosorok, M. R., Sverdrup, E., Wager, S., and Zhu, R. (2023). Estimating heterogeneous treatment effects with right-censored data via causal survival forests. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 85(2):179–211.
- Curth, A., Lee, C., and van der Schaar, M. (2021). Survite: Learning heterogeneous treatment effects from time-to-event data. *Advances in Neural Information Processing Systems*, 34:26740–26753.
- Feagan, B. G., Sands, B. E., Sandborn, W. J., Germinaro, M., Vetter, M., Shao, J., Sheng, S., Johanns, J., Panés, J., and VEGA Study Group (2023). Guselkumab plus golimumab combination therapy versus guselkumab or golimumab monotherapy in patients with ulcerative colitis (vega): A randomised, double-blind, controlled, phase 2, proof-of-concept trial. *The*

- Lancet Gastroenterology & Hepatology*, 8(4):307–320.
- Foster, D. J. and Syrgkanis, V. (2023). Orthogonal statistical learning. *The Annals of Statistics*, 51(3):879–908.
- Gelber, R. P., Launer, L. J., and White, L. R. (2012). The honolulu-asia aging study: Epidemiologic and neuropathologic research on cognitive impairment. *Current Alzheimer Research*, 9(6):664–672.
- Hahn, P. R., Murray, J. S., and Carvalho, C. M. (2020). Bayesian regression tree models for causal inference: Regularization, confounding, and heterogeneous effects (with discussion). *Bayesian Analysis*, 15(3):965–1056.
- Henderson, N. C., Louis, T. A., Rosner, G. L., and Varadhan, R. (2020). Individualized treatment effects with censored data via fully nonparametric bayesian accelerated failure time models. *Biostatistics*, 21(1):50–68.
- Hill, J. L. (2011). Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics*, 20(1):217–240.
- Hu, L., Lin, J.-Y., Sigel, K., and Kale, M. (2021). Estimating heterogeneous survival treatment effects of lung cancer screening approaches: a causal machine learning analysis. *Annals of Epidemiology*, 62:36–42.
- Imai, K. and Ratkovic, M. (2013). Estimating treatment effect heterogeneity in randomized program evaluation. *The Annals of Applied Statistics*, 7(1):443 – 470.
- Jaroszewicz, S. and Rzepakowski, P. (2014). Uplift modeling with survival data. In *Acm Sigkdd Workshop on Health Informatics (hi-kdd-14)*, New York City.
- Kang, S., Lu, W., and Zhang, J. (2018). On estimation of the optimal treatment regime with the additive hazards model. *Statistica Sinica*, 28(3):1539.
- Kennedy, E. H. (2023). Towards optimal doubly robust estimation of heterogeneous causal effects. *Electronic Journal of Statistics*, 17(2):3008–3049.
- Khan, N., Sundararajan, R., Patel, D., Patel, M., Bahirwani, J., and Mahmud, N. (2026). Comparative risk of serious infections associated with treatment of inflammatory bowel disease. *Inflammatory Bowel Diseases*, 32(2):199–206.
- Kirchgesner, J., Lemaitre, M., Carrat, F., Zureik, M., Carbonnel, F., and Dray-Spira, R. (2018). Risk of serious and opportunistic infections associated with treatment of inflammatory bowel diseases. *Gastroenterology*, 155(2):337–346.e10.

- Kruger, A. J., Dormont, F., Capit, N., Schilsky, S., Roberts, A., Lin, Y., Badalamenti, S., Wiekowski, M., Bewtra, M., and Colombel, J.-F. (2025). Biologic switch timing and risk of infection in patients with ulcerative colitis/crohn’s disease: A retrospective study. *Clinical Gastroenterology and Hepatology*, 23(12):2253–2262.
- Künzel, S. R., Sekhon, J. S., Bickel, P. J., and Yu, B. (2019). Metalearners for estimating heterogeneous treatment effects using machine learning. *Proceedings of the National Academy of Sciences*, 116(10):4156–4165.
- Li, D., Holmberg, D., Smith, T., Lewis, K., McCoy, R. G., Finney Rutten, L. J., and St. Sauver, J. L. (2020). Suitability of administrative claims databases for bariatric surgery research: Is the glass half-full or half-empty? *BMC Medical Research Methodology*, 20:274.
- National Heart, Lung, and Blood Institute (2024). Honolulu heart program. BioLINCC. Accessed 2026-07-17.
- Nie, X. and Wager, S. (2021). Quasi-oracle estimation of heterogeneous treatment effects. *Biometrika*, 108(2):299–319.
- OptumLabs (2019). Optumlabs and optumlabs data warehouse (oldw) descriptions and citation. Cambridge, MA: OptumLabs. Accessed 2026-07-17.
- Peila, R., White, L., Petrovich, H., Masaki, K., Ross, G., Havlik, R., and Launer, L. (2001). Joint effect of the apoe gene and midlife systolic blood pressure on late-life cognitive impairment: the honolulu-asia aging study. *Stroke*, 32(12):2882–2889.
- Picciotto, S., Hernán, M. A., Page, J. H., Young, J. G., and Robins, J. M. (2012). Structural nested cumulative failure time models to estimate the effects of interventions. *Journal of the American Statistical Association*, 107(499):886–900.
- Powers, S., Qian, J., Jung, K., Schuler, A., Shah, N. H., Hastie, T., and Tibshirani, R. (2018). Some methods for heterogeneous treatment effect estimation in high dimensions. *Statistics in Medicine*, 37(11):1767–1787.
- Robins, J. (1986). 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.
- Robins, J. and Hernan, M. (2008). Estimation of the causal effects of time-varying exposures. *Chapman & Hall/crc Handbooks of Modern Statistical Methods*, pages 553–599.
- Robins, J. M. (1989). The analysis of randomized and non-randomized aids treatment trials using a new approach to causal inference in longitudinal studies. *Health Service Research Methodology: a Focus on Aids*, pages 113–159.

- Robins, J. M. (1994). Correcting for non-compliance in randomized trials using structural nested mean models. *Communications in Statistics-theory and Methods*, 23(8):2379–2412.
- Robins, J. M. (2000). Marginal structural models versus structural nested models as tools for causal inference. In *Statistical Models in Epidemiology, the Environment, and Clinical Trials*, pages 95–133. Springer.
- Robins, J. M., Mark, S. D., and Newey, W. K. (1992). Estimating exposure effects by modelling the expectation of exposure conditional on confounders. *Biometrics*, pages 479–495.
- Rubin, D. B. (1980). Randomization analysis of experimental data: the fisher randomization test comment. *Journal of the American Statistical Association*, 75(371):591–593.
- Schrod, S., Schäfer, A., Solbrig, S., Lohmayer, R., Gronwald, W., Oefner, P. J., Beißbarth, T., Spang, R., Zacharias, H. U., and Altenbuchinger, M. (2022). Bites: Balanced individual treatment effect for survival data. *Bioinformatics*, 38(Supplement_1):i60–i67.
- Shalit, U., Johansson, F. D., and Sontag, D. (2017). Estimating individual treatment effect: generalization bounds and algorithms. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70, pages 3076–3085. PMLR.
- Singh, S., Heien, H. C., Sangaralingham, L., Shah, N. D., Lai, J. C., Sandborn, W. J., and Moore, A. A. (2021). Frailty and risk of serious infections in biologic-treated patients with inflammatory bowel diseases. *Inflammatory Bowel Diseases*, 27(10):1626–1633.
- Solitano, V., Ogunsakin, R. E., Yuan, Y., Bernstein, C. N., Bessissow, T., Bressler, B., Hoentjen, F., van Lierop, L., Leung, Y., Ma, C., Marshall, J. K., Narula, N., Alahmari, M., McCurdy, J. D., Murthy, S., Panaccione, R., Rosenfeld, G., Milgrom, R., Silverberg, M., and Jairath, V. (2026). Effectiveness and safety of advanced combination treatment in patients with refractory inflammatory bowel disease or concomitant immune-mediated disease or extraintestinal manifestations: A multicenter canadian study. *American Journal of Gastroenterology*, 121(2):410–423.
- Solitano, V., Yuan, Y., Singh, S., Ma, C., Nardone, O. M., Fiorino, G., Acosta Felquer, M. L., Barra, L., D’Agostino, M.-A., Pope, J., Peyrin-Biroulet, L., Danese, S., and Jairath, V. (2024). Efficacy and safety of advanced combination treatment in immune-mediated inflammatory disease: A systematic review and meta-analysis of randomized controlled trials. *Journal of Autoimmunity*, 149:103331.
- Syme, S. L., Marmot, M. G., Kagan, A., Kato, H., and Rhoads, G. (1975). Epidemiologic studies of coronary heart disease and stroke in japanese men living in japan, hawaii and california: Introduction. *American Journal of Epidemiology*, 102(6):477–480.
- Tabib, S. and Larocque, D. (2020). Non-parametric individual treatment effect estimation for

- survival data with random forests. *Bioinformatics*, 36(2):629–636.
- U.S. Food and Drug Administration (2022a). Crohn’s disease: Developing drugs for treatment. Guidance for Industry. Accessed 2026-07-17.
- U.S. Food and Drug Administration (2022b). Ulcerative colitis: Developing drugs for treatment. Guidance for Industry. Accessed 2026-07-17.
- Vansteelandt, S. and Joffe, M. (2014). Structural nested models and g-estimation: the partially realized promise. *Statistical Science*, 29(4):707 – 731.
- Vuuyuru, S. K., Archer, M., Nguyen, T. M., Beaton, M., and Jairath, V. (2023). Long washout periods between biologics for inflammatory bowel disease clinical trials are unnecessary: A canadian retrospective cohort study. *American Journal of Gastroenterology*, 118(12):2290–2293.
- Wager, S. and Athey, S. (2018). Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523):1228–1242.
- White, L. R., Petrovitch, H., Ross, G. W., Masaki, K. H., Abbott, R. D., Teng, E. L., Rodriguez, B. L., Blanchette, P. L., Havlik, R. J., Wergowske, G., Chiu, D., Foley, D. J., Murdaugh, C., and Curb, J. D. (1996). Prevalence of dementia in older japanese-american men in hawaii: The honolulu-asia aging study. *JAMA*, 276(12):955–960.
- Wils, P., Jairath, V., Sands, B. E., Reinisch, W., Danese, S., and Peyrin-Biroulet, L. (2024). Washout periods in inflammatory bowel disease trials: A systematic literature review and proposed solutions. *Clinical Gastroenterology and Hepatology*, 22(4):896–898.e13.
- Xu, S., Cobzaru, R., Finkelstein, S. N., Welsch, R. E., Ng, K., and Shahn, Z. (2024). Estimating heterogeneous treatment effects on survival outcomes using counterfactual censoring unbiased transformations. arXiv:2401.11263.
- Xu, Y., Ignatiadis, N., Sverdrup, E., Fleming, S., Wager, S., and Shah, N. (2023). Treatment heterogeneity for survival outcomes. In Zubizarreta, J. R., Stuart, E. A., Small, D. S., and Rosenbaum, P. R., editors, *Handbook of Matching and Weighting Adjustments for Causal Inference*. Routledge. Chapter preprint available on arXiv:2207.07758.
- Zhang, W., Le, T. D., Liu, L., Zhou, Z.-H., and Li, J. (2017). Mining heterogeneous causal effects for personalized cancer treatment. *Bioinformatics*, 33(15):2372–2378.
- Zhu, J. and Gallego, B. (2020). Targeted estimation of heterogeneous treatment effect in observational survival analysis. *Journal of Biomedical Informatics*, 107:103474.