



ALEXANDER MANGULAD CHRISTGAU

Model-free Methods for Event History Analysis and Efficient Adjustment

PHD THESIS

THIS THESIS HAS BEEN SUBMITTED TO THE PHD SCHOOL OF
THE FACULTY OF SCIENCE, UNIVERSITY OF COPENHAGEN

DEPARTMENT OF MATHEMATICAL SCIENCES
UNIVERSITY OF COPENHAGEN

AUGUST 2024

Alexander Mangelad Christgau
amc@math.ku.dk
Department of Mathematical Sciences
University of Copenhagen
Universitetsparken 5
2100 Copenhagen
Denmark

Thesis title: Model-free Methods for Event History Analysis and Efficient Adjustment

Supervisor: Professor Niels Richard Hansen
University of Copenhagen

Assessment Committee: Associate Professor Sebastian Weichwald (chair)
University of Copenhagen

Professor Ingrid Van Keilegom
KU Leuven

Professor Stijn Vansteelandt
Ghent University

Date of Submission: August 31,
2024

Date of Defense: November 15,
2024

ISBN: 978-87-7125-233-0

Chapter 1: © Christgau, A. M.
Chapter 2: © Institute of Mathematical Statistics.
Chapter 3 and 4: © Christgau, A. M. & Hansen, N. R.

This thesis has been submitted to the PhD School of The Faculty of Science, University of Copenhagen on 31 August 2024. It was supported by the Novo Nordisk Foundation (research grant NNF20OC0062897).

In loving memory of my grandparents, Ole and Sigrid.

Preface

This thesis is the culmination of three wonderful years from September 2021 to August 2024, under the supervision of Professor Niels Richard Hansen at the Department of Mathematical Sciences, University of Copenhagen. The thesis features three distinct manuscripts on varied, yet related, topics. Each manuscript is a stand-alone scientific contribution and can be read independently. Small discrepancies may occur between the contents of a chapter and its associated manuscript. Any typographical or mathematical errors are my own responsibility.

Acknowledgments

I am, first and foremost, grateful to Niels for giving me the opportunity to pursue a PhD, teaching me how to navigate academia, being an engaged co-author, and providing guidance, input, and feedback when necessary while allowing me to find my own path when appropriate. Most importantly, our collaboration over the past three years has been immensely enjoyable.

I extend my gratitude to my colleagues and now friends at the department. Coming to the office each day has been a delight, and our discussions about statistics and other topics have been truly entertaining, if not inspiring. Thank you for all the memories. In particular, I thank Anton for the numerous conversations, both serious and unserious, about mathematics, statistics, and everything beyond.

I thank Becca and Jake for their hospitality and enthusiasm during my stay at the University of Chicago. I am especially grateful to the Sjurson family for warmly welcoming me into their family and making my time in Chicago even more enjoyable.

A heartfelt thanks goes to all my friends for their encouragement and good times we shared during my studies, with a special mention to my friends from folkeskolen, GHG, 4. Indre, my study group from KU, and my PhD cohort. To avoid turning this into a lengthy acknowledgment, I will just say: You know who you are; thank you for your friendship. A special thanks goes to Christian, Mikkel, and Rasmus for being amazing roommates and for helping me maintain my sanity throughout my PhD studies.

Finally, but most importantly of all, I want to extend my deepest thanks to my entire family, especially Far and Mor, for their unwavering encouragement and support throughout my life. You have my deepest gratitude and love. To my grandparents, who never saw the end of this adventure, but whose inquisitive minds and kind hearts shaped me profoundly: this thesis is dedicated to you.

Alexander Mangelad Christgau
August, 2024

Abstract

This thesis contains a series of independent contributions to mathematical statistics, unified by a model-free perspective. The first chapter elaborates on how a model-free perspective can be used to formulate flexible methods that leverage prediction techniques from machine learning. Mathematical insights are obtained from concrete examples, and these insights are generalized to principles that permeate the rest of the thesis.

The second chapter studies the concept of local independence, which describes whether the evolution of one stochastic process is directly influenced by another. To test local independence, we define a model-free parameter called the Local Covariance Measure (LCM). We formulate an estimator for the LCM, from which a test of local independence is proposed. We discuss how the size and power of the proposed test can be controlled uniformly and investigate the test in a simulation study.

The third chapter focuses on covariate adjustment, a method used to estimate the effect of a treatment by accounting for observed confounding. We formulate a general framework that facilitates adjustment for any subset of covariate information. We identify the optimal covariate information for adjustment and, based on this, introduce the Debiased Outcome-adapted Propensity Estimator (DOPE) for efficient estimation of treatment effects. An instance of DOPE is implemented using neural networks, and we demonstrate its performance on both simulated and real data.

The fourth and final chapter introduces a model-free measure of the conditional association between an exposure and a time-to-event, which we call the Aalen Covariance Measure (ACM). The ACM serves as an assumption-lean generalization of the exposure coefficient in the Aalen additive hazards model. We develop a model-free estimation method and show that it is doubly robust, ensuring \sqrt{n} -consistency provided that the nuisance functions can be estimated with modest rates. A simulation study demonstrates the use of our estimator in several settings.

Sammenfatning

Denne Ph.D.-afhandling indeholder en række selvstændige bidrag inden for matematisk statistik med det til fælles, at de har et model-frit perspektiv. Det første afsnit uddyber, hvordan et model-frit perspektiv kan bruges til at formulere fleksible metoder, som benytter prædiktions-teknikker fra maskinlæring. Matematisk indsigt opnås via konkrete eksempler, og denne indsigt generaliseres til principper, der er gennemgående for resten af afhandlingen.

Det andet afsnit omhandler konceptet lokal uafhængighed, hvilket beskriver om udviklingen af en stokastisk proces er direkte påvirket af en anden proces. For at teste hypotesen om lokal uafhængighed definerer vi en model-fri parameter kaldet Local Covariance Measure (LCM). Vi beskriver en estimator af LCM, og baseret på denne foreslår vi nye test af lokal uafhængighed. Vi forklarer, hvordan størrelsen og styrken af disse test kan kontrolleres uniformt, og vi udforsker dem i et simulationsstudie.

Det tredje afsnit vedrører justering for kovariater, hvilket er en metode til at estimere effekten af en behandling ved at tage højde for observeret confounding. Vi udvikler en generel teori, der kan beskrive justering for enhver delmængde af information i kovariaterne. Vi identificerer den optimale information at justere for, og baseret på denne, introducerer vi Debiased Outcome-adapted Propensity Estimatoren (DOPE) for efficient estimering af behandlingseffekter. En version af DOPE implementeres med neurale netværk, og vi demonstrerer dens præstation på både simuleret og virkelig data.

Det fjerde og sidste afsnit introducerer en model-fri størrelse, der måler den betingede sammenhæng mellem en eksponering og en overlevelsestid, som vi kalder Aalen Covariance Measure (ACM). Man kan betragte ACM som en fleksibel generalisering af koefficienten for eksponering i en Aalen additive hazard model. Vi udvikler en model-fri estimationsmetode og viser at den er dobbelt robust, hvilket garanterer \sqrt{n} -konsistens under beskedne betingelser. Et simulationsstudie demonstrerer brugen af vores estimator i adskillige scenarier.

Contributions and Structure

Chapter 1 is an introduction, which briefly motivates model-free statistical inference based on machine learning. The introduction is not meant to be a literature review, but rather an effort to build intuition around some key principles that permeate the rest of the thesis. It is followed by 3 chapters, each of which corresponds to a paper. For reference within this thesis, we give each paper an acronym, for example [LCM]. Notation established in each chapter should be considered specific to the chapter where it is introduced.

Chapter 2 discusses hypothesis testing for time-to-event responses using the *Local Covariance Measure (LCM)* and corresponds to the paper:

[LCM] [Christgau et al., 2023b]. A. M. Christgau, L. Petersen, and N. R. Hansen. Non-parametric conditional local independence testing. *Annals of Statistics*, 51(5): 2116–2144, 2023b.

Chapter 3 formulates a general framework for the method of covariate adjustment, which leads to the *Debiased Outcome-adapted Propensity Estimator (DOPE)*, and corresponds to the paper:

[DOPE] [Christgau and Hansen, 2024]. A. M. Christgau and N. R. Hansen. Efficient adjustment for complex covariates: Gaining efficiency with DOPE. *arXiv preprint arXiv:2402.12980*, 2024.

Chapter 4 introduces a measure of the conditional association between an exposure and a time-to-event, namely the *Aalen Covariance Measure (ACM)*, and corresponds to the following paper in preparation:

[ACM] [Christgau and Hansen, 2024+] A. M. Christgau and N. R. Hansen. Assumption-lean Aalen regression. In preparation, 2024+.

Contributions and Structure

During my PhD studies, I also worked on the following two publications, which, however, are not included in this thesis.

1. A. M. Christgau, A. Arnaudon, and S. Sommer. Moment evolution equations and moment matching for stochastic image EPDiff. *Journal of Mathematical Imaging and Vision*, 65(4):563–576, 2023a.
2. A. Bianchi, A. M. Christgau, and J. S. Pedersen. On generators and relations of the rational cohomology of Hilbert schemes. *Journal of Algebraic Combinatorics*, 57(3):829–857, 2023

Contents

Preface	iv
Abstract	vi
Contributions and Structure	ix
1 Introduction	1
1.1 Model-free statistical objectives	2
1.2 Estimation	4
2 Nonparametric conditional local independence testing	15
2.1 Introduction	15
2.2 The Local Covariance Measure	20
2.3 Interpretations of the LCM estimator	28
2.4 General asymptotic results	32
2.5 The Local Covariance Test	37
2.6 Simulation study	40
2.7 Discussion	45
2.A Proofs of results in the main text	49
2.B Uniform stochastic convergence	68
2.C The Functional Martingale CLT	80
2.D Estimation of λ and G	83
2.E Relation to semiparametric survival models	86
2.F Details on Neyman orthogonality	87
2.G Additional details of simulation study	90
3 Efficient adjustment for complex covariates: Gaining efficiency with DOPE	95
3.1 Introduction	95
3.2 Generalized adjustment concepts	99
3.3 Efficiency bounds for adjusted means	101
3.4 Estimation based on outcome-adapted representations	107
3.5 Experiments	114
3.6 Discussion	123
3.A Auxiliary results and proofs	125
3.B Details of simulation study	136
3.C Extension to cross-fitting	139

Contents

4 Assumption-lean Aalen regression	141
4.1 Introduction	141
4.2 Target estimand and interpretations	145
4.3 Estimation	151
4.4 Asymptotic theory	153
4.5 Simulation study	158
4.6 Discussion	162
4.A Auxillary results and proofs	164
4.B Multiplicative hazards models	175
Bibliography	179

1 Introduction

This thesis delves into a range of statistical challenges, including event history analysis, statistical efficiency, hypothesis testing, effect estimation, and assumption-lean inference. Despite the variety of topics, a central theme is that the problems and methods are studied from a *model-free* perspective.

To understand what this means, suppose we are given a sample consisting of n observations Z_1, \dots, Z_n , and that our objective is to infer properties about the mechanisms that generated this data. To simplify matters, we may start by assuming that the observations are independent and identically distributed (i.i.d.) according to a distribution P , referred to as the *data-generating process (DGP)*. Whether this is reasonable assumption depends on the application, but it will be assumed throughout this thesis. Under the i.i.d. assumption, our objective can be reformulated in terms of inferring properties about P . To this end, it is convenient to let Z denote an auxiliary independent observation with distribution P .

Before we can understand and appreciate a model-free approach to statistics, we first follow the thought process of Rhonda Fisher, a (fictional) model-based statistician.

Example 1.0.1. Rhonda has obtained a sample in which each observation is an independent copy of the template observation $Z = (X, Y) \in \mathbb{R}^d \times \mathbb{R}$. Here $X = (X^1, \dots, X^d)$ is a d -dimensional covariate and Y is a real-valued response variable. She knows that without further assumptions, she can hardly infer anything about the relationship between the covariate and the response. With this consideration, Rhonda begins by positing a linear model:

$$Y = \alpha + \beta^\top X + \varepsilon, \quad \text{with } \varepsilon \sim \mathcal{N}(0, \sigma^2) \text{ and } \varepsilon \perp\!\!\!\perp X, \quad (1)$$

where $\alpha \in \mathbb{R}$, $\beta = (\beta_1, \dots, \beta_d) \in \mathbb{R}^d$, and $\sigma^2 > 0$ are unknown parameters. Conveniently, her favorite statistical software can ‘fit’ the linear model (1).

She now recalls that her primary objective is to describe the effect of X^1 on Y . In view of the linear model (1), she concludes that β_1 summarizes such an effect. Thus she decides to report the summary of the fitted model, which in particular includes an estimate of β_1 and a valid confidence interval. ♠

While Rhonda is a caricature, her line of reasoning is ubiquitous and applies to many applications using generalized linear models (GLMs) and Cox models, among other semi-parametric models. Whereas Rhonda is ‘lucky’ that her model conveniently summarizes her target of interest, others may choose a model with this objective in mind.

In some applications, for example in physics, a thorough understanding of the data-generating process may justify the belief that the data follow a (semi)parametric model.

1 Introduction

However, in more complex systems, e.g., those analyzed in the physiological and social sciences, (semi)parametric models are often chosen for simplicity and convenience, rather than being a consequence of principled reasoning.

Because the data-generating process is often more complicated than desired, this leads to a difficult trade-off between model misspecification due to simplicity, and impracticality due to complexity. Model misspecification may lead to bias in the results of the analysis, but a potentially more serious problem is the lack of interpretability of the target parameter for distributions that do not conform to the model. Other issues with model-based reasoning were pointed out in the seminal paper by Breiman [2001] and are also discussed by Hines et al. [2022], among others.

We will proceed to discuss how to define a model-free target parameter.

1.1 Model-free statistical objectives

The importance of explicitly defining a target of interest at the onset of statistical analyses has been increasingly highlighted, particularly within the field of causal inference [van der Laan and Rose, 2011]. A starting point is to consider a general collection \mathcal{P} of possible data-generating processes, that is, a collection of probability distributions over the sample space for a single observation. Although causal analyses require additional structure, this thesis will concentrate on observational quantities and their estimation.

We call a function an *estimand* if its domain is \mathcal{P} . For example, if each observation is real-valued, a possible estimand is the population average outcome

$$\mu: \mathcal{P} \rightarrow \mathbb{R}, \quad \mu(P) := \int zP(dz) = \mathbb{E}_P[Z]. \quad (2)$$

The statistician is responsible for defining an estimand that effectively summarizes the target of interest.

Note that the estimand μ in (2) requires the mild condition that the possible DGPs are integrable. Most estimands require similar assumptions to be well-defined, such as the existence of moments, the positivity of certain probabilities, absolute continuity, or regularity conditions on conditional means. These are properties that can be justified based on knowledge of the underlying DGP.

We say that an estimand is *model-based* if it explicitly requires a parametrization of a component of the DGPs and reduces the dimension of the component.¹ If an estimand is not model-based, we say that it is *model-free*, also commonly referred to as nonparametric in other literature.

Returning to the setting of Example 1.0.1, Rhonda defined her target parameter by restricting \mathcal{P} to the linear Gaussian model

$$\mathcal{P}^{\text{lin}} := \{P \in \mathcal{P} \mid \exists(\alpha, \beta, \sigma^2) \in \mathbb{R} \times \mathbb{R}^d \times \mathbb{R}_{>0}: \text{Equation (1) holds for } (X, Y) \sim P\}.$$

Rhonda's target can be interpreted as the model-based estimand $\tau_{\text{mb}}: \mathcal{P}^{\text{lin}} \rightarrow \mathbb{R}$ that inverts the parametrization in β_1 (technically, this requires that \mathcal{P} is initially restricted

¹This is, admittedly and intentionally, a vague definition.

enough to ensure that β_1 is *identifiable*). Critically, her statistical analysis becomes meaningless if the model is misspecified. That is, for a distribution $P \notin \mathcal{P}^{\text{lin}}$, the estimand τ_{mb} is undefined and it is *a priori* unclear what any estimator of τ_{mb} , or equivalently, β , will target. In view of this, model-free estimands have been increasingly advocated over model-based ones.

It turns out, however, that the model-based estimand τ_{mb} can be extended to a model-free estimand. For convenience, we let $X = (D, W)$, where $D := X^1$ denotes the covariate of interest, and $W := (X^2, \dots, X^d)$ represents the other covariates. Then we define the estimand $\tau_{\text{mf}}: \mathcal{P} \rightarrow \mathbb{R}$ given by

$$\tau_{\text{mf}}(P) := \frac{\mathbb{E}_P[(D - \mathbb{E}_P[D | W])Y]}{\mathbb{E}_P[(D - \mathbb{E}_P[D | W])^2]} = \frac{\mathbb{E}_P[\text{Cov}_P(D, Y | W)]}{\mathbb{E}_P[\text{Var}_P(D | W)]}, \quad (3)$$

where we use a subscript P to indicate that an operator is computed under the distribution P . This estimand unambiguously defines a target parameter without reference to a (semi)parametric model, and the conditions for being well-defined are clear from the expression, unlike the model-based estimand τ_{mb} .²

To see how τ_{mf} is a generalization of τ_{mb} , consider a distribution $P \in \mathcal{P}$ that follows the *partially linear model*:

$$\mathbb{E}_P[Y | D, W] = \beta_P D + g_P(W), \quad (4)$$

where $\beta_P \in \mathbb{R}$ and where g_P is an integrable function with respect to the distribution of W under P . This semiparametric model generalizes the linear model by allowing a more flexible dependency on the covariates W . By applying iterated expectations conditionally on W , we obtain that

$$\tau_{\text{mf}}(P) = \frac{\mathbb{E}_P[(D - \mathbb{E}_P[D | W])(\beta_P D + g_P(W))]}{\mathbb{E}_P[(D - \mathbb{E}_P[D | W])^2]} = \beta_P \frac{\mathbb{E}_P[D^2 - (\mathbb{E}_P[D | W])^2]}{\mathbb{E}_P[(D - \mathbb{E}_P[D | W])^2]} = \beta_P.$$

This shows that τ_{mf} reduces to the coefficient β_P in the (partially) linear model. Thus, we can more generally make sense of estimators that target τ_{mf} , while maintaining the same interpretation within the (partially) linear model.

There can exist multiple extensions of a model-based estimand. For example, another possible extension of τ_{mb} is the *average partial effect*, defined by

$$\tau_{\text{ape}}(P) = \mathbb{E}_P \left[\frac{\partial}{\partial d} \mathbb{E}_P[Y | D = d, W] \Big|_{d=D} \right]. \quad (5)$$

Within the partially linear model (4), we observe that the derivative inside the expectation equals β , and, consequently, τ_{ape} is also a model-free extension of τ_{mb} . Arguably, the estimand τ_{ape} offers more direct interpretations for causal analyses than τ_{mf} . However, τ_{mf} can also serve as a reasonable summary of conditional association, and estimators of τ_{mf} may be more robust and well-behaved than those of τ_{ape} [Lundborg and Pfister, 2023].

²Both estimands are well-defined under the moment conditions $\mathbb{E}_P[|DY|] < \infty$ and $\mathbb{E}_P[D^2] < \infty$ and the positivity condition $\mathbb{E}_P[\text{Var}_P(D | W)] > 0$.

1.2 Estimation

In this section, we explore how \sqrt{n} -consistent estimation of an estimand can be achieved, even when dealing with complex or high-dimensional underlying DGPs. A substantial body of work exists on nonparametric estimation using *influence functions*, with methods such as *one-step estimation* [Pfanzagl and Wefelmeyer, 1985, Bickel et al., 1998] and *Targeted Maximum Likelihood Estimation (TMLE)* [van der Laan and Rose, 2011]. This literature is well summarized in tutorial papers by, e.g., Hines et al. [2022] and Kennedy [2022].

While the theory of influence functions is useful for defining an effective estimator, rigorously establishing the distributional properties of the estimator is a task that is typically done through case-specific analysis. In this introduction, we focus on this task through direct and concrete analyses, without invoking influence functions.

Drawing from Chernozhukov et al. [2018], we begin with a motivating example, emphasizing an estimand-based perspective rather than a semiparametric (score-based) approach. After discussing the example, we will explore how these insights can be extended to more general contexts.

1.2.1 An instructive example

Consider again the setting where the observed sample Z_1, \dots, Z_n consists of i.i.d. copies of $Z = (D, W, Y) \in \mathbb{R} \times \mathbb{R}^{d-1} \times \mathbb{R}$ with $Z \sim P \in \mathcal{P}$. Suppose our target of interest is the numerator of (3), i.e., the estimand $\rho: \mathcal{P} \rightarrow \mathbb{R}$ given by

$$\rho(P) := \mathbb{E}_P[(D - \mathbb{E}_P[D | W])Y]. \quad (6)$$

This estimand has been studied for the purpose of quantifying and testing conditional (mean) independence [Shah and Peters, 2020, Lundborg, 2023]. We shall assume that D, Y , and DY are square-integrable under each $P \in \mathcal{P}$, which in particular ensures that ρ is well-defined.

Expression (6) combines two components of the distribution: the expectation operator, $\mathbb{E}_P[\cdot]$, and the conditional mean function, $m_P(\cdot) := \mathbb{E}_P[D | W = \cdot]$. According to the central limit theorem (CLT), the expectation operator is pointwise \sqrt{n} -approximated by the empirical mean operator, denoted $\mathbb{P}_n[\cdot]$ and given by $\mathbb{P}_n[f(Z)] := \frac{1}{n} \sum_{i=1}^n f(Z_i)$, for any function $f \in \mathcal{L}^2(P)$.

Distributional quantities that are used to compute the target estimand, but are not of primary interest, such as m_P , are called *nuisance parameters*. Estimating conditional means can be done using various regression methods, including model-free prediction algorithms from machine learning (ML), which we discuss further in Section 1.2.3. Assume, for now, that an estimate \hat{m} of m_P is given. This suggests the estimator

$$\hat{\rho}^{\text{plug-in}} := \mathbb{P}_n[(D - \hat{m}(W))Y] = \frac{1}{n} \sum_{i=1}^n (D_i - \hat{m}(W_i))Y_i.$$

Estimators like $\hat{\rho}^{\text{plug-in}}$, which are derived by substituting nuisance estimates directly into the expression for the estimand, are often referred to as *plug-in* estimators. However,

it is important to note that such estimators are not necessarily of the form $\rho(\hat{P})$ for an estimate $\hat{P} \in \mathcal{P}$ of the entire distribution P , which is a commonly used representation for plug-in estimators. As we will see shortly, there can be multiple valid expressions for the same estimand, resulting in different plug-in estimators.

To understand the behavior of this estimator, consider the following decomposition

$$\begin{aligned} \sqrt{n}(\hat{\rho}^{\text{plug-in}} - \rho(P)) &= \sqrt{n}\mathbb{P}_n[(D - \hat{m}(W))Y] - \sqrt{n}\mathbb{E}_P[(D - m_P(W))Y] \\ &= \underbrace{\sqrt{n}(\mathbb{P}_n - \mathbb{E}_P)[(D - m_P(W))Y]}_{=: U^{\text{plug-in}}} + \underbrace{\sqrt{n}\mathbb{P}_n[(m_P(W) - \hat{m}(W))Y]}_{=: R^{\text{plug-in}}}. \end{aligned}$$

The term $U^{\text{plug-in}}$ is an *oracle term* that describes the error of the plug-in estimator based on the true conditional mean $m_P(\cdot)$. The CLT asserts that $U^{\text{plug-in}}$ converges in distribution to a Gaussian distribution with mean zero.

The term $R^{\text{plug-in}}$ is a sum of n conditionally i.i.d. terms that are controlled by the estimation error of $m_P(\cdot)$. As we discuss in Section 1.2.3, we can generally expect the root mean squared error (RMSE) of $\hat{m}(\cdot)$ to converge at a rate of $n^{-\alpha}$, where $\alpha \in (0, \frac{1}{2})$. Consequently, if $\mathbb{E}_P[Y | W] \neq 0$, then we can expect $R^{\text{plug-in}}$ to be of stochastic order $n^{\frac{1}{2}-\alpha}$, in which case

$$|\sqrt{n}(\hat{\rho}^{\text{plug-in}} - \rho(P))| \xrightarrow{P} \infty.$$

This is, of course, undesirable and makes inference about $\rho(P)$ based on $\hat{\rho}^{\text{plug-in}}$ intractable.

Fortunately, the bias in $R^{\text{plug-in}}$ can be fixed by adding a correction term in the expression. Motivated by the fact that $R^{\text{plug-in}}$ becomes centered when $\mathbb{E}_P[Y | W] = 0$, we rewrite the estimand as

$$\rho(P) = \mathbb{E}_P[(D - \mathbb{E}_P[D | W])(Y - \mathbb{E}_P[Y | W])]. \quad (7)$$

This introduces a new nuisance parameter, $g_P(\cdot) := \mathbb{E}_P[Y | W = \cdot]$, which can be estimated similarly to $m_P(\cdot)$, e.g., using an ML algorithm, yielding another estimate $\hat{g}(\cdot)$. This suggests the following estimator,

$$\hat{\rho}^{\text{double}} = \mathbb{P}_n[(D - \hat{m}(W))(Y - \hat{g}(W))] = \frac{1}{n} \sum_{i=1}^n (D_i - \hat{m}(W_i))(Y_i - \hat{g}(W_i)). \quad (8)$$

To understand why $\hat{\rho}^{\text{double}}$ might be preferable to $\hat{\rho}^{\text{plug-in}}$, we can decompose its estimation error as follows:

$$\sqrt{n}(\hat{\rho}^{\text{double}} - \rho(P)) = U + R_m + R_g + R_{mg} \quad (9)$$

where

$$\begin{aligned} U &= \sqrt{n}(\mathbb{P}_n - \mathbb{E}_P)[(D - m_P(W))(Y - g_P(W))], \\ R_m &= \sqrt{n}\mathbb{P}_n[(m_P(W) - \hat{m}(W))(Y - g_P(W))], \\ R_g &= \sqrt{n}\mathbb{P}_n[(D - m_P(W))(g_P(W) - \hat{g}(W))], \\ R_{mg} &= \sqrt{n}\mathbb{P}_n[(m_P(W) - \hat{m}(W))(g_P(W) - \hat{g}(W))]. \end{aligned}$$

1 Introduction

The first term, U , is again an oracle term that is asymptotically Gaussian, that is, $U \xrightarrow{d} \mathsf{N}(0, \sigma_P^2)$ with asymptotic variance $\sigma_P^2 := \text{Var}_P((D - m_P(W))(Y - g_P(W)))$. We will argue that, under suitable conditions, the error terms R_m, R_g , and R_{mg} converge to zero in probability, in which case Slutsky's theorem asserts that

$$\sqrt{n}(\hat{\rho}^{\text{double}} - \rho(P)) \xrightarrow{d} \mathsf{N}(0, \sigma_P^2). \quad (10)$$

There are several strategies for controlling a term like R_m , which we elaborate further in Section 1.2.4. For simplicity, let us assume that $\hat{m}(\cdot)$ is estimated on an auxiliary sample that is independent from the sample $(Z_i)_{i \in [n]}$ used in \mathbb{P}_n , where $[n] := \{1, \dots, n\}$. Then, the summands in R_m are independent conditionally on \hat{m} and $(W_i)_{i \in [n]}$, and their conditional means are zero since $\mathbb{E}[Y_i | \hat{m}, (W_j)_{j \in [n]}] = \mathbb{E}[Y_i | W_i] = g_P(W_i)$ for $i \in [n]$. Note that the first equality crucially relies on \hat{m} being estimated on an independent sample.

It follows that the conditional variance of R_m is given by

$$\text{Var}_P(R_m | \hat{m}, (W_i)_{i \in [n]}) = \frac{1}{n} \sum_{i=1}^n (\hat{m}(W_i) - m_P(W_i))^2 \text{Var}_P(Y_i | W_i).$$

Thus, if there exists a bound $C > 0$ such that $\text{Var}_P(Y | W) \leq C$ almost surely and if \hat{m} is (out-of-sample) RMSE consistent for m_P , then it holds that $\text{Var}_P(R_m | \hat{m}, (W_i)_{i \in [n]}) \xrightarrow{P} 0$. In this case, an application of Chebyshev's inequality, conditionally on \hat{m} and $(W_i)_{i \in [n]}$, lets us to conclude that $R_m \xrightarrow{P} 0$.

Since R_m and R_g are symmetric in (m_P, D) and (g_P, Y) , the same argument can be applied to show that $R_g \xrightarrow{P} 0$ under analogous conditions on \hat{g} and $\text{Var}_P(D | W)$.

For the product error term, R_{mg} , the Cauchy-Schwarz inequality yields that

$$R_{mg} \leq \sqrt{n} \cdot \sqrt{\mathbb{P}_n[(\hat{m}(W) - m_P(W))^2]} \cdot \sqrt{\mathbb{P}_n[(\hat{g}(W) - g_P(W))^2]}. \quad (11)$$

As a consequence, we see that $R_{mg} \xrightarrow{P} 0$ if the product of the RMSE for $\hat{m}(\cdot)$ and the RMSE for $\hat{g}(\cdot)$ decays at an order of $o_P(n^{-1/2})$. When the rate requirements for nuisance estimators reduce to this condition, we say that the estimator exhibits *rate double robustness*. There are several settings where this requirement can be achieved by nonparametric regression estimators, as we elaborate in Section 1.2.3.

We proceed to discuss, from a more general perspective, conditions under which an estimator may exhibit (rate) double robustness.

1.2.2 Double robustness and orthogonality

In this section, we generalize some of the insights obtained from the example in the preceding section, and we derive a condition required for rate double robustness.

Suppose that $\tau: \mathcal{P} \rightarrow \mathbb{R}$ is an estimand that can be written of the form

$$\tau(P) = \frac{\mathbb{E}_P[\psi(Z, \eta_0(P))]}{\mathbb{E}_P[\phi(Z, \eta_0(P))]}, \quad (12)$$

where $\eta_0: \mathcal{P} \rightarrow \mathcal{E}$ is a nuisance parameter, and where $\phi, \psi: \mathcal{Z} \times \mathcal{E} \rightarrow \mathbb{R}$ are measurable maps, with \mathcal{Z} denoting the sample space for Z . We assume that for all $\eta \in \mathcal{E}$ and $P \in \mathcal{P}$, the random variables $\psi(Z, \eta)$ and $\phi(Z, \eta)$ are integrable under P with $\mathbb{E}_P[\phi(Z, \eta)] \neq 0$. Note that all the estimands considered so far have been of this form. An estimand may be written of the form (12) in more than one way, as was highlighted in Section 1.2.1.

Given a nuisance estimate $\hat{\eta} \in \mathcal{E}$ of $\eta_0(P)$, we can define a plug-in estimator more generally as:

$$\hat{\tau}(\hat{\eta}) = \frac{\mathbb{P}_n[\psi(Z, \hat{\eta})]}{\mathbb{P}_n[\phi(Z, \hat{\eta})]} \quad (13)$$

Both $\hat{\rho}^{\text{plug-in}}$ and $\hat{\rho}^{\text{double}}$ were conceived in this way, so why does $\hat{\rho}^{\text{double}}$ exhibit rate double robustness while $\hat{\rho}^{\text{plug-in}}$ does not?

To answer this question, consider first an estimand that can be written of the form (12) with $\phi(\cdot, \cdot) \equiv 1$, and let $\eta_P := \eta_0(P)$. We can then decompose the general plug-in estimator (13) similarly to the example in Section 1.2.1:

$$\sqrt{n}(\hat{\tau}(\hat{\eta}) - \tau(P)) = \sqrt{n}(\mathbb{P}_n - \mathbb{E}_P)[\psi(Z, \eta_P)] + \sqrt{n}\mathbb{P}_n[\psi(Z, \hat{\eta}) - \psi(Z, \eta_P)]$$

The first term is again an oracle term, whose behavior is governed by the CLT.

To analyze the second term, we may appeal to a Taylor expansion. Assume that \mathcal{E} is a convex set, and assume that the map $r \mapsto \psi(z, \eta_P + r(\hat{\eta} - \eta_P))$ is two times continuously differentiable for all $(z, \eta) \in \mathcal{Z} \times \mathcal{E}$. Then, a Taylor expansion yields that

$$\sqrt{n}\mathbb{P}_n[\psi(Z, \hat{\eta}) - \psi(Z, \eta_P)] \quad (14)$$

$$= \underbrace{\sqrt{n}\mathbb{P}_n[\partial_r \psi(Z, \eta_P + r(\hat{\eta} - \eta_P))|_{r=0}]}_{=: T_1} + \underbrace{\sqrt{n}\mathbb{P}_n[R_2(Z, \eta_P, \hat{\eta})]}_{=: T_2}, \quad (15)$$

where $R_2(Z, \eta_P, \hat{\eta})$ is a second-order remainder term that is left unspecified for now.

As we will discuss in Section 1.2.3, we can generally expect a nuisance estimator to converge at a rate of $n^{-\alpha}$, where $\alpha \in (0, \frac{1}{2})$. To simplify the analysis, suppose that $\hat{\eta} - \eta_P = n^{-\alpha} \Delta\eta$ for a fixed $\Delta\eta \in \mathcal{E}$. Although this might seem like an oversimplification, if the plug-in estimator is \sqrt{n} -consistent for general nuisance estimates, we can reasonably expect it to be \sqrt{n} -consistent under this simplifying assumption.

Then, by the chain rule,

$$T_1 = n^{\frac{1}{2}-\alpha}(\mathbb{P}_n - \mathbb{E}_P)[\partial_r \psi(Z, \eta_P + r\Delta\eta)|_{r=0}] + n^{\frac{1}{2}-\alpha} \partial_r \mathbb{E}_P[\psi(Z, \eta_P + r\Delta\eta)]|_{r=0}$$

provided we can interchange derivative and expectation. The first term vanishes in probability, whereas the absolute value of the second term is either diverging to infinity or equal to zero for all n . In conclusion, we can expect T_1 to vanish in probability if for all $\eta \in \mathcal{E}$:

$$\partial_r \mathbb{E}_P[\psi(Z, \eta_P + r(\eta - \eta_P))|_{r=0}] = 0 \quad (16)$$

We can interpret the condition (16) as the bias of the plug-in estimator not being sensitive to small estimation errors of the nuisance parameter η_0 .

1 Introduction

The condition (16) is similar to a condition at the core of *Double Machine Learning (DML)* [Chernozhukov et al., 2018]. Indeed, we can explicitly translate (12) into the DML framework by considering $(\tau_P, \eta_P) := (\tau(P), \eta_0(P))$ as solutions to the *score equation*

$$\mathbb{E}_P[S(Z; \theta, \eta)] = 0, \quad \text{where } S(Z; \theta, \eta) := \theta \cdot \phi(Z, \eta) - \psi(Z, \eta). \quad (17)$$

Definition 2.1 in Chernozhukov et al. [2018] states that the score S satisfies the *Neyman orthogonality condition* at (τ_P, η_P) if $\partial_r \mathbb{E}_P[S(Z; \tau_P, \eta_P + r(\eta - \eta_P))]$ exists for all $\eta \in \mathcal{E}$ and $r \in [0, 1)$, and if

$$\partial_r \mathbb{E}_P[S(Z; \tau_P, \eta_P + r(\eta - \eta_P))] |_{r=0} = 0. \quad (18)$$

This is essentially equivalent to (16) and $\partial_r \mathbb{E}_P[\phi(Z, \eta_P + r(\eta - \eta_P))] |_{r=0} = 0$. Orthogonality is paramount in order for an estimator $\hat{\tau}$ of the form (13) to satisfy that $\sqrt{n}(\hat{\tau} - \tau_P) \xrightarrow{d} N(0, \sigma_P^2)$ for some $\sigma_P^2 > 0$. It is straightforward to verify that the expression in (7) corresponds to a score that satisfies the orthogonality condition, whereas the expression (6) does not, unless $g_P \equiv 0$. Heuristically, this explains why $\hat{\psi}^{\text{double}}$, which is based on (7), exhibits rate double robustness.

Orthogonality alone is not sufficient to ensure asymptotic normality, and Chernozhukov et al. [2018] provide additional conditions that ensure this; see their Assumptions 3.1 and 3.2. However, these conditions essentially entail a reformulation of the condition that $T_2 \xrightarrow{P} 0$, cf. Assumption 3.2(c). Obtaining a general bound for T_2 with minimal assumptions is a complex task, so T_2 is typically analyzed case-by-case. When the nuisance parameter $\eta_P = (\eta_{P,1}, \dots, \eta_{P,q})$ consists of several nuisance functions, it is often possible to establish a bound in the form:

$$T_2 \leq O_P\left(\sqrt{n} \sum_{(i,j) \in \mathcal{S}} \|\hat{\eta}_i - \eta_{P,i}\| \cdot \|\hat{\eta}_j - \eta_{P,j}\|\right), \quad (19)$$

for a set $\mathcal{S} \subseteq [q] \times [q]$. This represents a general form of rate double robustness, which is met, for example, when $\|\hat{\eta}_i - \eta_{P,i}\| = o_P(n^{-1/4})$ for each $i \in [q]$. The bound (11), for instance, is of this form.

Keeping the orthogonality condition – or analogous conditions in other frameworks – in mind when constructing an estimator is crucial. This approach is taken in [LCM], [ACM], and [DOPE], though the details of the latter are more intricate.

1.2.3 Estimation of nuisance parameters using machine learning

The estimation procedures discussed up to this point have relied on the ability to estimate the nuisance parameter, and the analysis suggested that an estimation error of order $n^{-1/4}$ would be sufficient. Often, the nuisance parameter consists of functions given by conditional expectations, as in Section 1.2.1 with $\eta(P) = (m_P, g_P)$.

To discuss estimation of such nuisance functions, consider a generic regression setting where $Z = (X, Y) \in \mathbb{R}^d \times \mathbb{R}$ and suppose that $f_P(\cdot) := \mathbb{E}_P[Y | X = \cdot]$ is the nuisance

parameter. Assuming that $\mathbb{E}_P[Y^2] < \infty$, recall that

$$f_P \in \arg \min_{f \in \mathcal{F}} \mathbb{E}_P[(Y - f(X))^2]$$

for any collection of measurable functions \mathcal{F} that contains f_P . This connects the estimation of the conditional mean function, f_P , with the minimization of the mean squared prediction error over a sufficiently rich function class.

When f_P is known to be linear, this naturally leads to ordinary least squares, and the resulting estimator for f_P converges at a rate of $O_P(n^{-1/2})$ under a few additional conditions. Similarly, other classical regression techniques based on semiparametric models that represent \mathcal{F} based on a finite-dimensional representation, such as GLMs, generally yield \sqrt{n} -consistent estimators of f_P when the model is correctly specified.

However, contemporary datasets and regression competitions, such as those hosted by Kaggle [Bojer and Meldgaard, 2021], indicate that ML methods, e.g. random forests, gradient boosting, and neural nets, routinely and significantly outperform classical regression methods in terms of prediction accuracy. This suggests that the function classes corresponding to classical methods are not rich enough to capture good approximations of the truth.

Considering larger nonparametric function classes can come with two challenges: it can make the resulting optimization problem difficult and usually leads to overfitting if implemented naively. Remarkable advancements in machine learning have resulted in powerful and efficient tools for handling optimization tasks with few assumptions on the function class. To prevent overfitting, these methods often depend on advanced (implicit) regularization techniques, making them difficult to analyze theoretically.

In view of the above, it is often left as a general assumption that the nuisance parameter can be estimated with a sufficiently fast rate. This assumption can be justified based on (i) the empirical observation that machine learning estimators perform well in practice and (ii) by comparing with achievable (minimax) rates for reasonable function classes. To elaborate on (ii), suppose, for example, that f_P is known to be differentiable up to order $\lfloor s \rfloor$, with partial derivatives of order $\lfloor s \rfloor$ that are $(s - \lfloor s \rfloor)$ -Hölder continuous, for some $s > 0$. In the absence of additional constraints on f_P , the optimal uniform estimation rate is $O_P(n^{-\frac{s}{2s+d}})$, which can be attained using local polynomials [Györfi et al., 2002, Tsybakov, 2009]. The key insight is not necessarily to use local polynomials, but to understand that effective estimators of f_P can attain an estimation rate, e.g., of order $n^{-1/4}$, when f_P has a smoothness of order $s > d/2$.

While the discussion above revolves around conditional mean estimation, other estimation tasks might also be of interest. Estimation of the average partial effect (5) with DML requires density and derivative estimation [Klyne and Shah, 2023], and in both [LCM] and [ACM] we require estimation of conditional hazards. Such estimation methods are discussed within each manuscript, see for example Section 2.D in [LCM].

1.2.4 Controlling cross-terms by cross-fitting

In this section we elaborate the discussion of how to control *cross-terms* such as R_m and R_g from the decomposition (9) of $\sqrt{n}(\hat{\rho} - \rho(P))$. More generally, in Section 1.2.2 we

1 Introduction

encountered the ‘first-order’ term

$$T_1 = \sqrt{n} \mathbb{P}_n [\partial_r \psi(Z, \eta_P + r(\hat{\eta} - \eta_P))|_{r=0}].$$

By simplifying the analysis and considering the deterministic sequence of nuisance estimates $\hat{\eta} = \eta_P + n^{-\alpha} \Delta \eta$ for a fixed $\Delta \eta \in \mathcal{E}$, we derived that the orthogonality condition (16) was essential for achieving $T_1 \xrightarrow{P} 0$. We now discuss how to proceed when $\hat{\eta}$ is a stochastic estimate, e.g., obtained by machine learning as described in Section 1.2.3.

When $\hat{\eta}$ is estimated independently of the sample in \mathbb{P}_n , the orthogonality condition (16) implies that T_1 has mean zero conditionally on $\hat{\eta}$. In many concrete examples, we can then control the conditional variance of T_1 by the error in nuisance estimation, in which case the convergence $T_1 \xrightarrow{P} 0$ follows from consistency of the nuisance estimator and Chebyshev’s inequality. For example, this was the case for the estimator $\hat{\rho}^{\text{double}}$ in (8), where it is straightforward to show that $T_1 = R_m + R_g$.

Obtaining an independent nuisance estimate can be achieved by initially splitting the sample into two *folds*, and then using first part for nuisance estimation and the second part for plug-in estimation. This comes at the cost of reducing the effective sample size. However, it is possible to recover full-sample efficiency by swapping the roles of the two folds and then aggregating the resulting estimates. More generally, we can perform *K-fold cross-fitting* for any integer $K \geq 2$ as follows:

- i) Partition the observation indices into K disjoint folds: $[n] = I_1 \cup \dots \cup I_K$.
- ii) For each $k \in [K]$, use the holdout data $(Z_i)_{i \in [n] \setminus I_k}$ to compute an ML estimate $\hat{\eta}_k$.
- iii) Compute the plug-in estimates on each fold and aggregate the estimates by

$$\check{\tau} := \frac{1}{K} \sum_{k=1}^K \underbrace{\frac{\sum_{i \in I_k} \psi(Z_i, \hat{\eta}_k)}{\sum_{i \in I_k} \phi(Z_i, \hat{\eta}_k)}}_{\hat{\tau}_k}, \quad \text{or} \quad \tilde{\tau} := \frac{\frac{1}{K} \sum_{k=1}^K \sum_{i \in I_k} \psi(Z_i, \hat{\eta}_k)}{\frac{1}{K} \sum_{k=1}^K \sum_{i \in I_k} \phi(Z_i, \hat{\eta}_k)}. \quad (20)$$

The estimators $\check{\tau}$ and $\tilde{\tau}$ are known as the DML1 and DML2 estimators, respectively (see Definitions 3.1 and 3.2 in Chernozhukov et al. [2018]). The estimators are asymptotically equivalent (under conditions) and they are equal when ϕ is a constant, which was the case in (7). In small samples, however, DML2 may perform better due to the increased numerical stability achieved by pooling the estimates before division.

The DML1 estimator is considered in all the manuscripts in this thesis; therefore, we present a simple result for its asymptotics.

Lemma 1.2.1 (Asymptotics of DML1 estimator). *Let $\check{\tau}$ and $(\hat{\tau}_k)_{k \in [K]}$ be the estimators defined in (20) and let $\sigma_P^2 > 0$. Suppose that each $k \in [K]$,*

$$\lim_{n \rightarrow \infty} n/|I_k| = K \quad \text{and} \quad \sqrt{|I_k|} \cdot (\hat{\tau}_k - \tau_P) = U_k + R_k,$$

where U_k and R_k are variables satisfying that $U_k \xrightarrow{d} \mathcal{N}(0, \sigma_P^2)$ and $R_k \xrightarrow{P} 0$ as $n \rightarrow \infty$. Assume also the joint independence statement $U_1 \perp \dots \perp U_K$. Then it holds that

$$\sqrt{n}(\check{\tau} - \tau_P) \xrightarrow{d} \mathcal{N}(0, \sigma_P^2)$$

as $n \rightarrow \infty$.

Proof. The conditions on I_k, U_k , and R_k imply that

$$\tilde{U}_k := \frac{\sqrt{n}}{\sqrt{K|I_k|}} U_k \xrightarrow{d} \mathcal{N}(0, \sigma_P^2) \quad \text{and} \quad \tilde{R}_k := \frac{\sqrt{n}}{\sqrt{K|I_k|}} R_k \xrightarrow{P} 0.$$

By joint independence, it follows that $(\tilde{U}_1, \dots, \tilde{U}_K) \xrightarrow{d} \mathcal{N}(0, \sigma_P^2)^{\otimes K}$ [Billingsley, 2013, Theorem 2.8(ii)]. By continuous mapping and the convolution property of the Gaussian distribution, we conclude that

$$\sqrt{n}(\check{\tau} - \tau_P) = \frac{1}{\sqrt{K}} \sum_{k=1}^K (\tilde{U}_k + \tilde{R}_k) \xrightarrow{d} \mathcal{N}(0, \sigma_P^2)$$

as $n \rightarrow \infty$. □

In practice, we can apply the lemma with oracle terms (U_k) corresponding to using the oracle nuisance estimate η_P . In this case, it simplifies the analysis of the DML1 estimator to that of a single sample split estimator $\hat{\tau}_k$.

Note that the DML1 estimator indeed recovers full-sample efficiency, as it is scaled by \sqrt{n} rather than $\sqrt{|I_k|}$. In particular, the asymptotic distribution is independent of the number of folds K . The choice of K will have a practical impact in smaller samples, as it will affect the remainder terms $(R_k)_{k \in [K]}$, which are controlled by the nuisance estimation error. Larger values of K will allocate more data to the estimation of the nuisance parameter and are thus likely to result in smaller error terms. However, this also comes at the cost of higher computational demand since the nuisance estimator has to be fitted K times. Choosing $K = 4$ or $K = 5$ might be reasonable in practice, cf. Remark 3.1. in Chernozhukov et al. [2018], which was also consistent with the findings in the simulation studies conducted for the work in this thesis.

1.2.5 Alternatives to cross-fitting

In the preceding discussion, we have covered how cross-fitting can be used to control the cross-term. This raises the question:

$$\textit{Is cross-fitting necessary?} \tag{21}$$

There is no simple or general answer to this question, and it is instead typically discussed on a case-specific basis. In [LCM] and [ACM], we found that cross-fitting was necessary, and for [DOPE] we found that cross-fitting was neither necessary nor significantly detrimental.

1 Introduction

To understand how a cross-term might be controlled without cross-fitting, consider again the term R_m from (9), but now as a map

$$R_m(\bar{m}) := \frac{1}{\sqrt{n}} \sum_{i=1}^n (m_P(W_i) - \bar{m}(W_i))(Y_i - g_P(W_i))$$

which takes as input any measurable function $\bar{m}: \mathbb{R}^{d-1} \rightarrow \mathbb{R}$. Assuming for simplicity that $|Y| \leq C$ for some constant $C > 0$, the CLT implies that for each $\bar{m} \in \mathcal{L}^2(P_W)$,

$$R_m(\bar{m}) \xrightarrow{d} N(0, \nu_P(\bar{m})) \quad (22)$$

as $n \rightarrow \infty$, where $\nu_P(\bar{m}) := \mathbb{E}_P[(m_P(W) - \bar{m}(W))^2(Y - g_P(W))^2]$.

The question (21) is now whether $R_m(\hat{m}) \xrightarrow{P} 0$ can be established from a consistency condition of \hat{m} , e.g., $\nu_P(\hat{m}) \xrightarrow{P} 0$. Unfortunately, this is not generally true, but it can be concluded under additional assumptions on \hat{m} .

Donsker class conditions: Let $\mathcal{F}_m \subset \mathcal{L}^2(P_W)$ denote a function space such that $\hat{m} \in \mathcal{F}_m$ almost surely. A sufficient condition for the implication

$$\nu_P(\hat{m}) \xrightarrow{P} 0 \quad \implies \quad R_m(\hat{m}) \xrightarrow{P} 0 \quad (23)$$

to hold is that the collection of random variables

$$\{(m_P(W) - \bar{m}(W))(Y - g_P(W)) : \bar{m} \in \mathcal{F}_m\}$$

is a Donsker class. See, for example, Lemma 19.24 in van der Vaart [2000]. Informally, this Donsker class condition means that the weak convergence (22) holds uniformly over $\bar{m} \in \mathcal{F}_m$. There are numerous ways to establish the Donsker condition in itself, which typically involve technical bounds on the ‘size’ of the function class \mathcal{F}_m . However, these conditions are unsuitable for high-dimensional settings [Chernozhukov et al., 2018] or when the nuisance estimates belong to a complex function class.

Algorithmic stability: More recently, Chen et al. [2022] have shown that cross-fitting can be avoided when the algorithm for producing the nuisance estimates is *stable*. Informally, this means that the learned function does change significantly when the training data is perturbed, and there are several different definitions that formalize this idea.

For the cross-term R_m , Proposition S15 in Lundborg et al. [2022a] describes algorithmic stability conditions that guarantee $R_m(\hat{m}) \xrightarrow{d} 0$ without the need for cross-fitting. It is possible to verify that certain algorithms are stable, see, for example, Hardt et al. [2016], Bousquet and Elisseeff [2002], and recent work has shown that *bagging* can increase the stability of any algorithm [Soloff et al., 2024b,a]. These works, however, are not all working under the same notion of stability, and it is therefore unclear to what extent bagging can be used as a replacement for cross-fitting.

Donsker class conditions and algorithmic stability have not been investigated in depth in this thesis. However, the analysis of cross-terms plays a central role in the asymptotic theory, and these terms are controlled by cross-fitting or the aforementioned alternative conditions. Since the answer to question (21) appears to be yes in both [LCM] and [ACM], these alternative conditions must somehow be violated. A deeper understanding of why these conditions are violated could provide valuable insights into our methods.

2 Nonparametric conditional local independence testing

ALEXANDER MANGULAD CHRISTGAU, LASSE PETERSEN AND NIELS RICHARD HANSEN

Abstract

Conditional local independence is an asymmetric independence relation among continuous time stochastic processes. It describes whether the evolution of one process is directly influenced by another process given the histories of additional processes, and it is important for the description and learning of causal relations among processes. We develop a model-free framework for testing the hypothesis that a counting process is conditionally locally independent of another process. To this end, we introduce a new functional parameter called the Local Covariance Measure (LCM), which quantifies deviations from the hypothesis. Following the principles of double machine learning, we propose an estimator of the LCM and a test of the hypothesis using nonparametric estimators and sample splitting or cross-fitting. We call this test the (cross-fitted) Local Covariance Test ((X)-LCT), and we show that its level and power can be controlled uniformly, provided that the nonparametric estimators are consistent with modest rates. We illustrate the theory by an example based on a marginalized Cox model with time-dependent covariates, and we show in simulations that when double machine learning is used in combination with cross-fitting, then the test works well without restrictive parametric assumptions.

2.1 Introduction

Notions of how one variable influences a target variable are central to both predictive and causal modeling. Depending on the objective, the relevant notion of influence can be variable importance in a predictive model of the target, but it can also be the causal effect of the variable on the target. In either case, we can investigate influence conditionally on a third variable – to quantify the added predictive value, the direct causal effect or the causal effect adjusted for a confounder. Our interests are in an asymmetric notion of direct influence among stochastic processes, which is not adequately captured by classical (symmetric) notions of conditional dependence. The objective of this paper is therefore to quantify this notion of asymmetric influence and specifically to develop a

2 Nonparametric conditional local independence testing

new nonparametric test of the hypothesis that one stochastic process does not directly influence another.

The hypothesis we consider is that of *local independence* – a concept introduced by Schweder [1970] for Markov processes as a continuous time formalization of the phenomenon that the past of one stochastic process does not directly influence the evolution of another stochastic process. Generalizations to other continuous time processes were given by Aalen [1987] and studied by Commenges and Gégout-Petit [2009], who systematically used the term *conditional local independence* for the general concept. We will in this paper follow that convention whenever we want to emphasize the conditional nature of the local independence. We note that (conditional) local independence is a continuous time version of the discrete time concept of Granger non-causality [Granger, 1969, Aalen, 1987].

To illustrate the concept of conditional local independence we will in this introduction consider an example involving three processes: X , Z and N – see Figure 2.1.1. The process N is the indicator of death, $N_t = \mathbb{1}(T \leq t)$, for an individual with survival time T , and X_t denotes the total pension savings of the individual at time t . The process Z is a covariate process, e.g., health variables or employment status, that may directly affect both the pension savings and the survival time. This is indicated in Figure 2.1.1 by edges pointing from Z to X and N . Edges pointing from N to X and Z indicate that a death event directly affects both X and Z (which take the values X_T and Z_T , respectively, after time T , see Section 2.2.2).

To define conditional local independence let $\mathcal{F}_t^{N,Z} = \sigma(N_s, Z_s; s \leq t)$ denote the filtration generated by the N - and Z -processes. The σ -algebra $\mathcal{F}_t^{N,Z}$ represents the information contained in the N - and the Z - processes before time t . Informally, the process N_t is conditionally locally independent of the process X_t given $\mathcal{F}_t^{N,Z}$ if $(X_s)_{s \leq t}$ does not add predictable information to $\mathcal{F}_{t-}^{N,Z}$ about the infinitesimal evolution of N_t . For this particular example this means that the conditional hazard function of T does not depend on $(X_s)_{s \leq t}$ given $\mathcal{F}_t^{N,Z}$. In Figure 2.1.1 the hypothesis of interest, that N_t is conditionally locally independent of X_t given $\mathcal{F}_t^{N,Z}$, is represented by the lack of an edge from X to N .

A systematic investigation of algebraic properties of conditional local independence was initiated by Didelez [2006, 2008, 2015]. She also introduced local independence graphs, such as the directed graph in Figure 2.1.1, to graphically represent all conditional local independencies among several processes, and she studied the semantics of these graphs. This work was extended further by Mogensen and Hansen [2020] to graphical representations of partially observed systems. While we will not formally discuss local independence graphs, the problem of learning such graphs from data was an important motivation for us to develop a nonparametric test of conditional local independence. A constraint based learning algorithm of local independence graphs was given by Mogensen et al. [2018] in terms of a conditional local independence oracle, but a practical algorithm requires that the oracle is replaced by conditional local independence tests.

Another important motivation for considering conditional local independence arises from causal models. With a structural assumption about the stochastic process specifi-




Figure 2.1.1: Local independence graph illustrating a dependence structure among the three processes X , Z and N . Here N is the indicator of death for an individual, X is their cumulative pension savings and Z is a covariate process. All nodes in this graph have implicit self-loops. There is no edge from X to N , which indicates that death is not directly influenced by pension savings. This can be formalized as N being conditionally locally independent of X , which is the hypothesis we aim to test.

cation, a conditional local independence has a causal interpretation [Aalen, 1987, Aalen et al., 2012, Commenges and Gégout-Petit, 2009], and if the causal stochastic system is completely observed, a test of conditional local independence is a test of no *direct* causal effect. See also Røysland et al. [2022], who use local independence graphs to formulate criteria for identification of causal effects in continuous-time survival models. If the causal stochastic system is only partially observed, a conditional local dependency need not correspond to a direct causal effect due to unobserved confounding, but the projected local independence graph, as introduced by Mogensen and Hansen [2020], retains a causal interpretation, and its Markov equivalence class can be learned by conditional local independence testing. In addition, within the framework of structural nested models, testing the hypothesis of no *total* causal effect can also be cast as a test of conditional local independence [Lok, 2008, Thm. 9.2].

To appreciate what conditional local independence means – and, in particular, what it does not mean – it is useful to compare with ordinary conditional independence. In our example, N_t is conditionally locally independent of X_t given $\mathcal{F}_t^{N,Z}$, but this implies neither that $N \perp\!\!\!\perp X \mid Z$ (as processes), nor that $N_t \perp\!\!\!\perp X_t \mid \mathcal{F}_t^Z$. In fact, these conditional independencies cannot hold in this example where $X_t = X_T$ for $t \geq T$ – except in special cases such as T being a deterministic function of Z . Theorem 2 in Didelez [2008] gives a sufficient condition for $N_t \perp\!\!\!\perp X_t \mid \mathcal{F}_t^Z$ to hold in terms of the local independence graph, but this condition is also not fulfilled by the graph in Figure 2.1.1 due to the edge from N to X . Moreover, conditional local independence is in general also different from the baseline conditional independence $T \perp\!\!\!\perp X_0 \mid Z_0$ unless both X and Z are time-independent, see Section 2.3.2. In Section 2.E in the supplementary material [Christgau et al., 2023c], we elaborate further upon the connection to semiparametric survival models. Didelez [2008] argues that N_t being conditionally locally independent of X_t given $\mathcal{F}_t^{N,Z}$ heuristically means that $N_t \perp\!\!\!\perp \mathcal{F}_{t-}^X \mid \mathcal{F}_{t-}^{N,Z}$, but this is technically problematic in continuous time. If T has a continuous distribution, then for any fixed t , $N_t = N_{t-}$ almost surely, whence N_t is almost surely $\mathcal{F}_{t-}^{N,Z}$ -measurable and conditionally

2 Nonparametric conditional local independence testing

independent of anything given $\mathcal{F}_{t-}^{N,Z}$. This heuristic can thus not be used to formally define conditional local independence in continuous time. See instead the formal Definition 2 by Didelez [2008] or our Definition 2.2.1.

Several examples from health sciences given by Didelez [2008] demonstrate the usefulness of conditional local independence for multivariate event systems, and more recent attention to event systems in the machine learning community [Zhou et al., 2013, Xu et al., 2016, Achab et al., 2017, Bacry et al., 2018, Cai et al., 2022] testifies to the relevance of conditional local independence. This line of research relies primarily on the linear Hawkes process model, which is effectively used to infer local independence graphs – sometimes even interpreted causally. The Hawkes model is attractive because conditional local independencies can be inferred from corresponding kernel functions being zero – and statistical tests can readily be based on parametric or nonparametric estimation of kernels.

A less attractive property of the Hawkes model is that it is not closed under marginalization. As with any model based statistical test, the validity of the test is jeopardized by model misspecification, hence even within a subsystem of a linear Hawkes process, a test of conditional local independence based on a Hawkes model may be invalid.

The challenges resulting from model misspecification and marginalization is investigated further in Sections 2.2.2 and 2.6 based on an extension of our introductory example and Cox’s survival model. Both the Hawkes model and the Cox model illustrate that conditional local independence might be expressed and tested within a (semi-)parametric model, but model misspecification makes us question the validity of any such model based test. Thus there is a need for a nonparametric test of the hypothesis of conditional local independence. Moreover, since we cannot translate the hypothesis into an equivalent hypothesis about classical conditional independence, we cannot directly use existing nonparametric tests, such as the GCM [Shah and Peters, 2020] or the GHCM [Lundborg et al., 2022b], of conditional independence.

We propose a new nonparametric test when the target process N is a counting process and X is a real valued process, and where the hypothesis is that N is conditionally locally independent of X given a filtration \mathcal{F}_t . In the context of the introductory example, $\mathcal{F}_t = \mathcal{F}_t^{N,Z}$. We consider a counting process target primarily because the theory of conditional local independence is most complete in this case, but generalizations are possible – we refer to the discussion in Section 2.7. Within our framework we base our test on an infinite dimensional parameter, which we call the Local Covariance Measure (LCM). It is a function of time, which is constantly equal to zero under the hypothesis. Our main result is that the LCM can be estimated by using the ideas of double machine learning [Chernozhukov et al., 2018] in such a way that the estimator converges uniformly at a \sqrt{n} -rate to a mean zero Gaussian martingale under the hypothesis of conditional local independence. We use the LCM to develop the (cross-fitted) Local Covariance Test ((X)-LCT), for which we derive uniform level and power results.

2.1.1 Organization of the paper

In Section 2.2 we introduce the general framework for formulating the hypothesis of conditional local independence. This includes the introduction in Section 2.2.1 of an abstract residual process, which is used to define the LCM as a functional target parameter indexed by time. The LCM equals the zero-function under the hypothesis of conditional local independence, and to test this hypothesis we introduce an estimator of the LCM in Section 2.2.3. The estimator is a stochastic process, and we describe how sample splitting is to be used for its computation via the estimation of two unknown components.

In Section 2.3 we give interpretations of the LCM and its estimator. We show that the LCM estimator is a Neyman orthogonalized score statistic in Section 2.3.1, and in Section 2.3.2 we relate LCM to the partial copula when X is time-independent.

In Section 2.4 we state the main results of the paper. We establish in Section 2.4.1 that the LCM estimator generally approximates the LCM with an error of order $n^{-1/2}$. Under the hypothesis of conditional local independence, we show that the (scaled) LCM estimator converges weakly to a mean zero Gaussian martingale. The estimator requires a model of the target process N as well as the process X conditionally on \mathcal{F}_t to achieve the orthogonalization at the core of double machine learning. The model of X is in this paper expressed indirectly in terms of the residual process, and we show that if we can learn the residual process at rate $g(n)$ and the model of N at rate $h(n)$ such that $g(n), h(n) \rightarrow 0$ and $\sqrt{ng(n)h(n)} \rightarrow 0$ for $n \rightarrow \infty$ then we achieve a \sqrt{n} -rate convergence of the LCM estimator. We also show that the variance function of the Gaussian martingale can be estimated consistently, and we give a general result on the asymptotic distribution of univariate test statistics based on the LCM estimator. All asymptotic results are presented in the framework of uniform stochastic convergence.

Section 2.5 gives explicit examples of univariate test statistics, including the Local Covariance Test based on the normalized supremum of the LCM estimator. Its asymptotic distribution is derived and we present results on uniform asymptotic level and power. In Section 2.5.2 we present the generalization from the sample split estimator to the cross-fit estimator. Though this estimator and the corresponding cross-fit Local Covariance Test (X-LCT) are a bit more involved to compute and analyze, X-LCT is more powerful and thus our recommended test for practical usage.

The survival example from the introduction is used and elaborated upon throughout the paper. We introduce a Cox model in terms of the time-varying covariate processes, and we report in Section 2.6 the results from a simulation study based on this model.

The paper is concluded by a discussion in Section 2.7.

The supplementary material [Christgau et al., 2023c], henceforth referred to as the supplement, consists of Sections 2.A through 2.G and contains: proofs of results in this paper (2.A); definitions and results on uniform asymptotics (2.B); a uniform version of Rebolledo’s martingale CLT (2.C); an overview of achievable rate results for estimation of nuisance parameters that enter into the LCM estimator (2.D); a comparison with semiparametric survival models (2.E); details on Neyman orthogonality (2.F); and additional results from the simulation study (2.G).

2.2 The Local Covariance Measure

In this section we present the general framework of the paper, we define conditional local independence and we introduce the Local Covariance Measure as a means to quantify deviations from conditional local independence. In Section 2.2.3 we outline how the Local Covariance Measure can be estimated using double machine learning and sample splitting. We illustrate the central concepts and methods by an example based on Cox's survival model with time-varying covariates.

We consider a counting process $N = (N_t)$ and another real value process $X = (X_t)$, both defined on the probability space $(\Omega, \mathbb{F}, \mathbb{P})$. All processes are assumed to be defined on a common compact time interval. We assume, without loss of generality, that the time interval is $[0, 1]$. We will assume that N is adapted w.r.t. a right continuous and complete filtration \mathcal{F}_t , and we denote by \mathcal{G}_t the right continuous and complete filtration generated by \mathcal{F}_t and X_t . We assume throughout that X is càglàd (that is, has sample paths that are continuous from the left and with limits from the right), which will ensure bounded sample paths and that the process is \mathcal{G}_t -predictable.

In the survival example of the introduction, $N_t = \mathbb{1}(T \leq t)$ is the indicator of whether death has happened by time t , and there can only be one event per individual observed. Furthermore, $\mathcal{F}_t = \mathcal{F}_t^{N,Z}$ and $\mathcal{G}_t = \mathcal{F}_t^{N,X,Z}$. Our general setup works for any counting process, thus it allows for recurrent events and censoring, and the filtration \mathcal{F}_t can contain the histories of any number of processes in addition to the history of N itself.

2.2.1 The hypothesis of conditional local independence

The counting process N is assumed to have an \mathcal{F}_t -intensity λ_t , that is, λ_t is \mathcal{F}_t -predictable and with

$$\Lambda_t = \int_0^t \lambda_s ds$$

being the compensator of N ,

$$M_t = N_t - \Lambda_t \tag{2.2.1}$$

is a local \mathcal{F}_t -martingale. Within this framework we can define the hypothesis of conditional local independence precisely.

Definition 2.2.1 (Conditional local independence). We say that N_t is conditionally locally independent of X_t given \mathcal{F}_t if the local \mathcal{F}_t -martingale M_t defined by (2.2.1) is also a local \mathcal{G}_t -martingale. ♣

For simplicity, we may also refer to this hypothesis as *local independence* and write

$$H_0 : M_t = N_t - \Lambda_t \text{ is a local } \mathcal{G}_t\text{-martingale.} \tag{2.2.2}$$

As argued in the introduction, the hypothesis of local independence is the hypothesis that observing X on $[0, t]$ does not add any information to \mathcal{F}_{t-} about whether an N -event will happen in an infinitesimal time interval $[t, t+dt)$. Definition 2.2.1 captures this

interpretation by requiring that the \mathcal{F}_t -compensator, Λ , of N is also the \mathcal{G}_t -compensator. Thus, λ is also the \mathcal{G}_t -intensity under H_0 .

If N has \mathcal{G}_t -intensity λ , the innovation theorem, Theorem II.T14 in Brémaud [1981], gives that the predictable projection $\lambda_t = \mathbb{E}(\lambda_t | \mathcal{F}_{t-})$ is the (predictable) \mathcal{F}_t -intensity. Local independence follows if λ is \mathcal{F}_t -predictable. Intensities are, however, only unique almost surely, and we can have local independence even if λ is not *a priori* \mathcal{F}_t -predictable but have an \mathcal{F}_t -predictable version. When N has \mathcal{G}_t -intensity λ , H_0 is thus equivalent to λ having an \mathcal{F}_t -predictable version. We find Definition 2.2.1 preferable because it directly gives an operational criterion for determining whether N has an \mathcal{F}_t -predictable version of a \mathcal{G}_t -intensity.

Remark 2.2.2 (Censoring). Suppose that the data is censored such that $(N_t, X_t, \mathcal{F}_t) = (N_{t \wedge C}^*, X_{t \wedge C}^*, \mathcal{F}_{t \wedge C}^*)$, where $(N^*, X^*, \mathcal{F}^*)$ are uncensored data and where C is the censoring time. The hypothesis regarding the uncensored data,

$$H_0^*: N_t^* \text{ is locally independent of } X_t^* \text{ given } \mathcal{F}_t^*,$$

might then be the hypothesis of interest. If $\mathbb{1}(C \geq t)$ happens to be \mathcal{F}_t^* -predictable, it is straightforward to show that H_0^* implies H_0 , and consequently a test of H_0 is also a test of H_0^* . However, \mathcal{F}_t^* may not *a priori* contain information about the censoring process. Suppose instead that the common condition of *independent censoring* [Andersen et al., 1993] holds, which is equivalent to N_t^* being locally independent of $\mathcal{C}_t := \mathbb{1}(C \leq t)$ given \mathcal{G}_t^* [Røysland et al., 2022]. Then H_0^* implies that N_t is locally independent of X_t given $\mathcal{F}_t \vee \mathcal{F}_t^c$. Thus, in order to test H_0^* , we replace \mathcal{F}_t by the enlarged filtration $\mathcal{F}_t \vee \mathcal{F}_t^c$ and proceed *mutatis mutandis* with testing H_0 using the observed data. \diamond

Since X is assumed càglàd, and thus especially \mathcal{G}_t -predictable, the stochastic integral

$$\int_0^t X_s dM_s, \quad (2.2.3)$$

is under H_0 a local \mathcal{G}_t -martingale. A test could be based on detecting whether (2.2.3) is, indeed, a local martingale. We will take a slightly different approach where we replace the integrand X by a residual process as defined below. We do so for two reasons. First, to achieve a \sqrt{n} -rate via double machine learning we need the integrand to fulfill (2.2.4) below. Second, other choices of integrands than X could potentially lead to more powerful tests.

Definition 2.2.3 (Residual Process). A residual process $G = (G_t)_{t \in [0,1]}$ of X_t given \mathcal{F}_t is a càglàd stochastic process that is \mathcal{G}_t -adapted and satisfies

$$\mathbb{E}(G_t | \mathcal{F}_{t-}) = 0 \quad (2.2.4)$$

for $t \in [0, 1]$. ♣

The geometric interpretation is that the residual process evolves such that G_t is orthogonal to $L_2(\mathcal{F}_{t-})$ within $L_2(\mathcal{G}_{t-})$ at each time t . One obvious residual process is the *additive residual process* given by

$$G_t = X_t - \Pi_t = X_t - \mathbb{E}(X_t | \mathcal{F}_{t-}),$$

2 Nonparametric conditional local independence testing

where $\Pi_t = \mathbb{E}(X_t \mid \mathcal{F}_{t-})$ denotes the predictable projection of the càglàd process X_t , see Theorem VI.19.6 in Rogers and Williams [2000]. The additive residual projects X_t onto the orthogonal complement of $L_2(\mathcal{F}_{t-})$, but this may not necessarily remove all \mathcal{F}_t -predictable information from X_t . An alternative choice that does so under sufficient regularity conditions is the *quantile residual process* given by

$$G_t = F_t(X_t) - \frac{1}{2},$$

where F_t is the conditional distribution function given by $F_t(x) = \mathbb{P}(X_t \leq x \mid \mathcal{F}_{t-})$. The quantile residual process satisfies (2.2.4) provided that $(t, x) \mapsto F_t(x)$ is continuous. In Section 2.3.1 we discuss additional transformations of X that can also be applied before any residualization procedure.

We will formulate the general results in terms of an abstract residual process, but we focus on the additive residual process in the examples. Any non-degenerate residual process will contain a predictive model of (aspects of) X_t given \mathcal{F}_{t-} in order to satisfy (2.2.4). We use \hat{G}_t to denote the residual obtained by plugging in an estimate of that predictive model. For the additive residual process, the predictive model is Π_t and $\hat{G}_t = X_t - \hat{\Pi}_t$. For the quantile residual process, the predictive model is F_t and $\hat{G}_t = \hat{F}_t(X_t) - \frac{1}{2}$.

We can now define our functional target parameter of interest, which we call the *Local Covariance Measure*.

Definition 2.2.4 (Local Covariance Measure). With G_t a residual process, define for $t \in [0, 1]$

$$\gamma_t = \mathbb{E}(I_t), \quad \text{where} \quad I_t = \int_0^t G_s dM_s, \quad (2.2.5)$$

whenever the expectation is well defined. We call the function $t \mapsto \gamma_t$ the Local Covariance Measure (LCM). ♣

The following propositions illuminate how γ relates to the null hypothesis of N_t being conditionally locally independent of X_t given \mathcal{F}_t .

Proposition 2.2.5. *Under H_0 , the process $I = (I_t)$ is a local \mathcal{G}_t -martingale with $I_0 = 0$. If I is a martingale, then $\gamma_t = 0$ for $t \in [0, 1]$.*

To interpret γ in the alternative, we assume that N has \mathcal{G}_t -intensity λ .

Proposition 2.2.6. *If $\int_0^1 \mathbb{E}(|G_s|(\lambda_s + \lambda_s)) ds < \infty$, then for every $t \in [0, 1]$,*

$$\gamma_t = \int_0^t \text{Cov}(G_s, \lambda_s - \lambda_s) ds.$$

In particular, γ is the zero-function if and only if $\text{Cov}(G_s, \lambda_s - \lambda_s) = 0$ for almost all $s \in [0, 1]$.

We note that under H_0 , the condition $\int_0^1 \mathbb{E}(|G_s| \lambda_s) ds < \infty$ is sufficient to ensure that I is a martingale and $\gamma_t = 0$ for all $t \in [0, 1]$. By Proposition 2.2.6, the LCM quantifies deviations from H_0 in terms of the covariance between the residual process and the difference of the \mathcal{F}_t - and \mathcal{G}_t -intensities. To this end, note that if X happens to be \mathcal{F}_t -adapted, then $\mathcal{G}_t = \mathcal{F}_t$ and N is trivially locally independent of X . The hypothesis of local independence is only of interest when \mathcal{G}_t is a strictly larger filtration than \mathcal{F}_t , that is, when X provides information not already in \mathcal{F}_t .

For the additive residual process, where $G_t = X_t - \Pi_t$,

$$\gamma_t = \mathbb{E} \left(\int_0^t G_s dM_s \right) = \mathbb{E} \left(\int_0^t X_s dM_s \right) - \mathbb{E} \left(\int_0^t \Pi_s dM_s \right)$$

provided that the expectations are well defined. Since the predictable projection Π_t has a càglàd version and is \mathcal{F}_t -predictable, and since M_t is a local \mathcal{F}_t -martingale, $\int_0^t \Pi_s dM_s$ is a local \mathcal{F}_t -martingale. If it is a martingale, it is a mean zero martingale, and

$$\gamma_t = \mathbb{E} \left(\int_0^t X_s dM_s \right) = \mathbb{E} \left(\sum_{\tau \leq t: \Delta N_\tau = 1} X_\tau - \int_0^t X_s \lambda_s ds \right). \quad (2.2.6)$$

The computation above shows that the additive residual process defines the same functional target parameter γ_t as the stochastic integral (2.2.3) would. It is, however, the representation of γ_t as the expectation of the residualized stochastic integral that will allow us to achieve a \sqrt{n} -rate of convergence of the estimator of γ_t in cases where the estimator of λ_t converges at a slower rate.

2.2.2 A Cox model with a partially observed covariate process

To further illustrate the hypothesis of conditional local independence and the Local Covariance Measure we consider an example based on Cox's survival model with time dependent covariates. This is an extension of the example from the introduction with T being the time to death of an individual, and with X and Z being time-varying processes. There is, moreover, one additional time-varying process Y in the full model.

An interpretation of the processes is as follows:

$$\begin{aligned} X &= \text{Pension savings} \\ Y &= \text{Blood pressure} \\ Z &= \text{BMI} \end{aligned}$$

Periods of overweight or obesity may influence blood pressure in the long term, and due to, e.g., job market discrimination, high BMI could influence pension savings negatively. Death risk is influenced directly by BMI and blood pressure but not the size of your pension savings. Figure 2.2.2 illustrates two possible dependence structures among the three processes and the death time as local independence graphs, and we will use these two graphs to discuss the concept of conditional local independence of pension savings on time to death.

2 Nonparametric conditional local independence testing

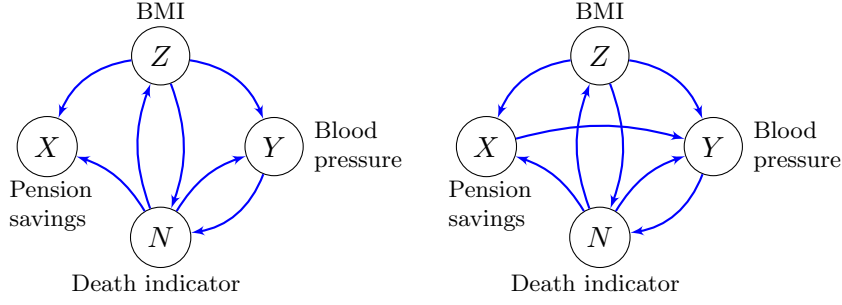


Figure 2.2.2: Local independence graphs illustrating how the three processes X , Y , and Z could affect each other and time of death in the Cox example. There is no direct influence of X (pension savings) on time of death in either of the two graphs, but in the left graph the death indicator is furthermore conditionally locally independent of X given the history of Z and N . In the right graph, Z and N does not block all paths from X to N , thus conditioning on the history of Z and N only would not render N conditionally locally independent of X .

We assume that $T \in [0, 1]$ and that X , Y and Z have continuous sample paths. Recall also that $N_t = \mathbb{1}(T \leq t)$ is the death indicator process. To maintain some form of realism, all processes are stopped at time of death, that is, $X_t = X_T$, $Y_t = Y_T$ and $Z_t = Z_T$ for $t \geq T$. This feedback from the death event to the other processes is reflected in Figure 2.2.2 by the edges pointing out of N . Recall also that

$$\mathcal{F}_t^{N,Z} = \sigma(N_s, Z_s; s \leq t)$$

is the filtration generated by the N - and Z -processes. We use a similar notation for other processes and combinations of processes. For example, $\mathcal{F}_t^{N,X,Y,Z}$ is the filtration generated by N and all three X -, Y -, and Z -processes. With λ_t^{full} denoting the $\mathcal{F}_t^{N,X,Y,Z}$ -intensity of time of death based on the history of all processes, we assume in this example a Cox model given by

$$\lambda_t^{\text{full}} = \mathbb{1}(T \geq t) \lambda_t^0 e^{Y_t + \beta Z_t} \quad (2.2.7)$$

with λ_t^0 a deterministic baseline intensity. It is not important that λ_t^{full} is a Cox model for our general theory, but it allows for certain theoretical computations in this example.

The fact that λ_t^{full} does not depend upon X_t implies that λ_t^{full} is also the $\mathcal{F}_t^{N,Y,Z}$ -intensity, and according to Definition 2.2.1, N_t is conditionally locally independent of X_t given $\mathcal{F}_t^{N,Y,Z}$. This is in agreement with the local independence graphs in Figure 2.2.2 where there is no edge in either of them from X to N .

We will take an interest in the case where Y is unobserved and test the hypothesis:

$$H_0 : N_t \text{ is conditionally locally independent of } X_t \text{ given } \mathcal{F}_t^{N,Z}.$$

That is, with Y unobserved we want test if the intensity of time to death given the history of N , X and Z depends on X . To simplify notation let $\mathcal{F}_t = \mathcal{F}_t^{N,Z}$ and $\mathcal{G}_t = \mathcal{F}_t^{N,X,Z}$

in accordance with the general notation. The \mathcal{G}_t -intensity is by the innovation theorem given as

$$\lambda_t = \mathbb{E}(\lambda_t^{\text{full}} \mid \mathcal{G}_{t-}) = \mathbb{1}(T \geq t) \lambda_t^0 e^{\beta Z_t} \mathbb{E}(e^{Y_t} \mid \mathcal{G}_{t-}), \quad (2.2.8)$$

while the \mathcal{F}_t -intensity is

$$\lambda_t = \mathbb{E}(\lambda_t^{\text{full}} \mid \mathcal{F}_{t-}) = \mathbb{1}(T \geq t) \lambda_t^0 e^{\beta Z_t} \mathbb{E}(e^{Y_t} \mid \mathcal{F}_{t-}), \quad (2.2.9)$$

and H_0 is equivalent to $\lambda_t = \boldsymbol{\lambda}_t$ almost surely. Comparing (2.2.8) and (2.2.9) we see that H_0 holds in this example if $\mathbb{E}(e^{Y_t} \mid \mathcal{G}_{t-}) = \mathbb{E}(e^{Y_t} \mid \mathcal{F}_{t-})$, and a sufficient condition for this is

$$\mathcal{F}_t^X \perp\!\!\!\perp \mathcal{F}_t^Y \mid \mathcal{F}_t. \quad (2.2.10)$$

The condition (2.2.10) is in concordance with the left graph in Figure 2.2.2, see Theorem 2 in Didelez [2008], but not the right, and it implies H_0 . We will in Section 2.6.1 elaborate on condition (2.2.10) and give explicit examples.

We recall that H_0 can be reformulated as $\boldsymbol{\lambda}_t$ not depending on X , and we could investigate the hypothesis via a marginal Cox model

$$\boldsymbol{\lambda}_t^{\text{cox}} = \mathbb{1}(T \geq t) \boldsymbol{\lambda}_t^0 e^{\alpha_1 X_t + \alpha_2 Z_t} \quad (2.2.11)$$

and test if $\alpha_1 = 0$. The Cox model is, however, not closed under marginalization and the semi-parametric model (2.2.11) is quite likely misspecified. Consequently, a test of $\alpha_1 = 0$ is not equivalent to a test of H_0 .

Our proposed nonparametric test of H_0 does not rely on a specific (semi-)parametric model of $\boldsymbol{\lambda}_t$. To test H_0 we consider the LCM using the additive residual process. Then (2.2.6) implies that

$$\gamma_t = \mathbb{E} \left(X_T N_t - \int_0^t X_s \lambda_s ds \right),$$

By Proposition 2.2.5, $\gamma_t = 0$ for $t \in [0, 1]$ under H_0 , whence conditional local independence implies $\gamma_t = 0$, and we test H_0 by estimating γ_t and testing if it is constantly equal to 0.

Before introducing a general estimator of the LCM in Section 2.2.3 we outline how to estimate the end point parameter γ_1 in this example. Due to $T \leq 1$ and the appearance of the indicator $\mathbb{1}(T \geq t)$ in (2.2.9),

$$\gamma_1 = \mathbb{E} \left(X_T - \int_0^T X_s \lambda_s ds \right).$$

With i.i.d. observations $(T_1, X_1, Z_1), \dots, (T_n, X_n, Z_n)$ and (nonparametric) estimates, $\hat{\lambda}_{j,t}$, based on $(T_1, Z_1), \dots, (T_n, Z_n)$, we could compute the plug-in estimate

$$\hat{\gamma}_{1, \text{plug-in}}^{(n)} = \frac{1}{n} \sum_{j=1}^n \left(X_{j, T_j} - \int_0^{T_j} X_{j,s} \hat{\lambda}_{j,s} ds \right).$$

2 Nonparametric conditional local independence testing

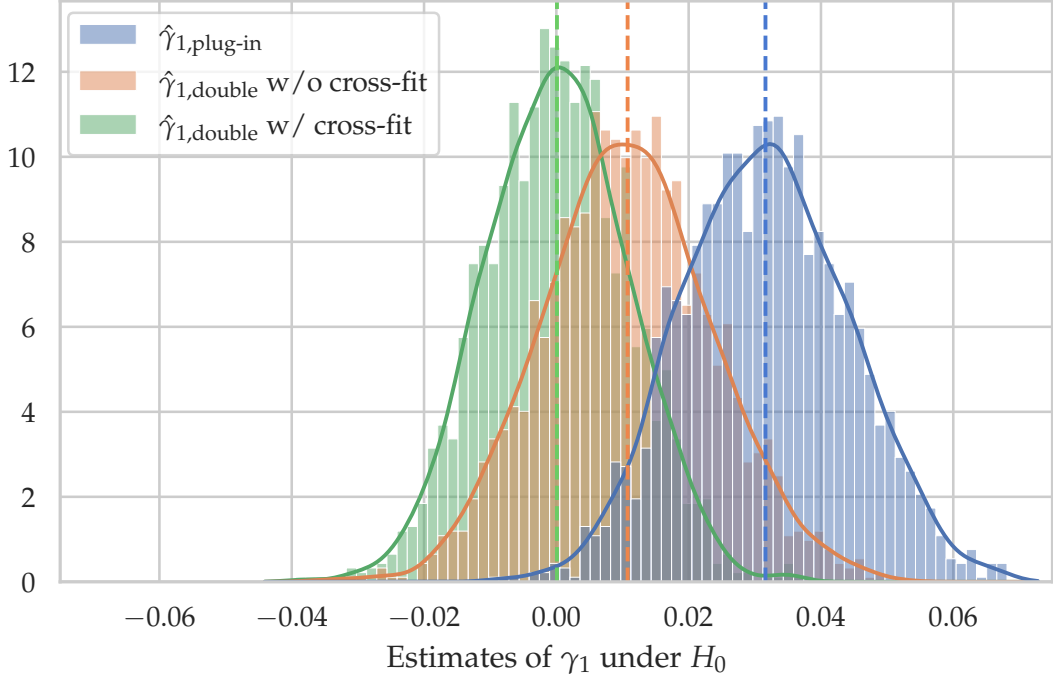


Figure 2.2.3: Histograms of the distributions of three different estimators of γ_1 . Each histogram contains 1000 estimates fitted to samples of size $n = 500$. The samples were sampled from a model that satisfies the hypothesis of conditional local independence and hence the ground truth is $\gamma_1 = 0$. See Section 2.6.2 for further details of the data generating process.

However, we cannot expect the plug-in estimator to have a \sqrt{n} -rate unless $\hat{\lambda}$ has \sqrt{n} -rate, which effectively requires parametric model assumptions on the intensity. Using the definition of γ_1 in terms of the additive residual process $G_t = X_t - \Pi_t$, we also have that

$$\gamma_1 = \mathbb{E} \left(X_T - \Pi_T - \int_0^T (X_s - \Pi_s) \lambda_s ds \right). \quad (2.2.12)$$

A double machine learning estimator based on the ideas by Chernozhukov et al. [2018] is therefore obtained by plugging in two nonparametric estimators:

$$\hat{\gamma}_{1,\text{double}}^{(n)} = \frac{1}{n} \sum_{j=1}^n \left(X_{j,T_j} - \hat{\Pi}_{j,T_j} - \int_0^{T_j} (X_{j,s} - \hat{\Pi}_{j,s}) \hat{\lambda}_{j,s} ds \right).$$

To achieve a small bias and a \sqrt{n} -rate of convergence, we use sample splitting. The nonparametric estimates $\hat{\Pi}_j$ and $\hat{\lambda}_j$ are based on one part of the sample only, and are thus independent of the other part of the sample used for testing, see Section 2.2.3. To obtain a fully efficient estimator, multiple sample splits can be combined, e.g., via cross-fitting, see Section 2.5.2.

Figure 2.2.3 shows the distributions of $\hat{\gamma}_{1,\text{plug-in}}^{(500)}$ and $\hat{\gamma}_{1,\text{double}}^{(500)}$ for the Cox example with $\gamma_1 = 0$, see Section 2.6.2 for details on the full model specification. The latter estimator was computed using cross-fitting but also without using any form of sample splitting. The figure illustrates the bias of $\hat{\gamma}_{1,\text{plug-in}}^{(500)}$, which is somewhat diminished by double machine learning without sample splitting and mostly eliminated by double machine learning in combination with cross-fitting.

2.2.3 Estimating the Local Covariance Measure

To estimate the LCM we assume that we have observed n i.i.d. replications of the processes, $(N_1, X_1, \mathcal{F}_1), \dots, (N_n, X_n, \mathcal{F}_n)$, where observing $\mathcal{F}_j = (\mathcal{F}_{j,t})$ signifies that anything adapted to the j -th filtration is computable from observations. The process N_j is adapted to \mathcal{F}_j , while X_j is not, and \mathcal{G}_j denotes the smallest right continuous and complete filtration generated by X_j and \mathcal{F}_j .

For each n , we consider a sample split corresponding to a partition $J_n \cup J_n^c = \{1, \dots, n\}$ of the indices into two disjoint sets. We let $\hat{\lambda}^{(n)}$ and $\hat{G}^{(n)}$ be estimates of the intensity and the residualization map, respectively, fitted on data indexed by J_n^c only. By an estimate, $\hat{\lambda}^{(n)}$, of λ we mean a (stochastic) function that can be evaluated on the basis of $\mathcal{F}_{j,t}$ for $j \in J_n$, and its value, denoted by $\hat{\lambda}_{j,t}^{(n)}$, is interpreted as a prediction of $\lambda_{j,t}$. The stochasticity in $\hat{\lambda}^{(n)}$ arises from its dependence on data indexed by J_n^c , from which its functional form is completely determined. Similarly, $\hat{G}^{(n)}$ is a function that can be evaluated on the basis of $\mathcal{G}_{j,t}$ for $j \in J_n$ to give a prediction $\hat{G}_{j,t}^{(n)}$ of $G_{j,t}$. In Section 2.6.1 we illustrate through the Cox example how $\hat{\lambda}^{(n)}$ and $\hat{G}^{(n)}$ are to be computed in practice when we use sample splitting. In Section 2.D in the supplement we give more examples of such estimation procedures and discuss their statistical properties in greater detail.

To ease notation, we will throughout assume that (N, X, \mathcal{F}) denotes one additional process and filtration – independent of and with the same distribution as the observed processes. Then the estimated intensity $\hat{\lambda}^{(n)}$ and estimated residual process $\hat{G}^{(n)}$ can be evaluated on (N, X, \mathcal{F}) , and thus we may write $\hat{\lambda}_t^{(n)}$ and $\hat{G}_t^{(n)}$ to denote template copies of $\hat{\lambda}_{j,t}^{(n)}$ and $\hat{G}_{j,t}^{(n)}$ for $j \in J_n$.

In terms of the estimates $\hat{\lambda}^{(n)}$ and $\hat{G}^{(n)}$ we estimate LCM by the stochastic process $\hat{\gamma}^{(n)}$ given by

$$\hat{\gamma}_t^{(n)} = \frac{1}{|J_n|} \sum_{j \in J_n} \int_0^t \hat{G}_{j,s}^{(n)} d\widehat{M}_{j,s}^{(n)}, \quad (2.2.13)$$

where $\widehat{M}_{j,t}^{(n)} = N_{j,t} - \int_0^t \hat{\lambda}_{j,s}^{(n)} ds$. We can regard $\hat{\gamma}_t^{(n)}$ as a double machine learning estimator of γ_t , with the observations indexed by J_n^c used to learn models of λ and G , and with observations indexed by J_n used to estimate γ_t based on these models. In Section 2.5.2 we define the more efficient estimator that uses cross-fitting, but it is instructive to study the simpler estimator based on sample splitting first.

In practical applications, we do not directly observe the filtration \mathcal{F}_j , but rather samples from the stochastic processes generating the filtration. In accordance with the introductory Cox example, consider \mathcal{F}_j and \mathcal{G}_j given by $\mathcal{F}_{j,t} = \sigma(Z_{j,s}, N_{j,s}; s \leq t)$ and

2 Nonparametric conditional local independence testing

$\mathcal{G}_{j,t} = \sigma(X_{j,s}, Z_{j,s}, N_{j,s}; s \leq t)$ for a third stochastic process Z_j , with Z_j possibly being multivariate. Within this setup, a general procedure for numerically computing the LCM is described in Algorithm 1. Here, historical regression refers to any method which regresses the outcome at a given time on the history of the regressors up to that time. For example, historical linear regression is discussed in Section 2.6 and various alternative methods are discussed in Section 2.D in the supplement. The choice of sample split will be discussed further in Section 2.5.2 in the context of cross-fitting.

Algorithm 1: Sample split estimator of LCM

- 1 **input:** processes $(N_j, X_j, Z_j)_{j=1, \dots, n}$, partition $J_n \cup J_n^c$ of indices ;
- 2 **options:** historical regression methods for estimation of λ and G given N and Z ,
- 3 discrete time grid $0 = t_0 < \dots < t_k \leq 1$;
- 4 **begin**
- 5 historically regress $(X_j)_{j \in J_n^c}$ on $(N_j, Z_j)_{j \in J_n^c}$ to obtain a fitted model $\widehat{G}^{(n)}$;
- 6 historically regress $(N_j)_{j \in J_n^c}$ on $(N_j, Z_j)_{j \in J_n^c}$ to obtain a fitted model $\widehat{\lambda}^{(n)}$;
- 7 compute out of sample residuals $\widehat{G}_{j,t_i}^{(n)}$ and $\widehat{M}_{j,t_i}^{(n)}$ for $j \in J_n$ and $i = 0, \dots, k$;
- 8 for each $i = 1, \dots, k$, compute

$$\widetilde{\gamma}_{t_i}^{(n)} = \frac{1}{|J_n|} \sum_{j \in J_n} \sum_{1 \leq l \leq i} \widehat{G}_{j,t_l}^{(n)} (\widehat{M}_{j,t_l}^{(n)} - \widehat{M}_{j,t_{l-1}}^{(n)})$$

- 9 **output:** Local Covariance Measure $\widetilde{\gamma}^{(n)}$ numerically approximated on grid;
-

As in Section 2.2.2 we could suggest estimating the entire function $t \mapsto \gamma_t$ by a simple plug-in estimator of λ using the representation (2.2.6). Figure 2.2.4 illustrates the distribution of estimators of the entire time dependent LCM for this plug-in estimator together with the double machine learning estimator with and without using cross-fitting. The figure also shows the distribution of the endpoint being the same distribution shown in Figure 2.2.3. The simulation is under H_0 , and we see that only the double machine learning estimator with cross-fitting results in estimated sample paths centered around 0.

2.3 Interpretations of the LCM estimator

In this section we provide some additional perspectives on and interpretations of the LCM. First we show that the LCM estimator can be seen as a Neyman orthogonalization of the score statistic for a particular one-parameter family. The abstract formulation of the residual process (G_t) permits that we transform X into another \mathcal{G}_t -predictable processes. Using this perspective, we may optimize the choice of the process X in terms of power.

Next we show that when X is independent of time, the test statistic reduces in a

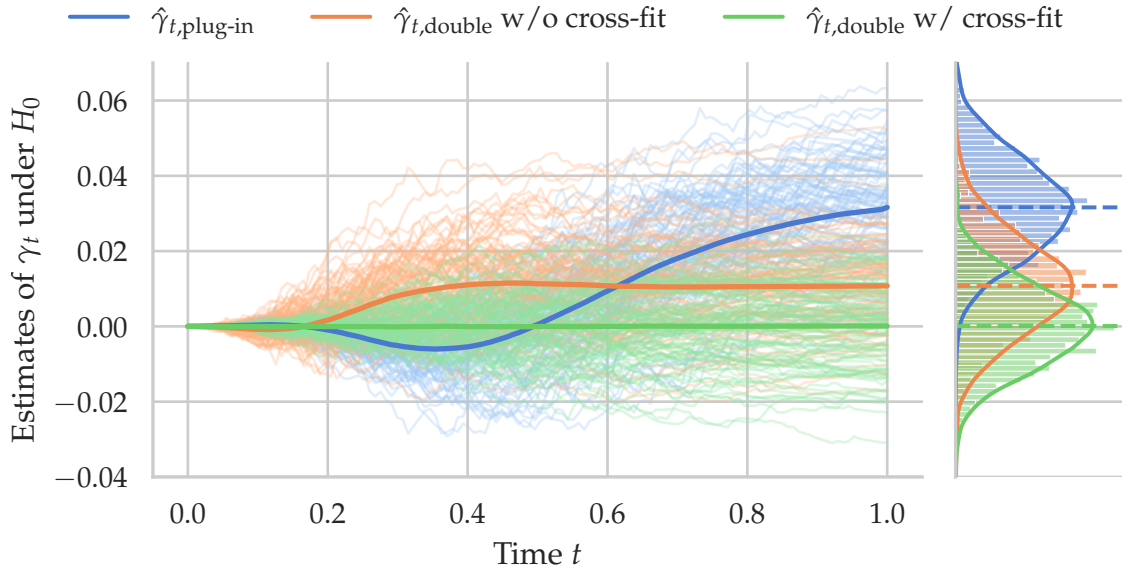


Figure 2.2.4: A time dependent extension of Figure 2.2.3 showing the distribution of the sample paths $t \mapsto \hat{\gamma}_{t,\text{plug-in}}^{(500)}$ and $t \mapsto \hat{\gamma}_{t,\text{double}}^{(500)}$, the latter with and without using cross-fitting. The data were simulated under H_0 where $t \mapsto \gamma_t$ is the zero function. See Section 2.6.2 for further details of the data generating process.

survival context to certain covariance measures between X -residuals and Cox-Snell-residuals, which we can link to existing test statistics for ordinary conditional independence.

2.3.1 Neyman orthogonalization of a score statistic

Consider the one-parameter family of \mathcal{G}_t -intensities

$$\lambda_t^\beta = e^{\beta X_t} \lambda_t$$

for $\beta \in \mathbb{R}$. Within this one-parameter family, the hypothesis of conditional local independence is equivalent to $H_0 : \beta = 0$. The normalized log-likelihood with n i.i.d. observations in the interval $[0, t]$ is

$$\begin{aligned} \ell_t(\beta) &= \frac{1}{n} \sum_{j=1}^n \left(\int_0^t \log(\lambda_{j,s}^\beta) dN_{j,s} - \int_0^t \lambda_{j,s}^\beta ds \right) \\ &= \frac{1}{n} \sum_{j=1}^n \left(\int_0^t \beta X_{j,s} + \log(\lambda_{j,s}) dN_{j,s} - \int_0^t e^{\beta X_{j,s}} \lambda_{j,s} ds \right). \end{aligned}$$

2 Nonparametric conditional local independence testing

Straightforward computations show that

$$\partial_\beta \ell_t(0) = \frac{1}{n} \sum_{j=1}^n \int_0^t X_{j,s} dM_{j,s} \quad \text{and} \quad -\partial_\beta^2 \ell_t(0) = \frac{1}{n} \sum_{j=1}^n \int_0^t X_{j,s}^2 \lambda_{j,s} ds.$$

If λ were known, the score statistic $\partial_\beta \ell_t(0)$ satisfies $\mathbb{E}(\partial_\beta \ell_t(0)) = \gamma_t$. Moreover, under $H_0 : \beta = 0$ we have that $-\partial_\beta^2 \ell_t(0) = \langle \partial_\beta \ell_t(0) \rangle$ is a consistent estimate of the asymptotic variance of the mean zero martingale $\partial_\beta \ell_t(0)$. The hypothesis of local independence – with λ known – could thus be tested using the score test statistic $-\partial_\beta \ell_t(0)^2 / \partial_\beta^2 \ell_t(0)$.

The nuisance parameter λ is, however, unknown and we want to avoid restrictive parametric assumptions about λ . Replacing $X_{j,t}$ by the residual process $G_{j,t}$ in the score statistic $\partial_\beta \ell_t(0)$ gives a *Neyman orthogonalized* score

$$\frac{1}{n} \sum_{j=1}^n \int_0^t G_{j,s} dM_{j,s}.$$

This score is linear in λ , which is used in supplementary Section 2.F to show that it satisfies the Neyman orthogonality condition under H_0 , cf. Definition 2.1 in Chernozhukov et al. [2018]. In this section, it is also shown that the act of replacing X_t with $G_t = X_t - \Pi_t$ can, in fact, be viewed as *concentrating out* the intensity of the score statistic in the sense of Newey [1994]. While Neyman orthogonality is never invoked explicitly, it is implicitly a central part of the asymptotic results for the LCM estimator (in particular Lemma 2.A.7 in the supplement).

The perspective on the LCM from a Neyman orthogonalized score statistic suggests that a test based on the LCM has most power against alternatives in the one-parameter family $\boldsymbol{\lambda}^\beta$. If it happens that the most important alternatives are of the form

$$\boldsymbol{\lambda}_t^\beta = e^{\beta \bar{X}_t} \boldsymbol{\lambda}_t$$

for some \mathcal{G}_t -predictable process \bar{X}_t different from X_t , then we should replace X_t by \bar{X}_t in our test statistic, that is, in the residualization procedure. Examples of processes \bar{X}_t are:

- transformations, $\bar{X}_t = f(X_t)$ for a function f
- time-shifts, $\bar{X}_t = X_{t-s}$ for $s > 0$
- linear filters, $\bar{X}_t = \int_0^t \kappa(t-s) X_s ds$ for a kernel κ
- non-linear filters, $\bar{X}_t = \varphi \left(\int_0^t \kappa(t-s) f(X_s) ds \right)$ for a kernel κ and functions f and φ .

Any finite number of such processes could, of course, also be combined into a vector process, and we could, indeed, generalize the LCM estimator (2.2.13) to a vector process. The generalization is straightforward.

2.3.2 Survival time with time-independent covariates

A different perspective on the test statistic is obtained if X is constant over time, if $N_t = \mathbb{1}(T \leq t)$ is the counting process of a survival time T , and if $\mathcal{F}_t = \sigma(N_s, Z; s \leq t)$ where Z is a vector of additional baseline variables. Then the \mathcal{F}_t -intensity is

$$\lambda_t = \mathbb{1}(T \geq t)h(t, Z),$$

where $h(t, Z)$ is the hazard function for T given the baseline Z . In this special case, the hypothesis of conditional local independence is equivalent to the ordinary conditional independence

$$X \perp\!\!\!\perp T \mid Z. \quad (2.3.14)$$

We also find that

$$\gamma_t = \mathbb{E}(X(\mathbb{1}(T \leq t) - \Lambda_{t \wedge T})),$$

and in particular $\gamma_1 = \mathbb{E}(X(1 - \Lambda_T))$ as $T \in [0, 1]$ by assumption. Since Λ_T is exponentially distributed with mean 1, we may write

$$\gamma_1 = -\text{cov}(X, \Lambda_T).$$

Testing if $\gamma_1 \neq 0$ in this particular setup is effectively a test of the conditional independence (2.3.14). When (2.3.14) is true, it further holds that $\Pi_t = \mathbb{E}(X \mid \mathcal{F}_{t-}) = \mathbb{E}(X \mid Z) = \Pi_0$ is independent of t , and if we incorporate this into our model of Π , the LCM estimator of γ_1 equals

$$\hat{\gamma}_1^{(n)} = \frac{1}{|J_n|} \sum_{j \in J_n} (X_j - \hat{\Pi}_{j,0})(1 - \hat{\Lambda}_{T_j}). \quad (2.3.15)$$

This is a (non-normalized) generalized covariance measure (GCM), see [Shah and Peters, 2020], which is simply the (negative) empirical covariance between the additive residuals $X_j - \hat{\Pi}_{j,0}$ and the Cox-Snell residuals $\hat{\Lambda}_{T_j}$.

Alternatively, consider the quantile residual process $G_{j,t} = F_t(X_j) - \frac{1}{2}$ where $F_t(x) = \mathbb{P}(X \leq x \mid \mathcal{F}_{t-})$. If (2.3.14) is true, it holds again that $F_t(x) = F_0(x) = \mathbb{P}(X \leq x \mid Z)$ is independent of t , and our LCM estimator becomes

$$\hat{\gamma}_1^{(n)} = \frac{1}{|J_n|} \sum_{j \in J_n} \hat{G}_{j,0} (1 - \hat{\Lambda}_{T_j}). \quad (2.3.16)$$

This is likewise an empirical covariance, but now between the generalized residuals and the Cox-Snell residuals. This is closely related to the partial copula between X and T given Z , which can be estimated as

$$\frac{1}{|J_n|} \sum_{j \in J_n} \hat{G}_{j,0} \left(\frac{1}{2} - \exp(-\hat{\Lambda}_{T_j}) \right).$$

See Petersen and Hansen [2021] for further details on the partial copula and how this statistic can be used to test conditional independence.

2 Nonparametric conditional local independence testing

Under a combined rate condition on estimation of G and Λ , the endpoint statistics above are known to be asymptotically Gaussian with mean zero when the hypothesis of conditional independence in (2.3.14) holds. Within this survival setting, the endpoint statistics (2.3.15) can, furthermore, be seen as a score test derivable from a semiparametric efficient score function. Section 2.E in the supplement gives the details for two specific semiparametric survival models.

Whenever $\widehat{G}_{j,t} = \widehat{G}_{j,0}$ is independent of time, e.g., if we incorporate (2.3.14) into the residual model, the t -indexed LCM estimator is

$$\widehat{\gamma}_t^{(n)} = \frac{1}{|J_n|} \sum_{j \in J_n} \widehat{G}_{j,0} \left(\mathbb{1}(T_j \leq t) - \widehat{\Lambda}_{T_j \wedge t} \right),$$

which can be seen as a t -indexed extension of (2.3.16). For a general, time-dependent residual process, the full t -indexed LCM estimator is

$$\widehat{\gamma}_t^{(n)} = \frac{1}{|J_n|} \sum_{j \in J_n} \mathbb{1}(T_j \leq t) \widehat{G}_{j,T_j} - \int_0^{t \wedge T_j} \widehat{G}_{j,s} \widehat{\lambda}_s ds.$$

The general results of this paper show that the t -indexed LCM estimator is asymptotically distributed as a mean zero Gaussian martingale under H_0 . This appears to be a novel result even when X is constant over time. However, the main contributions of this paper is to the case where X and Z are stochastic processes varying with time – where the hypothesis of conditional local independence is also distinct from (2.3.14).

2.4 General asymptotic results

In this section we derive uniform asymptotic results regarding the general LCM estimator as a stochastic process. In Section 2.5 we discuss how to construct tests of H_0 based on the asymptotic results.

We assume that N has a \mathcal{G}_t -intensity λ_t , we let $\mathbf{\Lambda}_t = \int_0^t \lambda_s ds$ denote the \mathcal{G}_t -compensator of N and let $\mathbf{M}_t = N_t - \mathbf{\Lambda}_t$ be the compensated local \mathcal{G}_t -martingale. We also recall that $\widehat{\gamma}^{(n)}$ denotes the LCM estimator based on sample splitting as defined in Section 2.2.3. Within this framework we consider the decomposition

$$\sqrt{|J_n|} \widehat{\gamma}^{(n)} = U^{(n)} + R_1^{(n)} + R_2^{(n)} + R_3^{(n)} + D_1^{(n)} + D_2^{(n)}, \quad (2.4.17)$$

where the processes $U^{(n)}$, $R_1^{(n)}$, $R_2^{(n)}$, $R_3^{(n)}$, $D_1^{(n)}$, and $D_2^{(n)}$ are given by

$$U_t^{(n)} = \frac{1}{\sqrt{|J_n|}} \sum_{j \in J_n} \int_0^t G_{j,s} dM_{j,s}, \quad (2.4.18)$$

$$R_{1,t}^{(n)} = \frac{1}{\sqrt{|J_n|}} \sum_{j \in J_n} \int_0^t G_{j,s} (\lambda_{j,s} - \hat{\lambda}_{j,s}^{(n)}) ds, \quad (2.4.19)$$

$$R_{2,t}^{(n)} = \frac{1}{\sqrt{|J_n|}} \sum_{j \in J_n} \int_0^t (\hat{G}_{j,s}^{(n)} - G_{j,s}) dM_{j,s}, \quad (2.4.20)$$

$$R_{3,t}^{(n)} = \frac{1}{\sqrt{|J_n|}} \sum_{j \in J_n} \int_0^t (\hat{G}_{j,s}^{(n)} - G_{j,s}) (\lambda_{j,s} - \hat{\lambda}_{j,s}^{(n)}) ds, \quad (2.4.21)$$

$$D_{1,t}^{(n)} = \frac{1}{\sqrt{|J_n|}} \sum_{j \in J_n} \int_0^t G_{j,s} (\lambda_{j,s} - \lambda_{j,s}) ds, \quad (2.4.22)$$

$$D_{2,t}^{(n)} = \frac{1}{\sqrt{|J_n|}} \sum_{j \in J_n} \int_0^t (\hat{G}_{j,s}^{(n)} - G_{j,s}) (\lambda_{j,s} - \lambda_{j,s}) ds. \quad (2.4.23)$$

Note that the processes D_1 and D_2 are (almost surely) the zero-process under H_0 , since the null is equivalent to λ_t being a version of $\boldsymbol{\lambda}_t$. We proceed to show that:

- the processes $U^{(n)}$ and $D_1^{(n)} - \sqrt{|J_n|}\gamma$ each converge in distribution,
- and the processes R_1, R_2, R_3, D_2 converge to the zero-process.

For the analysis of each of R_1 , R_2 , and D_2 , sample splitting is used to render the summands conditionally independent.

These asymptotic properties imply that $\sqrt{|J_n|}(\hat{\gamma}^{(n)} - \gamma)$ is stochastically bounded in general, so the LCM estimator will asymptotically detect if the LCM is non-zero. Moreover, it will follow that $U^{(n)}$ drives the asymptotic limit of the LCM estimator under H_0 . Based on these general asymptotic results we derive in Section 2.5 asymptotic error control for tests based on the LCM estimator.

2.4.1 Asymptotics of the LCM estimator

Our asymptotic results are formulated in terms of uniform stochastic convergence, which has also been discussed extensively in the recent literature on hypothesis testing [Shah and Peters, 2020, Lundborg et al., 2022b,a, Scheidegger et al., 2022, Neykov et al., 2021]. Uniform convergence allows us to establish uniform asymptotic level of our proposed test, as well as power under local alternatives. We have collected key definitions and results related to uniform convergence in Section 2.B in the supplement.

To state uniform assumptions and asymptotic results we need to indicate a range of possible sampling distributions for which the assumptions apply and the results hold. For this purpose, we extend our setup and allow all data to be parametrized by a fixed parameter set Θ . The set Θ is not *a priori* assumed to have any structure, and $\theta \in \Theta$

2 Nonparametric conditional local independence testing

simply indicates that N^θ , X^θ , λ^θ , G^θ etc. have θ -dependent distributions. We generally denote evaluation of processes or derived quantities for a specific θ -value by a superscript, with the LCM, γ^θ , in particular, depending on θ . The LCM estimator is likewise written as $\hat{\gamma}^{(n),\theta} = (\hat{\gamma}_t^{(n),\theta})$ for $\theta \in \Theta$ to denote its dependence on the sampling distribution. The superscript notation is, however, heavy and unnecessary in many cases and we will suppress the dependency on $\theta \in \Theta$ whenever it is not needed. Any result that does not explicitly involve Θ should be understood as a pointwise result for each $\theta \in \Theta$.

The parametrization allows us to express convergence in distribution and probability uniformly over Θ , which are denoted by $\xrightarrow{D/\Theta}$ and $\xrightarrow{P/\Theta}$, respectively. These concepts are defined rigorously in Definition 2.B.2 in the supplement. We note that uniform convergence reduces to classical (pointwise) convergence if Θ is a singleton, which corresponds to fixing the sampling distribution. We also introduce the parameter subset

$$\Theta_0 := \{\theta \in \Theta \mid H_0 \text{ is valid}\}, \quad (2.4.24)$$

consisting of all parameter values for which the hypothesis of conditional local independence holds. Correspondingly, we will use $\xrightarrow{D/\Theta_0}$ and $\xrightarrow{P/\Theta_0}$ to denote stochastic convergences uniformly over Θ_0 .

We are now ready to formulate the underlying assumptions on the data required for our asymptotic results. These assumptions may appear strong, but we argue in the discussion in Section 7 that they are not unreasonable from a practical viewpoint.

Assumption 2.4.1. *There exist constants $C, C' > 0$, such that for any parameter value $\theta \in \Theta$:*

- i) *The \mathcal{G}_t^θ -intensity λ_t^θ of N^θ is càglàd with $\sup_{0 \leq t \leq 1} \lambda_t^\theta \leq C$ almost surely.*
- ii) *The residual process G^θ is càglàd with $\sup_{0 \leq t \leq 1} |G_t^\theta| \leq C'$ almost surely.*

The estimator, $\hat{\lambda}_t^{(n)}$, of λ_t and the estimator, $\hat{G}_t^{(n)}$, of the residual process are assumed to satisfy the same bounds as λ_t and G_t . We note that Assumption 2.4.1 i) implies that M_t is a true \mathcal{G}_t -martingale, and by the innovation theorem, $\lambda_t = \mathbb{E}[\lambda_t \mid \mathcal{F}_{t-}]$. As a consequence, the \mathcal{F}_t -intensity λ_t inherits the boundedness from the \mathcal{G}_t -intensity λ_t , and M_t is an \mathcal{F}_t -martingale. More generally, we have the following proposition ensuring that stochastic integrals are true martingales, e.g., that I_t is a martingale under H_0 .

Proposition 2.4.1. *Under Assumption 2.4.1 it holds that each of the processes*

$$\left(\int_0^t f(G_s) dM_s \right)_{t \in [0,1]} \quad \text{and} \quad \left(\int_0^t f(\hat{G}_s^{(n)}) dM_s \right)_{t \in [0,1]}$$

are mean zero, square integrable \mathcal{G}_t -martingales for any $f \in C(\mathbb{R})$.

To express the asymptotic distribution of $U^{(n)}$ we need its *variance function*.

Definition 2.4.2. We define the variance function $\mathcal{V}: [0, 1] \rightarrow [0, \infty]$ as

$$\mathcal{V}(t) = \mathbb{E} \left(\int_0^t G_s^2 dN_s \right) \quad (2.4.25)$$

for $t \in [0, 1]$. ♣

As everything else, the variance function, $\mathcal{V} = \mathcal{V}^\theta$, is also indexed by the parameter θ , which we, for notational simplicity, suppress unless explicitly needed.

By taking $f(x) = x^2$ in Proposition 2.4.1, Assumption 2.4.1 implies that for each $t \in [0, 1]$,

$$\mathcal{V}(t) = \mathbb{E} \left(\int_0^t G_s^2 \lambda_s ds \right) < \infty.$$

Moreover, $\mathcal{V}(t)$ is the variance of $\int_0^t G_s dM_s$, which under H_0 is the same as the variance of $I_t = \int_0^t G_s dM_s$.

With the assumptions above we can prove the following proposition about the uniform distributional limit of the process $U^{(n)}$ in the Skorokhod space $D[0, 1]$, the space of càdlàg functions from $[0, 1]$ to \mathbb{R} endowed with the Skorokhod topology. A corresponding pointwise result is an application of Rebolledo's classical martingale CLT. Our generalization to uniform convergence is based on a uniform extension of Rebolledo's theorem, see Theorem 2.C.4 in Section 2.C in the supplement.

Proposition 2.4.3. *Under Assumption 2.4.1 it holds that*

$$U^{(n), \theta} \xrightarrow{D/\Theta} U^\theta$$

in $D[0, 1]$ as $n \rightarrow \infty$, where for each $\theta \in \Theta$, U^θ is a mean zero continuous Gaussian martingale on $[0, 1]$ with variance function \mathcal{V}^θ .

To control the remainder terms in (2.4.17) we will bound the estimation errors in terms of the 2-norm, $\|\cdot\|_2$, on $L_2([0, 1] \times \Omega)$, i.e.,

$$\|W\|_2^2 = \mathbb{E} \left(\int_0^1 W_s^2 ds \right)$$

for any process $W \in L_2([0, 1] \times \Omega)$. We will make the following consistency assumptions on $\hat{\lambda}^{(n)}$ and $\hat{G}^{(n)}$.

Assumption 2.4.2. *Assume that $|J_n| \rightarrow \infty$ when $n \rightarrow \infty$ and let*

$$g^\theta(n) = \left\| G^\theta - \hat{G}^{(n), \theta} \right\|_2 \quad \text{and} \quad h^\theta(n) = \left\| \lambda^\theta - \hat{\lambda}^{(n), \theta} \right\|_2.$$

Then each of the sequences $g^\theta(n)$, $h^\theta(n)$, and $\sqrt{|J_n|} g^\theta(n) h^\theta(n)$ converge to zero uniformly over Θ as $n \rightarrow \infty$, i.e.,

$$\lim_{n \rightarrow \infty} \sup_{\theta \in \Theta} \max\{g^\theta(n), h^\theta(n), \sqrt{|J_n|} g^\theta(n) h^\theta(n)\} = 0.$$

With this assumption we can establish that the remainder terms also converge uniformly to the zero-process.

Proposition 2.4.4. *Under Assumptions 2.4.1 and 2.4.2, it holds that*

$$\sup_{t \in [0, 1]} |R_{i,t}^{(n), \theta}| \xrightarrow{P/\Theta} 0$$

as $n \rightarrow \infty$ for $i = 1, 2, 3$.

2 Nonparametric conditional local independence testing

To control the asymptotic behavior of the LCM estimator in the alternative we need to control the two terms $D_1^{(n)}$ and $D_2^{(n)}$.

Proposition 2.4.5. *Let Assumptions 2.4.1 and 2.4.2 hold true.*

- i) *The stochastic process $D_1^{(n),\theta} - \sqrt{|J_n|}\gamma^\theta$ converges in distribution in $(C[0,1], \|\cdot\|_\infty)$ uniformly over Θ as $n \rightarrow \infty$.*
- ii) *If $G_t^\theta = X_t^\theta - \Pi_t^\theta$ is the additive residual process, then $D_2^{(n),\theta} \xrightarrow{P/\Theta} 0$ in $D[0,1]$ as $n \rightarrow \infty$.*

We note that $D_2^{(n)}$ might not vanish without an assumption like G_t being the additive residual process, and it is not clear if $D_2^{(n)}$ will even converge in general. We will not pursue an analysis of the asymptotic behavior of $D_2^{(n)}$ in the general case. We note, however, that if we can estimate G with a parametric rate, that is, $\sqrt{|J_n|}g(n) = O(1)$, then it follows from the Cauchy-Schwarz inequality that $D_2^{(n)}$ is stochastically bounded, and $D_1^{(n)}$ still dominates in the alternative where $\gamma \neq 0$.

We can combine all of the propositions into a single theorem regarding the asymptotics of the LCM estimator, which we consider as our main result.

Theorem 2.4.6. *Let Assumptions 2.4.1 and 2.4.2 hold true.*

- i) *It holds that*

$$\sqrt{|J_n|}\hat{\gamma}^{(n),\theta} \xrightarrow{D/\Theta_0} U^\theta$$

in $D[0,1]$ as $n \rightarrow \infty$, where for each $\theta \in \Theta_0$, U^θ is a mean zero continuous Gaussian martingale on $[0,1]$ with variance function \mathcal{V}^θ .

- ii) *For the additive residual process it holds that for every $\varepsilon > 0$ there exists $K > 0$ such that*

$$\limsup_{n \rightarrow \infty} \sup_{\theta \in \Theta} \mathbb{P} \left(\sqrt{|J_n|} \cdot \|\hat{\gamma}^{(n),\theta} - \gamma^\theta\|_\infty > K \right) < \varepsilon. \quad (2.4.26)$$

Thus we have established the weak asymptotic limit of $\sqrt{|J_n|}\hat{\gamma}^{(n)}$ under H_0 . However, the variance function \mathcal{V} of the limiting Gaussian martingale is unknown and must be estimated from data. We propose to use the empirical version of (2.4.25),

$$\hat{\mathcal{V}}_n(t) = \frac{1}{|J_n|} \sum_{j \in J_n} \int_0^t \left(\hat{G}_{j,s}^{(n)} \right)^2 dN_{j,s} = \frac{1}{|J_n|} \sum_{j \in J_n} \sum_{\tau \leq t: \Delta N_{j,s}=1} \left(\hat{G}_{j,\tau}^{(n)} \right)^2, \quad (2.4.27)$$

for which we have the following consistency result.

Proposition 2.4.7. *Under Assumptions 2.4.1 and 2.4.2 it holds that*

$$\sup_{t \in [0,1]} |\hat{\mathcal{V}}_n^\theta(t) - \mathcal{V}^\theta(t)| \xrightarrow{P/\Theta} 0,$$

as $n \rightarrow \infty$.

We emphasize that \mathcal{V} is only the asymptotic variance function of the LCM estimator under H_0 . It is always the asymptotic variance function of $U^{(n)}$, but in the alternative the asymptotic distribution of $\hat{\gamma}^{(n)}$ also involves the asymptotic distribution of $D_1^{(n)}$ and is thus more complicated.

Tests of conditional local independence can now be constructed in terms of univariate functionals of $\hat{\gamma}^{(n)}$ and $\hat{\mathcal{V}}_n$ that quantify the magnitude of the LCM. The asymptotics of such test statistics under H_0 are described in the following corollary, which is essentially an application of the continuous mapping theorem.

Corollary 2.4.8. *Let $\mathcal{J}: D[0, 1] \times D[0, 1] \rightarrow \mathbb{R}$ be a functional that is continuous on the closed subset $C[0, 1] \times \{\mathcal{V}^\theta: \theta \in \Theta_0\}$ with respect to the uniform topology, i.e., the topology generated by the norm $\|(f_1, f_2)\| = \max\{\|f_1\|_\infty, \|f_2\|_\infty\}$ for $f_1, f_2 \in D[0, 1]$. Define the test statistic*

$$\hat{D}_n^\theta = \mathcal{J}\left(\sqrt{|J_n|}\hat{\gamma}^{(n),\theta}, \hat{\mathcal{V}}_n^\theta\right).$$

Under Assumptions 2.4.1 and 2.4.2, it holds that

$$\hat{D}_n^\theta \xrightarrow{\mathcal{D}/\Theta_0} \mathcal{J}(U^\theta, \mathcal{V}^\theta), \quad n \rightarrow \infty, \quad (2.4.28)$$

where U^θ is a mean zero continuous Gaussian martingale with variance function \mathcal{V}^θ .

2.5 The Local Covariance Test

In this section we introduce a practically applicable test based on the LCM estimator. Using the asymptotic distribution of the LCM estimator we show that the asymptotic distribution of our proposed test is independent of the sampling distribution under H_0 and has an explicit representation. We show, in addition, uniform asymptotic level of the test, and we give a uniform power result for the additive residual process. Finally, we modify the test to be based on a cross-fitted estimator of the LCM instead of using sample splitting, and we show uniform level of that test.

To construct a test statistic based on the LCM estimator it is beneficial that its distributional limit does not depend on the variance function. As a simple example, consider the endpoint test statistic:

$$\left(\hat{\mathcal{V}}_n(1)\right)^{-\frac{1}{2}}\sqrt{|J_n|}\hat{\gamma}_1^{(n)}, \quad (2.5.29)$$

which under H_0 converges in distribution to $\mathcal{V}(1)^{-\frac{1}{2}}U_1$ by Corollary 2.4.8. The distribution of the latter is the standard normal distribution, and in particular it does not depend on \mathcal{V} .

Any test statistic constructed from $\hat{\gamma}^{(n)}$ should capture deviations of γ_t away from 0. The test statistic in (2.5.29) does, however, only consider the endpoint of the process, and since γ is not necessarily monotone, γ_t may deviate more from 0 for other $t \in [0, 1]$. Thus in order to increase power against such alternatives we consider the test statistic

$$\hat{T}_n = \frac{\sqrt{|J_n|}}{\sqrt{\hat{\mathcal{V}}_n(1)}} \sup_{0 \leq t \leq 1} |\hat{\gamma}_t^{(n)}|. \quad (2.5.30)$$

2 Nonparametric conditional local independence testing

We refer to \widehat{T}_n as the *Local Covariance Test statistic* (LCT statistic). We proceed to show that the LCT statistic can be turned into a test of H_0 with asymptotic level α , and which has asymptotic power against any alternative with a non-zero LCM. This is the best we can hope for of any test based on the LCM estimator.

We note that it might be possible to establish similar results for other norms of the LCM, for example, a statistic based on a weighted L_2 -norm. However, since other norms of the distributional limit U will generally have a distribution with a complicated dependency on \mathcal{V} , we believe that the LCT statistic is the simplest to construct.

To establish uniform asymptotic level via Corollary 2.4.8 for tests based on test statistics such as (2.5.30) we need to assume that the asymptotic variances in $t = 1$ are uniformly bounded away from zero.

Assumption 2.5.1. *There exists $\delta_1 > 0$ such that for all $\theta \in \Theta$ it holds that $\mathcal{V}^\theta(1) \geq \delta_1$.*

2.5.1 Type I and type II error control

We proceed to show that under H_0 , the LCT statistic is distributed as $\sup_{0 \leq t \leq 1} |B_t|$, where (B_t) is a standard Brownian motion. From this point onwards, we let S denote a random variable with such a distribution and note that its CDF can be written as:

$$F_S(x) = \mathbb{P}(S \leq x) = \frac{4}{\pi} \sum_{k=0}^{\infty} \frac{(-1)^k}{2k+1} \exp\left(-\frac{\pi^2(2k+1)^2}{8x^2}\right), \quad x > 0. \quad (2.5.31)$$

See, for example, Section 12.2 in Schilling and Partzsch [2012] where the formula is derived from Lévy's triple law.

The p -value for a test of H_0 equals $1 - F_S(\widehat{T}_n)$, and since the series in (2.5.31) converges at an exponential rate, the p -value can be computed with high numerical precision by truncating the series. Given a significance level $\alpha \in (0, 1)$, we also let $z_{1-\alpha}$ denote the $(1 - \alpha)$ -quantile of F_S , which exists and is unique since the right-hand side of (2.5.31) is strictly increasing and continuous. The *Local Covariance Test* (LCT) with significance level α is then defined by

$$\Psi_n = \Psi_n^\alpha = \mathbb{1}(F_S(\widehat{T}_n) > 1 - \alpha) = \mathbb{1}(\widehat{T}_n > z_{1-\alpha}). \quad (2.5.32)$$

From Theorem 2.4.6 we can now deduce the asymptotic properties of the LCT under the hypothesis of conditional local independence. Recall that $\xrightarrow{\mathcal{D}/\Theta_0}$ denotes uniform convergence in distribution under H_0 .

Theorem 2.5.1. *Let Assumptions 2.4.1, 2.4.2 and 2.5.1 hold true. Then it holds that*

$$\widehat{T}_n^\theta \xrightarrow{\mathcal{D}/\Theta_0} S$$

as $n \rightarrow \infty$. As a consequence, for any $\alpha \in (0, 1)$,

$$\limsup_{n \rightarrow \infty} \sup_{\theta \in \Theta_0} \mathbb{P}(\Psi_n^{\alpha, \theta} = 1) \leq \alpha.$$

In other words, the *Local Covariance Test* defined in (2.5.32) has uniform asymptotic level α .

In general, we cannot expect that the test has power against alternatives to H_0 for which the LCM is the zero-function. This is analogous to other types of conditional independence tests based on conditional covariances, e.g., GCM [Shah and Peters, 2020]. However, we do have the following result that establishes power against local alternatives with $\|\gamma\|_\infty$ decaying at an order of at most $|J_n|^{-1/2}$.

Theorem 2.5.2. *Let Assumptions 2.4.1 and 2.4.2 hold true. Using the additive residual process it holds that for any $0 < \alpha < \beta < 1$ there exists $c > 0$ such that*

$$\liminf_{n \rightarrow \infty} \inf_{\theta \in \mathcal{A}_{c,n}} \mathbb{P}(\Psi_n^{\alpha,\theta} = 1) \geq \beta,$$

where $\mathcal{A}_{c,n} = \{\theta \in \Theta \mid \|\gamma^\theta\|_\infty \geq c|J_n|^{-1/2}\}$.

2.5.2 Extension to cross-fitting

In Section 2.4 we considered sample splitting with observations indexed by J_n^c used to estimate the two models and with observations indexed by J_n used to estimate γ . Following Chernozhukov et al. [2018], we can improve efficiency by cross-fitting, i.e., by flipping the roles of J_n and J_n^c to obtain a second equivalent estimator of γ . Heuristically, the two estimators are approximately independent, and thus their average should be a more efficient estimator. This procedure generalizes directly to a partition $J_n^1 \cup \dots \cup J_n^K = \{1, \dots, n\}$ of the indices into K disjoint folds. The partition is assumed to have a uniform asymptotic density, meaning that $|J_n^k|/n \rightarrow \frac{1}{K}$ as $n \rightarrow \infty$ for each k .

We estimate G and λ using $(J_n^k)^c = \{1, \dots, n\} \setminus J_n^k$ and subsequently estimate γ using J_n^k . Then the K -fold *Cross-fitted LCM estimator*, abbreviated as X-LCM, is defined as the average LCM estimator over the K folds, i.e.,

$$\check{\gamma}_t^{K,(n)} = \frac{1}{K} \sum_{k=1}^K \frac{1}{|J_n^k|} \sum_{j \in J_n^k} \int_0^t \hat{G}_{j,s}^{k,(n)} d\hat{M}_{j,s}^{k,(n)}, \quad (2.5.33)$$

where for each $j \in J_n^k$, the processes $\hat{G}_j^{k,(n)}$ and $\hat{M}_j^{k,(n)}$ are the model predictions of G_j and M_j , respectively, based on training data indexed by $(J_n^k)^c$. We also define a K -fold version of the variance estimator:

$$\check{V}_n^K(t) = \frac{1}{K} \sum_{k=1}^K \frac{1}{|J_n^k|} \sum_{j \in J_n^k} \int_0^t \left(\hat{G}_{j,s}^{k,(n)} \right)^2 dN_{j,s}. \quad (2.5.34)$$

Now, similarly to the LCT statistic, the cross-fitted estimator can be used to construct a test statistic,

$$\check{T}_n^K = \sqrt{\frac{n}{\check{V}_{K,n}(1)}} \sup_{0 \leq t \leq 1} \left| \check{\gamma}_t^{K,(n)} \right|, \quad (2.5.35)$$

from which we define the following test of conditional local independence.

2 Nonparametric conditional local independence testing

Definition 2.5.3. Let $\alpha \in (0, 1)$ and let \tilde{T}_n^K be the test statistic from (2.5.35). The K -fold *Cross-fitted Local Covariance Test* (X-LCT) with significance level α is defined by

$$\check{\Psi}_n^K = \mathbb{1}(F_S(\tilde{T}_n^K) > 1 - \alpha) = \mathbb{1}(\tilde{T}_n^K > z_{1-\alpha}),$$

where $z_{1-\alpha}$ is the $(1 - \alpha)$ -quantile of the distribution function F_S given in (2.5.31). ♣

We provide a summary of the computation of the X-LCT in Algorithm 2. The asymptotic analysis of $\hat{\gamma}^{(n)}$ generalizes to $\hat{\gamma}^{K,(n)}$, but we will refrain from restating all results for the K -fold cross-fitted estimator. For simplicity, we focus on the fact that the X-LCT has asymptotic level α .

Theorem 2.5.4. *Suppose that Assumption 2.4.2 is satisfied for every sample split $J_n^k \cup (J_n^k)^c, k = 1, \dots, K$. Under Assumptions 2.4.1 and 2.5.1, the X-LCT statistic satisfies*

$$\tilde{T}_n^{K,\theta} \xrightarrow{D/\Theta_0} S$$

for $n \rightarrow \infty$. In particular, the X-LCT has uniform asymptotic level α .

Note that cross-fitting recovers full efficiency in the sense that the scaling factor is \sqrt{n} rather than $\sqrt{|J_n|}$, which leads to a more powerful test. Moreover, the asymptotic distribution of \tilde{T}_n^K does not depend on the number of folds K , and any difference between various choices of K can thus be attributed to finite sample errors. Larger values of K will allocate more data to estimation of G and λ , which intuitively should be the harder estimation problem. Following Remark 3.1 in Chernozhukov et al. [2018], we believe that a default choice of $K = 4$ or $K = 5$ should be reasonable in practice.

2.6 Simulation study

In this section we present the results from a simulation study based on the Cox example introduced in Section 2.2.2. We elaborate in Section 2.6.1 on the full model specification used for the simulation study – which will also illuminate how Π and λ can be modeled and estimated. The results from the simulation study focus on the distribution of the X-LCT statistic \tilde{T}_n^K and validate the asymptotic level and power of the X-LCT $\check{\Psi}_n^K$. The latter is also compared to a hazard ratio test based on the marginal Cox model (2.2.11). The simulations were implemented in Python and the code is available¹.

2.6.1 Cox model continued

Consider the same setup as in Section 2.2.2. To fully specify the model we need to specify the distribution of the processes X, Y and Z . We suppose that X and Y can be written in terms of Z as

$$X_t = \int_0^t Z_s \rho_X(s, t) ds + V_t, \quad \text{and} \quad Y_t = \int_0^t Z_s \rho_Y(s, t) ds + W_t, \quad (2.6.36)$$

¹<https://github.com/AlexanderChristgau/nonparametric-cli-test>

Algorithm 2: K -fold cross-fitted local covariance test (X-LCT)

-
- 1 **input:** processes $(N_j, X_j, Z_j)_{j=1, \dots, n}$, partition $J_n^1 \cup \dots \cup J_n^K$ of indices into K folds;
 - 2 **options:** historical regression methods for estimation of λ and G given (N, Z) ,
 - 3 discrete time grid $\mathbb{T} \subset [0, 1]$, significance level $\alpha \in (0, 1)$;
 - 4 **begin**
 - 5 **for** $k = 1, \dots, K$ **do**
 - 6 apply Algorithm 1 on sample split $J_n^k \cup (J_n^k)^c$ to compute $\tilde{\gamma}^{k,(n)}$ on the grid \mathbb{T} ;
 - 7 use Equation (2.4.27) on sample split $J_n^k \cup (J_n^k)^c$ to compute $\tilde{V}_{k,n}(1)$;
 - 8 compute $\tilde{\gamma}^{K,(n)} = \frac{1}{K} \sum_{k=1}^K \tilde{\gamma}^{k,(n)}$ on grid \mathbb{T} ;
 - 9 compute $\tilde{V}_{K,n}(1) = \frac{1}{K} \sum_{k=1}^K \tilde{V}_{k,n}(1)$;
 - 10 compute the X-LCT statistic $\tilde{T}_n^K = \sqrt{n} \cdot \max_{t \in \mathbb{T}} |\tilde{\gamma}_t^{K,(n)}| / \sqrt{\tilde{V}_{K,n}(1)}$;
 - 11 compute p -value $\tilde{p} = 1 - F_S(\tilde{T}_n^K)$ by truncating the series in Equation (2.5.31).
 - 12 **output:** the X-LCT $\tilde{\Psi}_n^K = \mathbb{1}(\tilde{p} < \alpha)$, and optionally the p -value \tilde{p} ;
-

where ρ_X and ρ_Y are two functions defined on the triangle $\{(s, t) \in [0, 1]^2 \mid s \leq t\}$, and where $V = (V_t)_{0 \leq t \leq 1}$ and $W = (W_t)_{0 \leq t \leq 1}$ are two noise processes with mean zero. The processes Z , V and W are assumed independent, which implies (2.2.10) and thus that N is conditionally locally independent of X given $\mathcal{F}_t = \mathcal{F}_t^{N,Z}$.

The specific dependency of X and Y on Z is known as the *historical functional linear model* in functional data analysis [Malfait and Ramsay, 2003]. Within this model,

$$\Pi_t = \mathbb{E}(X_t \mid \mathcal{F}_{t-}) = \int_0^t Z_s \rho_X(s, t) ds, \quad (2.6.37)$$

and on $(T \geq t)$

$$\mathbb{E}(e^{Y_t} \mid \mathcal{F}_t) = e^{\int_0^t Z_s \rho_Y(s, t) ds} \mathbb{E}(e^{W_t} \mid T \geq t) = e^{\tilde{\beta}_0(t) + \int_0^t Z_s \rho_Y(s, t) ds},$$

where $\tilde{\beta}_0(t) = \log(\mathbb{E}(e^{W_t} \mid T \geq t))$. Since

$$\lambda_t = \mathbb{1}(T \geq t) \lambda_t^0 e^{\beta Z_t} \mathbb{E}(e^{Y_t} \mid \mathcal{F}_t),$$

it follows that on $(T \geq t)$,

$$\log(\lambda_t) = \beta_0(t) + \beta Z_t + \int_0^t Z_s \rho_Y(s, t) ds, \quad (2.6.38)$$

where the two baseline terms depending only on time have been merged into β_0 .

The computations above suggest how the estimators $\hat{\lambda}^{(n)}$ and $\hat{\Pi}^{(n)}$ could be constructed. That is, $\hat{\lambda}^{(n)}$ could be based on estimates of β , β_0 and ρ_Y from the observations

2 Nonparametric conditional local independence testing

$(T_j, Z_j)_{j \in J_n^c}$, and $\widehat{\Pi}^{(n)}$ could be based on estimates of ρ_X from $(X_j, Z_j)_{j \in J_n^c}$. We would then have

$$\widehat{\Pi}_{j,t}^{(n)} = \int_0^t Z_{j,s} \widehat{\rho}_X^{(n)}(s, t) ds$$

for $j \in J_n$ where $\widehat{\rho}_X^{(n)}$ denotes the estimate of ρ_X , and similarly for $\widehat{\lambda}_{j,t}^{(n)}$. Particular choices of estimators $\widehat{\rho}_X^{(n)}$ and $\widehat{\rho}_Y^{(n)}$ and their theoretical properties are reviewed in Section 2.D in the supplement. Our conclusion from this review is that for the historical functional linear model, sufficient rate results should be possible but have not yet been established rigorously.

2.6.2 Sampling scheme

The actual time-discretized simulations and computations were implemented using an equidistant grid $\mathbb{T} = (t_i)_{i=1}^q$ with $q = 128$ time points $0 = t_1 < \dots < t_q = 1$. Inspired by Harezlak et al. [2007], we generated the processes as follows: let $\xi \in \mathbb{R}^3$ and $V, W, \mathbb{W} \in \mathbb{R}^{\mathbb{T}}$ be independent random variables such that $\xi \sim \mathcal{N}(0, I_3)$ and such that V, W , and \mathbb{W} are identically distributed with $V_{t_1}, V_{t_2} - V_{t_1}, \dots, V_{t_q} - V_{t_{q-1}} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1/q)$. Then the process Z is determined by $Z_t = \xi_1 + \xi_2 t + \sin(2\pi\xi_3 t) + \mathbb{W}_t$ for $t \in \mathbb{T}$. The processes X and Y were then given by the historical linear model (2.6.36) with kernels ρ_X and ρ_Y being one of the following four kernels:

$$\begin{aligned} \text{zero: } (s, t) &\mapsto 0, & \text{constant: } (s, t) &\mapsto 1, \\ \text{Gaussian: } (s, t) &\mapsto e^{-2(t-s)^2}, & \text{sine: } (s, t) &\mapsto \sin(4t - 20s). \end{aligned}$$

To compute X and Y , we evaluated the kernels on $\{(s, t) \in \mathbb{T}^2 \mid s \leq t\}$ and approximated the integrals by Riemann sums. The full intensity for $N_t = \mathbb{1}(T \leq t)$ was specified with a Weibull baseline of the form $\lambda_t^{\text{full}} = \mathbb{1}(T \geq t) \beta_1 t^2 \exp(\beta_2 Z_t + Y_t)$, for $\beta_1 > 0$ and a choice of $\beta_2 \in \{-1, 1\}$. To sample T we applied the inverse hazard method, which utilizes that Λ_T^{full} is standard exponentially distributed. That is, we sampled $E \sim \text{Exp}(1)$ and numerically computed $T = \max\{t \in \mathbb{T} \mid \Lambda_t^{\text{full}} < E\}$ as a discretized approximation. For any given parameter setting, the baseline coefficient β_1 was chosen sufficiently large to ensure that $\Lambda_t^{\text{full}} \geq E$ would occur before time $t = 1$ in more than $\frac{q-1}{q} \cdot n$ samples.

The simulation setting used to sample the data for Figures 2.2.3 and 2.2.4 was $\beta_2 = -1$ and $\rho_X = \rho_Y = \text{constant}$.

With this setup, Assumption 2.4.1 is satisfied if V, W and \mathbb{W} were bounded. Since we use the Gaussian distribution, they are technically not bounded, but they could be made bounded by introducing a lower and upper cap. Due to the light tails of the Gaussian distribution such caps would have no noticeable effect on the simulation results, and the results we report are generated without a cap.

The implementation details for the X-LCT and the hazard ratio test are given in Section 2.G.1 in the supplement.

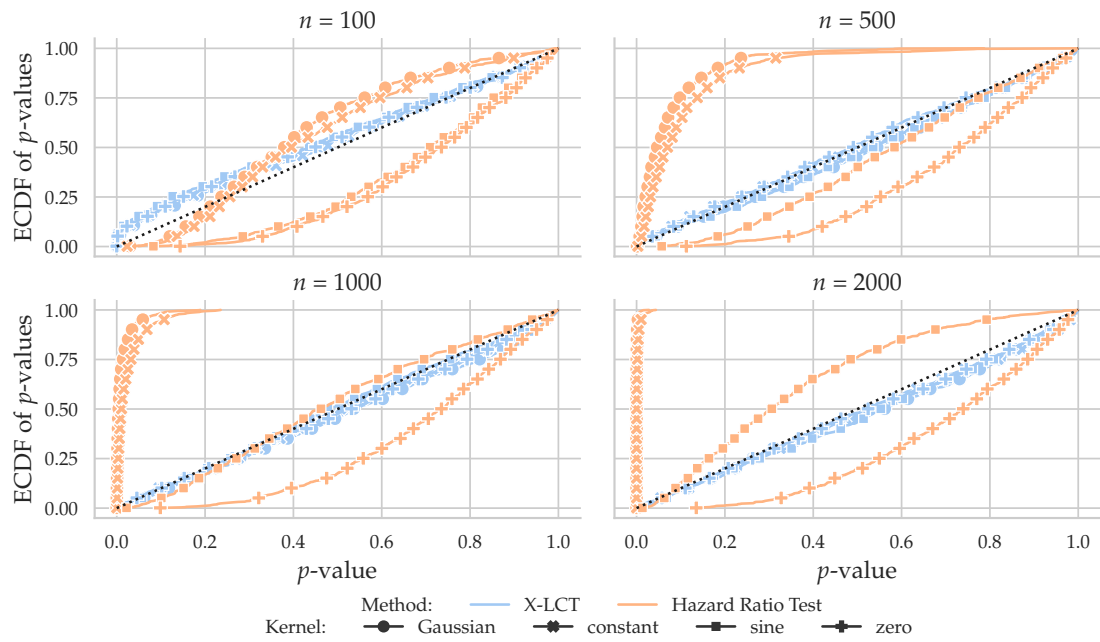


Figure 2.6.5: Empirical cumulative distribution functions of simulated p -values for the cross-fitted local covariance test and the hazard ratio test. The simulated data satisfies the hypothesis of conditional local independence, so the p -values are supposed to be uniformly distributed, and the CDF should fall on the diagonal dotted line.

2.6.3 Distributions of p -values under H_0

We examine the distributional approximation $\check{T}_n^K \stackrel{\text{as}}{\approx} S$, cf. Theorem 2.5.4, by comparing the p -values $1 - F_S(\check{T}_n^K)$ to a uniform distribution. Figure 2.6.5 shows the empirical distribution functions of the p -values computed from data simulated according to the scheme described in the previous section. The results are aggregated over the two choices of $\beta_2 \in \{-1, 1\}$ since these two settings were found to be similar. For more detailed results from the experiment, see Figure 2.G.1 in Section 2.G in the supplement, which also includes the p -values corresponding to the endpoint test statistic.

For the hazard ratio test, Figure 2.6.5 shows that the p -values are sub-uniform for the zero-kernel. In this case, the marginal Cox model is correct, and the non-uniformity of the p -values can be explained by the L_2 -penalization. For the constant and Gaussian kernels the hazard ratio test fails completely, whereas for the sine kernel, the mediated effect of Z on T through Y is more subtle, and the model misspecification only becomes apparent for $n = 2000$. Overall, these results are consistent with the reasoning in the Section 2.2.2: a test based on the misspecified Cox model will wrongly reject the hypothesis of conditional local independence.

For the proposed X-LCT, Figure 2.6.5 shows that the associated p -values are slightly anti-conservative for $n = 100$. This is to be expected, and can be explained by the finite sample errors leading to more extreme values of \check{T}_n^K than the approximation by S . As n increases, these errors become smaller – and for $n = 2000$ the p -values actually seem to be sub-uniform. The sub-uniformity may be explained by the time discretization, since the maximum of the process is taken over \mathbb{T} rather than $[0, 1]$. Figure 2.G.3 in Section 2.G in the supplement illustrates the asymptotic effect of the time discretization which supports this claim. Another support of this claim is that the endpoint test does not appear to give sub-uniform p -values for large n , see Figure 2.G.1. We finally note that the distributions of the p -values for our proposed test is largely unaffected by the kernel used to generate the data.

2.6.4 Power against local alternatives

To investigate the power of the X-LCT we construct local alternatives to H_0 in accordance with the right graph in Figure 2.2.2 by replacing Y_t by the process $Y_t + \frac{\rho_0}{\sqrt{n}}X_t$. That is, for $\rho_0 \neq 0$, blood pressure is then directly affected by pension savings, and N_t is no longer conditionally locally independent of X_t given \mathcal{F}_t . In terms of the full intensity, these local alternatives are equivalent to

$$\lambda_t^{\text{full}} = \mathbb{1}(T \geq t) \beta_1 t^2 \exp \left(\beta_2 Z_t + Y_t + \frac{\rho_0}{\sqrt{n}} X_t \right). \quad (2.6.39)$$

We simulated data for the dependency parameter $\rho_0 \in \{0, 5, 10\}$. Note that $\rho_0 = 0$ corresponds to our previous sampling scheme with conditional local independence. For each of the $96 = 4 \times 2 \times 4 \times 3$ choices of kernel, β_2 , n and ρ_0 we ran the tests 400 times and computed the p -values. For simplicity, we report the rejection rate at an $\alpha = 5\%$ significance level and the results are shown in Figure 2.6.6.

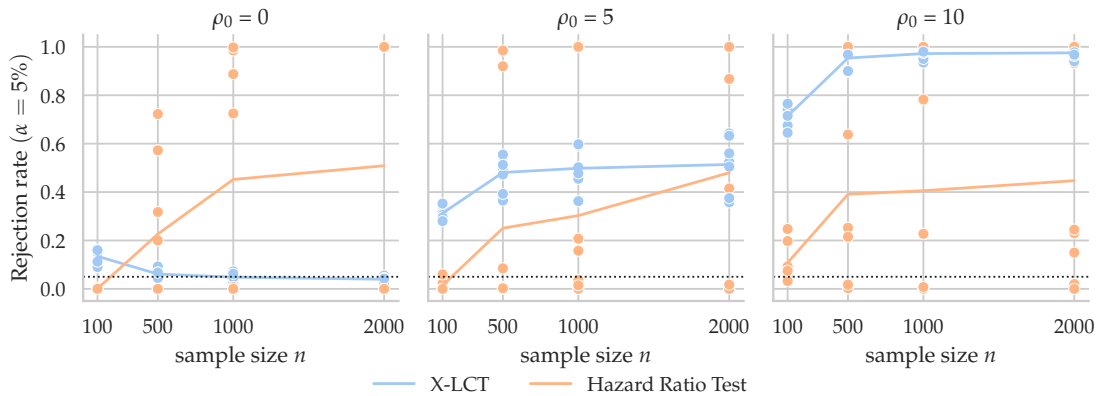


Figure 2.6.6: For each $\rho_0 \in \{0, 5, 10\}$, the lines show the average rejection rates of our proposed test X-LCT (blue) and the hazard ratio test (orange) as functions of sample size, with each average taken over 8 different settings. For each setting, the rejection rate is computed from 400 simulated datasets at a 5% significance level and the rejection rate is displayed with a dot.

In the leftmost panel, the data was generated under H_0 and the plot shows what we noted previously, namely that the X-LCT holds level for large n , whereas the hazard ratio test does not. For the local alternatives, $\rho_0 = 5$ and $\rho_0 = 10$, we note that the power of the hazard ratio test is quite sensitive to the simulation settings. For some settings it has no power, while for others it has some power.

In contrast, the proposed X-LCT has power against all of the local alternatives. The power increases with n initially but stabilizes from around $n = 1000$. This is similar to the behavior observed under the null hypothesis and is not surprising. We expect that the sample size needs to be sufficiently large for the nonparametric estimators to work sufficiently well, and we expect the sufficient sample size to be mostly unaffected by the value of ρ_0 . For fixed n , we also note that the power of $\check{\Psi}_n^K$ is fairly robust with respect to the choice of β_2 and the choice of kernel. Overall, we find that the X-LCT is applicable in these settings with historical effects: it has consistent power against the \sqrt{n} alternatives while controlling type I error for n reasonably large.

In Section 2.G in the supplement, we provide additional numerical results for time-varying alternatives, and we compare the X-LCT with its endpoint counterpart.

2.7 Discussion

The LCM was introduced as a functional parameter that quantifies deviations from the hypothesis H_0 of conditional local independence. We showed how the parameter may be expressed in several ways, but that it is the representation in terms of the residual process that allows us to estimate the LCM with a \sqrt{n} -rate under H_0 without parametric model assumptions. The residual process was introduced as an abstract model of X_t for each t given the history up to time t , and we showed that such a residualization could

be viewed as a form of orthogonalization. Similar ideas have been used recently for classical conditional independence testing, such as GCM [Shah and Peters, 2020], tests based on the partial copula [Petersen and Hansen, 2021], and GHCM [Lundborg et al., 2022b]. It is, however, not possible to use any of these to test H_0 , which cannot be expressed as a classical conditional independence. Our test based on the LCM is the first nonparametric test of conditional local independence with substantial theoretical support, and we propose to test H_0 in practice by using X-LCT based on the cross-fitted estimator of LCM.

Contrary to the tests of conditional independence mentioned above, we need sample splitting – even under H_0 – to achieve our asymptotic results. We do not believe that this can be avoided. The standard argument to avoid this uses classical conditional independence in a crucial way, which does not translate into our framework – basically because we condition on information that changes with time. Our simulation study also indicates that sample splitting or cross-fitting is needed in practice for the LCM estimator to be unbiased under H_0 .

While our cross-fitted estimator of the LCM, the X-LCM, share some of the general patterns of other double machine learning procedures – including the overall decomposition (2.4.17) – our analysis and results required a range of generalizations of known results and some novel ideas. The asymptotic distribution of the leading term, $U^{(n)}$, is also a well known consequence of Rebolledo’s CLT, see, e.g., Section V.4 in [Andersen et al., 1993] for related results in the context of survival analysis. However, we generalized this result to uniform convergence in the Skorokhod space $D[0, 1]$, and we introduced new techniques for handling the remainder terms. These novel techniques are made necessary by the decomposition (2.4.17) being a decomposition of stochastic processes indexed by time. We outline below the three most important technical contributions we made.

First, to obtain uniform control of level and power, all asymptotic results in Section 2.4 are formulated in terms of uniform stochastic convergence. Since this notion of convergence had not previously been considered on general metric spaces, and especially not on the Skorokhod space, we had to develop the necessary theory. This development could be of independent interest, and we have collected the general definitions and main results on uniform stochastic convergence in metric spaces in Section 2.B in the supplement. This framework also allowed us to show a uniform version of Rebolledo’s martingale CLT in Section 2.C in the supplement.

Second, to establish distributional convergence under H_0 , we need to control the remainder terms $R_{i,t}^{(n)}$ uniformly over t . The third term, $R_3^{(n)}$, is simple to bound, and by exploiting Doob’s submartingale inequality, the second term, $R_2^{(n)}$, can also be bounded. The most difficult first term, $R_1^{(n)}$, was controlled using stochastic equicontinuity via an exponential tail bound and the use of the chaining lemma. The necessary general uniform stochastic equicontinuity and chaining arguments are collected in Section 2.B.3 in the supplement.

Third, to achieve rate results in the alternative, the processes $D_1^{(n)}$ and $D_2^{(n)}$ must be controlled. The process $D_1^{(n)}$ does, like $U^{(n)}$, not involve any estimation, and its dis-

tributional convergence follows from a general CLT argument for continuous stochastic processes. The term $D_2^{(n)}$ is more difficult to handle, as it may not have mean zero if G_t is not the additive residual process. However, X_t cancels out in $\widehat{G}_t^{(n)} - G_t$ for the additive residual process, which makes the difference \mathcal{F}_t -predictable, and $D_2^{(n)}$ can then be bounded similarly to $R_1^{(n)}$. For a general residual process, it seems possible for $D_2^{(n)}$ to have a bias of order $\sqrt{|J_n|}g(n)$.

Our main result, Theorem 2.4.6, is stated under two assumptions. The second, Assumption 2.4.2, is a straightforward generalization to our setup of similar assumptions in the double machine learning literature on rates of convergence for the two estimators used. Both estimation errors are measured using a 2-norm, and it is plausible that we can relax one norm to a weaker form of convergence if we simultaneously strengthen the other norm. The first assumption, Assumption 2.4.1, requires uniform bounds on both λ and G . This is a strong assumption but perhaps not particularly problematic from a practical viewpoint. Indeed, G is a process we can choose, and we can thus make it bounded if necessary. And though many theoretically interesting counting process models have unbounded intensities, a large cap on the intensity will make no difference in practice. We believe, nevertheless, that it is possible to relax Assumption 2.4.1 to a weaker form of control on the magnitudes of λ and G as functions of time, e.g., moment bounds uniform in θ . However, such a generalization will come at the expense of considerably more technical proofs, and we did not pursue this line of research.

A major practical question is whether we can estimate λ and G with sufficient rates, e.g. $n^{-\frac{1}{4}+\epsilon}$. In Section 2.D in the supplement we give an overview of some known and some conjectured rate results for specific forms of λ and Π . Beyond parametric models we conclude that the existing rate results are scarce, and we regard it as an independent research project to establish rates for general historical regression methods.

Another question is whether we can replace the counting process N by a more general semimartingale. Commenges and Gégout-Petit [2009] define conditional local independence for a class of special semimartingales, and Mogensen et al. [2018] and Mogensen and Hansen [2022] show global Markov properties for local independence graphs of certain Itô processes, which are, in particular, special semimartingales. Thus conditional local independence is well defined beyond counting processes, and we believe that most definitions and results of this paper would generalize beyond N being a counting process. Besides some additional technical challenges, the major practical obstacle with such a generalization is that we cannot realistically assume to have completely observed sample paths of Itô processes, say. The discrete time nature of the observations should then be included in the analysis, and this is beyond the scope of the present paper.

Irrespectively of the remaining open problems, the simulation study demonstrated some important properties of our proposed test, the X-LCT. First, it was fairly simple to implement for the specific example considered using some standard estimation techniques that were not tailored to the specific model class. Second, it had good level and power properties and clearly outperformed the test based on the misspecified marginal Cox model. Third, both Neyman orthogonalization as well as cross-fitting were pivotal for achieving the good properties of the test.

2 Nonparametric conditional local independence testing

Acknowledgements The authors would like to thank the Associate Editor and the anonymous reviewers for constructive comments that lead to substantial improvements of the paper. The work was supported by a grant (NNF20OC0062897) from the Novo Nordisk Foundation.

Supplement The supplementary article [Christgau et al., 2023c] contains proofs of results from the main text, auxiliary results, additional discussions and figures.

Supplement to ‘Nonparametric conditional local independence testing’

The supplementary material is organized as follows: In Section 2.A, we give the proofs of the results of the main text. In Section 2.B, we formulate a general uniform asymptotic theory for metric spaces, whereafter we specialize the theory to the Skorokhod space $D[0, 1]$ and chaining of stochastic processes. In Section 2.C, we state Rebolledo’s martingale central limit theorem, and then we generalize the result to a uniform version that is used in the proofs. In Section 2.D, we discuss estimation of the intensity λ and the residual process G in practice. In particular, we compare known rate results with the rates required in Assumption 2.4.2. In Section 2.E, we compare the LCM estimator with existing work in semiparametric survival models. In Section 2.F, we provide mathematical details regarding Neyman orthogonality. Finally, Section 2.G contains additional details, numerical results, and figures related to the simulation study of Section 2.6.

2.A. Proofs of results in the main text

2.A.1 Proof of Proposition 2.2.5

The process G_t is càglàd and \mathcal{G}_t -predictable by assumption, and the process $I = (I_t)$ is a stochastic integral of G_t w.r.t. a local \mathcal{G}_t -martingale under the hypothesis H_0 . It is thus also a local \mathcal{G}_t -martingale under H_0 . By definition, $I_0 = 0$, and if I is a martingale, $\gamma_t = \mathbb{E}(I_t) = \mathbb{E}(I_0) = 0$. \square

2.A.2 Proof of proposition 2.2.6

Suppose that H is non-negative, càglàd and \mathcal{G}_t -predictable, then since $\int_0^t H_s dM_s$ is a local \mathcal{G}_t -martingale it follows by monotone convergence along a localizing sequence that

$$\mathbb{E} \left(\int_0^t H_s dN_s \right) = \mathbb{E} \left(\int_0^t H_s \lambda_s ds \right) = \int_0^t \mathbb{E}(H_s \lambda_s) ds \quad (2.A.1)$$

for all $t \in [0, 1]$. We can apply the identity above with H the positive and negative part of G , respectively, and the integrability assumption ensures that (2.A.1) also holds with $H = G$. It follows that

$$\gamma_t = \mathbb{E}(I_t) = \mathbb{E} \left(\int_0^t G_s (\lambda_s - \lambda_s) ds \right) = \int_0^t \mathbb{E}(G_s (\lambda_s - \lambda_s)) ds.$$

The latter expectation is indeed a covariance since $\mathbb{E}(G_s) = \mathbb{E}(\mathbb{E}(G_s | \mathcal{F}_{s-})) = 0$. \square

2.A.3 Proof of Lemma 2.4.1

Before proving Lemma 2.4.1, we first state general martingale criteria in the context of counting processes.

2 Nonparametric conditional local independence testing

Lemma 2.A.1. *Let (H_t) be a locally bounded \mathcal{G}_t -predictable process, let N be a counting process with a \mathcal{G}_t -intensity λ_t , and let $M_t = N_t - \int_0^t \lambda_s ds$.*

If $\int_0^1 \lambda_s ds$ (or equivalently N_1) is integrable, then M_t and $M_t^2 - \int_0^t \lambda_s ds$ are each \mathcal{G}_t -martingales. If, in addition, $\int_0^1 H_s^2 \lambda_s ds$ is integrable, then $\int_0^t H_s dM_s$ is a mean zero square integrable martingale.

Proof. The first part is Lemma 2.3.2 and Theorem 2.5.3 in Fleming and Harrington [2011]. For the second part, assume that $\int_0^1 \lambda_s ds$ and $\int_0^1 H_s^2 \lambda_s ds$ are both integrable. In this case, the \mathcal{G}_t -predictable quadratic variation of M is $\langle M \rangle(t) = \int_0^t \lambda_s ds$ by the first part. Then it remains to note that (H_t) is a locally bounded \mathcal{G}_t -predictable process, so the conditions of Theorem 2.4.4 in Fleming and Harrington [2011] are satisfied if $\int_0^1 H_s^2 \lambda_s ds$ is integrable. This establishes the second part. \square

We now return to the proof of Lemma 2.4.1. Let $f \in C(\mathbb{R})$, and we shall prove that $\int_0^t f(G_s) dM_s$ is a mean zero, square integrable \mathcal{G}_t -martingale. The proof for the integral with $f(\widehat{G}_s^{(n)})$ is identical.

Continuity of f implies that $C_f := \sup_{x \in [-C', C']} |f(x)| < \infty$ and that $(f(G_t))$ is a \mathcal{G}_t -predictable process. By Assumption 2.4.1, the process $(f(G_t))$ is almost surely bounded by C_f and therefore

$$\mathbb{E} \left(\int_0^1 f(G_s)^2 \lambda_s ds \right) \leq C_f^2 C < \infty.$$

Thus we can apply Lemma 2.A.1 to conclude that $\int_0^t f(G_s) dM_s$ is a mean zero, square integrable \mathcal{G}_t -martingale. \square

2.A.4 Proof of Proposition 2.4.3

As noted elsewhere, the explicit parametrization of all objects by θ is notationally heavy, and there will thus be an implicit parameter value $\theta \in \Theta$ in most of the subsequent constructions and arguments.

To simplify notation we write

$$U_t^{(n)} = \sum_{j \in J_n} \int_0^t H_{j,s}^{(n)} dM_{j,s}, \quad \text{where} \quad H_{j,s}^{(n)} = \frac{G_{j,s}}{\sqrt{|J_n|}}.$$

We will use a uniform extension of Rebolledo's martingale central limit theorem on the sequence $(U^{(n)})_{n \geq 1}$ to show the result. See Section 2.C for a discussion of Rebolledo's CLT and Theorem 2.C.4 for its uniform extension.

Define $\tilde{\mathcal{G}}_t^n$ be the smallest right continuous and complete filtration generated by the filtrations $\{\mathcal{G}_{j,t} \mid j \in J_n\}$. We can apply Lemma 2.4.1 to each of the terms of $U^{(n)}$ to conclude that the j -th term is a square integrable, mean zero $\mathcal{G}_{j,t}^n$ -martingale. By independence of the observations for each j , we can enlarge the filtration for each term and conclude that they are also square integrable, mean zero $\tilde{\mathcal{G}}_t^n$ -martingales. Thus $U^{(n)}$ is also a square integrable, mean zero $\tilde{\mathcal{G}}_t^n$ -martingale.

To apply Theorem 2.C.4 first establish that the conditions in Equation (2.C.2) are fulfilled. By Proposition 2.C.5, we have that

$$\langle U^{(n)} \rangle(t) = \sum_{j \in J_n} \int_0^t \left(H_{j,s}^{(n)} \right)^2 \lambda_{j,s} ds = \frac{1}{|J_n|} \sum_{j \in J_n} \int_0^t G_{j,s}^2 \lambda_{j,s} ds.$$

By directly applying the bounds from Assumption 2.4.1, we see that the square mean of $\int_0^t G_s^2 \lambda_s ds$ is bounded by $C^2(C')^4$. Thus, for fixed $t \in [0, 1]$, the uniform law of large numbers [Shah and Peters, 2020, Lemma 19] gives that

$$\langle U^{(n)} \rangle(t) = \frac{1}{|J_n|} \sum_{j \in J_n} \int_0^t G_{j,s}^2 \lambda_{j,s} ds \xrightarrow{P/\Theta} \mathbb{E} \left(\int_0^t G_s^2 \lambda_s ds \right) = \mathcal{V}(t)$$

for $n \rightarrow \infty$, since the integrals are i.i.d. with the same distribution as $\int_0^t G_s^2 \lambda_s ds$. This establishes the first part of the condition in Equation (2.C.2). For the second part, we also have from Proposition 2.C.5 that

$$\begin{aligned} \langle U_\epsilon^{(n)} \rangle(t) &= \sum_{j \in J_n} \int_0^t \left(H_{j,s}^{(n)} \right)^2 \mathbb{1} \left(|H_{j,s}^{(n)}| \geq \epsilon \right) d\Lambda_{j,s} \\ &= \frac{1}{|J_n|} \sum_{j \in J_n} \int_0^t G_{j,s}^2 \mathbb{1} \left(|G_{j,s}| \geq \epsilon \sqrt{|J_n|} \right) \lambda_{j,s} ds \end{aligned} \quad (2.A.2)$$

for each $t \in [0, 1]$ and $\epsilon > 0$. From Assumption 2.4.1, we note that for n sufficiently large such that $|J_n| > (C')^2/\epsilon^2$, it holds that $\mathbb{P} \left(|G_{j,s}| \geq \epsilon \sqrt{|J_n|} \right) = 0$ for all $j \in J_n$. As a consequence, the terms in (2.A.2) are almost surely zero for n sufficiently large uniformly over Θ . It follows that $\langle U_\epsilon^{(n)} \rangle(t) \xrightarrow{P/\Theta} 0$, which establishes the second part of (2.C.2).

We finally note that the collection of variance functions, $(\mathcal{V}^\theta)_{\theta \in \Theta}$, is uniformly equicontinuous and bounded above under Assumption 2.4.1. This is established in Lemma 2.A.2 below. We have thus verified all the conditions of Theorem 2.C.4, so we conclude that

$$U^{(n),\theta} \xrightarrow{D/\Theta} U^\theta$$

in $D[0, 1]$ as $n \rightarrow \infty$, where U^θ is a mean zero continuous Gaussian martingale with variance function \mathcal{V}^θ . \square

Note that the convergence of (2.A.2) is established directly from the uniform bounds in Assumption 2.4.1. However, the convergence could also be established under a milder conditions with alternative arguments. For example, under the weaker assumption of uniformly bounded variance functions, dominated convergence can be used to establish L_1 -convergence.

In the proof above we invoked the following lemma, which we will also use in several proofs in the sequel.

2 Nonparametric conditional local independence testing

Lemma 2.A.2. *Under Assumption 2.4.1, the collections $(\gamma^\theta)_{\theta \in \Theta}$ and $(\mathcal{V}^\theta)_{\theta \in \Theta}$ are each uniformly Lipschitz and in particular uniformly equicontinuous. Moreover, it holds almost surely that*

$$\sup_{t \in [0,1]} |\gamma_t| \leq 2CC' \quad \text{and} \quad \mathcal{V}(1) = \mathbb{E} \left(\int_0^1 G_s^2 \lambda_s ds \right) \leq C(C')^2.$$

Proof. For any $0 \leq s < t \leq 1$, a direct application of Assumption 2.4.1 and Proposition 2.4.1 yields

$$|\gamma_t - \gamma_s| \leq \mathbb{E} \left| \int_s^t G_u dN_u - \int_s^t G_u \lambda_u du \right| \leq \mathbb{E} \left(\int_s^t |G_u| (\lambda_u + \lambda_u) du \right) \leq 2CC'(t-s),$$

and similarly,

$$\mathcal{V}(t) - \mathcal{V}(s) = \mathbb{E} \left(\int_s^t G_u^2 \lambda_u du \right) \leq C(C')^2(t-s).$$

This establishes the first part. The bounds follow from inserting $(s, t) = (0, t)$ in the first inequality and $(s, t) = (0, 1)$ in the second inequality. \square

2.A.5 Proof of Proposition 2.4.4

We will divide the proof into three lemmas for each of the remainder terms $R_1^{(n)}$, $R_2^{(n)}$ and $R_3^{(n)}$, where we establish convergence to the zero-process uniformly over t and θ . However, note that the notion of uniform convergence differs for the process index, $t \in [0, 1]$, and the parameter, $\theta \in \Theta$, as we need to show that

$$\forall i \in \{1, 2, 3\} \forall \epsilon > 0 : \quad \lim_{n \rightarrow \infty} \sup_{\theta \in \Theta} \mathbb{P} \left(\sup_{t \in [0,1]} |R_{i,t}^\theta| > \epsilon \right) = 0.$$

For a general discussion of the relation between weak convergence and convergence in probability uniformly as a stochastic process, see Newey [1991]. For a general discussion of uniform stochastic convergence over a distribution parameter, see Section 2.B and the references contained therein. In Section 2.B.3, we discuss the combination of both convergences.

As in the proof of Proposition 2.4.3, $\tilde{\mathcal{G}}_t^n$ denotes the smallest right continuous and complete filtration generated by the filtrations $\{\mathcal{G}_{j,t} \mid j \in J_n\}$. Analogously, we let $\tilde{\mathcal{G}}_t^{n,c}$ be the smallest right continuous and complete filtration generated by the filtrations $\{\mathcal{G}_{j,t} \mid j \in J_n^c\}$. We start by considering $R_3^{(n)}$, since this is the easiest case.

Lemma 2.A.3. *Under Assumption 2.4.2 it holds that $\sup_{t \in [0,1]} |R_{3,t}^{(n)}| \xrightarrow{P/\Theta} 0$.*

Proof. We will show the result by showing that

$$\sup_{\theta \in \Theta} \mathbb{E} \left(\sup_{0 \leq t \leq 1} |R_{3,t}^{(n),\theta}| \right) \rightarrow 0$$

as $n \rightarrow \infty$. Using that the random variables

$$\sup_{0 \leq t \leq 1} |G_{j,t} - \widehat{G}_{j,t}^{(n)}| \cdot \sup_{0 \leq t \leq 1} |\lambda_{j,t} - \widehat{\lambda}_{j,t}^{(n)}|$$

for $j \in J_n$ are identically distributed for each fixed $n \geq 2$, we have that

$$\begin{aligned} & \mathbb{E} \left(\sup_{0 \leq t \leq 1} |R_{3,t}^{(n)}| \right) \\ &= \mathbb{E} \left(\sup_{0 \leq t \leq 1} \left| \frac{1}{\sqrt{|J_n|}} \sum_{j \in J_n} \int_0^t (G_{j,s} - \widehat{G}_{j,s}^{(n)}) (\lambda_{j,s} - \widehat{\lambda}_{j,s}^{(n)}) ds \right| \right) \\ &\leq \frac{1}{\sqrt{|J_n|}} \sum_{j \in J_n} \mathbb{E} \left(\sup_{0 \leq t \leq 1} \int_0^t |G_{j,s} - \widehat{G}_{j,s}^{(n)}| \cdot |\lambda_{j,s} - \widehat{\lambda}_{j,s}^{(n)}| ds \right) \\ &= \sqrt{|J_n|} \mathbb{E} \left(\int_0^1 |G_s - \widehat{G}_s^{(n)}| \cdot |\lambda_s - \widehat{\lambda}_s^{(n)}| ds \right) \\ &\leq \sqrt{|J_n|} \mathbb{E} \left(\sqrt{\int_0^1 (G_s - \widehat{G}_s^{(n)})^2 ds} \sqrt{\int_0^1 (\lambda_s - \widehat{\lambda}_s^{(n)})^2 ds} \right) \\ &\leq \sqrt{|J_n|} \sqrt{\mathbb{E} \left(\int_0^1 (G_s - \widehat{G}_s^{(n)})^2 ds \right)} \sqrt{\mathbb{E} \left(\int_0^1 (\lambda_s - \widehat{\lambda}_s^{(n)})^2 ds \right)} \\ &= \sqrt{|J_n|} g(n) h(n). \end{aligned}$$

By Assumption 2.4.2, $\sqrt{|J_n|} g(n) h(n) \rightarrow 0$ uniformly over Θ as $n \rightarrow \infty$, so the result follows. \square

Next we proceed to the remainder process $R_2^{(n)}$.

Lemma 2.A.4. *Under Assumptions 2.4.1 and 2.4.2, it holds that $\sup_{t \in [0,1]} |R_{2,t}^{(n)}| \xrightarrow{P/\Theta} 0$.*

Proof. We first write

$$R_{2,t}^{(n)} = \frac{1}{\sqrt{|J_n|}} \sum_{j \in J_n} \int_0^t (G_{j,s} - \widehat{G}_{j,s}^{(n)}) dM_{j,s},$$

and note that $R_{2,t}^{(n)}$ is a square integrable, mean zero $\widetilde{\mathcal{G}}_t^n$ -martingale conditionally on $\widetilde{\mathcal{G}}_1^{n,c}$. This follows by applying Lemma 2.4.1 to each of the terms, which are i.i.d. conditionally on $\widetilde{\mathcal{G}}_1^{n,c}$. We conclude that the squared process $(R_{2,t}^{(n)})^2$ is a $\widetilde{\mathcal{G}}_t^n$ -submartingale conditionally on $\widetilde{\mathcal{G}}_1^{n,c}$. By Doob's submartingale inequality we have that

$$\begin{aligned} \mathbb{P} \left(\sup_{0 \leq t \leq 1} |R_{2,t}^{(n)}| \geq \epsilon \right) &= \mathbb{P} \left(\sup_{0 \leq t \leq 1} (R_{2,t}^{(n)})^2 \geq \epsilon^2 \right) \\ &= \mathbb{E} \left(\mathbb{P} \left(\sup_{0 \leq t \leq 1} (R_{2,t}^{(n)})^2 \geq \epsilon^2 \mid \widetilde{\mathcal{G}}_1^{n,c} \right) \right) \\ &\leq \frac{\mathbb{E} \left(\text{Var} \left(R_{2,1}^{(n)} \mid \widetilde{\mathcal{G}}_1^{n,c} \right) \right)}{\epsilon^2} \end{aligned}$$

2 Nonparametric conditional local independence testing

for $\epsilon > 0$. The collection of random variables

$$\left(\int_0^1 (G_{j,s} - \widehat{G}_{j,s}^{(n)}) d\mathbf{M}_{j,s} \right)_{j \in J_n}$$

are i.i.d. conditionally on $\tilde{\mathcal{G}}_1^{n,c}$. Therefore,

$$\begin{aligned} \text{Var} \left(R_{2,1}^{(n)} \mid \tilde{\mathcal{G}}_1^{n,c} \right) &= \frac{1}{|J_n|} \sum_{j \in J_n} \text{Var} \left(\int_0^1 (G_{j,s} - \widehat{G}_{j,s}^{(n)}) d\mathbf{M}_{j,s} \mid \tilde{\mathcal{G}}_1^{n,c} \right) \\ &= \mathbb{E} \left(\int_0^1 (G_s - \widehat{G}_s^{(n)})^2 d\langle \mathbf{M} \rangle_s \mid \tilde{\mathcal{G}}_1^{n,c} \right) \\ &= \mathbb{E} \left(\int_0^1 (G_s - \widehat{G}_s^{(n)})^2 \boldsymbol{\lambda}_s ds \mid \tilde{\mathcal{G}}_1^{n,c} \right) \\ &\leq C \cdot \mathbb{E} \left(\int_0^1 (G_s - \widehat{G}_s^{(n)})^2 ds \mid \tilde{\mathcal{G}}_1^{n,c} \right) \end{aligned}$$

where we have used that $\boldsymbol{\lambda}_t$ is bounded by Assumption 2.4.1 (i). Thus

$$\mathbb{E} \left(\text{Var} \left(R_{2,1}^{(n)} \mid \tilde{\mathcal{G}}_1^{n,c} \right) \right) \leq C \cdot \mathbb{E} \left(\int_0^1 (G_s - \widehat{G}_s^{(n)})^2 ds \right) = C \cdot g(n)^2,$$

and we conclude that

$$\mathbb{P} \left(\sup_{0 \leq t \leq 1} |R_{2,t}^{(n)}| \geq \epsilon \right) \leq \frac{C \cdot g(n)^2}{\epsilon^2} \rightarrow 0,$$

as $n \rightarrow \infty$ uniformly over Θ by Assumption 2.4.2. \square

Before proving that $R_1^{(n)}$ converges weakly to the zero-process, we will need two auxiliary lemmas. The first is a conditional version of Hoeffding's lemma, which lets us conclude conditional sub-Gaussianity. Recall that a mean zero random variable A is sub-Gaussian with variance factor $\nu > 0$ if

$$\log \mathbb{E}(e^{xA}) \leq \frac{x^2 \nu}{2}$$

for all $x \in \mathbb{R}$. See, for example, Boucheron et al. [2013], Lemma 2.2, for the classical unconditional version.

Lemma 2.A.5 (conditional Hoeffding's lemma). *Let Y be a random variable taking values on a bounded interval $[a, b]$, satisfying $\mathbb{E}[Y \mid \mathcal{G}] = 0$ for a σ -algebra \mathcal{G} .*

Then $\log \mathbb{E}(e^{xY} \mid \mathcal{G}) \leq (b-a)^2 x^2 / 8$ almost surely for all $x \in \mathbb{R}$.

Proof. Fix $x \in \mathbb{R}$. By convexity of the exponential function we have

$$e^{xy} \leq \frac{b-y}{b-a} e^{xa} + \frac{y-a}{b-a} e^{xb}, \quad y \in [a, b].$$

Inserting Y in place of y and taking the conditional expectation yields

$$\mathbb{E}[e^{xY} \mid \mathcal{G}] \leq \frac{b}{b-a} e^{xa} - \frac{a}{b-a} e^{xb} = e^{L(x(b-a))}$$

almost surely, where $L(h) = \frac{ha}{b-a} + \log(1 + \frac{a-e^h a}{b-a})$. Standard calculations show that $L(0) = L'(0) = 0$, and the AM-GM inequality implies

$$L''(h) = -\frac{abe^h}{(b-ae^h)^2} \leq \frac{1}{4}.$$

Thus, a second order Taylor expansion yields that $L(h) \leq \frac{1}{8}h^2$, and it follows that $\log \mathbb{E}[e^{xY} \mid \mathcal{G}] \leq \frac{(b-a)^2}{8}x^2$ as desired. \square

For the next lemma define for $s, t \in [0, 1]$ with $s < t$

$$W^{s,t} = \frac{1}{t-s} \int_s^t G_u (\lambda_u - \hat{\lambda}_u^{(n)}) du.$$

Lemma 2.A.6. *Let Assumption 2.4.1 hold true. Then, for any $0 \leq s < t \leq 1$, it holds that $\mathbb{E}(W^{s,t} \mid \tilde{\mathcal{G}}_1^{n,c}) = 0$ and that $W^{s,t}$ is sub-Gaussian conditionally on $\tilde{\mathcal{G}}_1^{n,c}$ with variance factor $\nu = (2CC')^2$, that is,*

$$\log \mathbb{E}(e^{xW^{s,t}} \mid \tilde{\mathcal{G}}_1^{n,c}) \leq 2(xCC')^2$$

for all $s < t$ and $x \in \mathbb{R}$.

Proof. For fixed $u \in [0, 1]$, note that

$$\begin{aligned} \mathbb{E}\left(G_u \left(\lambda_u - \hat{\lambda}_u^{(n)}\right) \mid \tilde{\mathcal{G}}_1^{n,c}\right) &= \mathbb{E}\left(\mathbb{E}\left(G_u \left(\lambda_u - \hat{\lambda}_u^{(n)}\right) \mid \mathcal{F}_{s-} \vee \tilde{\mathcal{G}}_1^{n,c}\right) \mid \tilde{\mathcal{G}}_1^{n,c}\right) \\ &= \mathbb{E}\left(\mathbb{E}\left(G_u \mid \mathcal{F}_{s-}\right) \left(\lambda_u - \hat{\lambda}_u^{(n)}\right) \mid \tilde{\mathcal{G}}_1^{n,c}\right) = 0, \end{aligned} \quad (2.A.3)$$

where we have used that $\lambda_t - \hat{\lambda}_t^{(n)}$ is \mathcal{F}_t -predictable conditionally on $\tilde{\mathcal{G}}_1^{n,c}$, that G_t is independent of $\tilde{\mathcal{G}}_1^{n,c}$ since it is \mathcal{G}_t -predictable, and that $\mathbb{E}(G_s \mid \mathcal{F}_{s-}) = 0$ per definition. By applying the conditional Fubini theorem [Schilling, 2017, Theorem 27.17], we conclude that $\mathbb{E}(W^{s,t} \mid \tilde{\mathcal{G}}_1^{n,c}) = 0$.

We can now use the conditional version of Hoeffding's lemma formulated in Lemma 2.A.5. Indeed, we have that for all $s < t$

$$\begin{aligned} |W^{s,t}| &\leq \frac{1}{t-s} \int_s^t |G_u| |\lambda_u - \hat{\lambda}_u^{(n)}| du \\ &\leq \sup_{0 \leq u \leq 1} |G_u| \sup_{0 \leq u \leq 1} |\lambda_u - \hat{\lambda}_u^{(n)}| \leq 2CC' \end{aligned}$$

by Assumption 2.4.1. Hence, for all $s < t$, Lemma 2.A.5 lets us conclude that

$$\log \mathbb{E}(e^{xW^{s,t}} \mid \tilde{\mathcal{G}}_1^{n,c}) \leq 2(xCC')^2, \quad x \in \mathbb{R}. \quad \square$$

2 Nonparametric conditional local independence testing

Then we have the following regarding $R_1^{(n)}$.

Lemma 2.A.7. *Under Assumptions 2.4.2 and 2.4.1 it holds that $\sup_{t \in [0,1]} |R_{1,t}^{(n)}| \xrightarrow{P/\Theta} 0$.*

Proof. The proof consists of two parts. First we show that for each $t \in [0, 1]$ it holds that

$$R_{1,t}^{(n)} \xrightarrow{P/\Theta} 0$$

for $n \rightarrow \infty$. Then we show *stochastic equicontinuity* of the process $R_1^{(n)}$ uniformly over Θ , and by Lemma 2.B.16 it follows that

$$\sup_{t \in [0,1]} |R_{1,t}^{(n)}| \xrightarrow{P/\Theta} 0.$$

This is a direct generalization of Theorem 2.1 in Newey [1991]. The collection of random variables

$$\left(G_{j,s} \left(\lambda_{j,s} - \hat{\lambda}_{j,s}^{(n)} \right) \right)_{j \in J_n}$$

are i.i.d. conditionally on $\tilde{\mathcal{G}}_1^{n,c}$. Therefore, an application of the conditional Fubini theorem yields

$$\mathbb{E}(R_{1,t} | \tilde{\mathcal{G}}_1^{n,c}) = \frac{1}{\sqrt{|J_n|}} \sum_{j \in J_n} \int_0^t \mathbb{E} \left(G_{j,s} \left(\lambda_{j,s} - \hat{\lambda}_{j,s}^{(n)} \right) | \tilde{\mathcal{G}}_1^{n,c} \right) ds = 0$$

where the last equality follows from the computation in (2.A.3). Whence $\mathbb{E}(R_{1,t}^{(n)}) = 0$, and $\text{Var}(R_{1,t}^{(n)}) = \mathbb{E}(\text{Var}(R_{1,t}^{(n)} | \tilde{\mathcal{G}}_1^{n,c}))$, so

$$\begin{aligned} \text{Var}(R_{1,t}^{(n)}) &= \mathbb{E} \left(\frac{1}{|J_n|} \sum_{j \in J_n} \text{Var} \left(\int_0^t G_{j,s} \left(\lambda_{j,s} - \hat{\lambda}_{j,s}^{(n)} \right) ds | \tilde{\mathcal{G}}_1^{n,c} \right) \right) \\ &= \mathbb{E} \left(\mathbb{E} \left(\left(\int_0^t G_s \left(\lambda_s - \hat{\lambda}_s^{(n)} \right) ds \right)^2 | \tilde{\mathcal{G}}_1^{n,c} \right) \right) \\ &= \mathbb{E} \left(\left(\int_0^t G_s \left(\lambda_s - \hat{\lambda}_s^{(n)} \right) ds \right)^2 \right) \\ &\leq (C')^2 \mathbb{E} \left(\int_0^t \left(\lambda_s - \hat{\lambda}_s^{(n)} \right)^2 ds \right) \\ &\leq (C')^2 h(n)^2 \end{aligned}$$

where we have used Assumption 2.4.1 (ii). Hence by Chebychev's inequality, it holds for all $\epsilon > 0$ that

$$\mathbb{P}(|R_{1,t}^{(n)}| > \epsilon) \leq \frac{\text{Var}(R_{1,t}^{(n)})}{\epsilon^2} \leq \frac{(C')^2 h(n)^2}{\epsilon^2} \rightarrow 0$$

as $n \rightarrow \infty$ uniformly over Θ by Assumption 2.4.2. This completes the first part of the proof. For the second part, we use a chaining argument based on the exponential inequality in Lemma 2.A.6. We let

$$W_j^{s,t} = \frac{1}{t-s} \int_s^t G_{j,u} (\lambda_{j,u} - \hat{\lambda}_{j,u}^{(n)}) du$$

and

$$A = \frac{1}{\sqrt{|J_n|}} \sum_{j \in J_n} W_j^{s,t} = \frac{1}{t-s} (R_{1,t}^{(n)} - R_{1,s}^{(n)}).$$

Using that $(W_j^{s,t})_{j \in J_n}$ are i.i.d. conditionally on $\tilde{\mathcal{G}}_1^{n,c}$ we have by Lemma 2.A.6 that $\mathbb{E}(A) = 0$ and that

$$\begin{aligned} \log \mathbb{E}(e^{xA}) &= \log \mathbb{E} \left(\mathbb{E} \left(e^{xA} \mid \tilde{\mathcal{G}}_1^{n,c} \right) \right) \\ &= \log \mathbb{E} \left(\prod_{j \in J_n} \mathbb{E} \left(e^{\frac{x}{\sqrt{|J_n|}} W_j^{s,t}} \mid \tilde{\mathcal{G}}_1^{n,c} \right) \right) \\ &\leq \log \mathbb{E} \left(e^{\frac{x^2 \nu}{2}} \right) \\ &= \frac{x^2 \nu}{2}. \end{aligned}$$

Hence A is also sub-Gaussian with variance factor ν . This implies that

$$\mathbb{P}(|A| > \eta) \leq 2e^{-\frac{\eta^2 \nu}{2}}$$

for all $\eta > 0$. Rephrased in terms of $R_1^{(n)}$ this bound reads

$$\mathbb{P} \left(|R_{1,t}^{(n)} - R_{1,s}^{(n)}| > \eta(t-s) \right) \leq 2e^{-\frac{\eta^2 \nu}{2}}$$

for all $\eta > 0$ and $s < t$. It now follows from the chaining lemma, Pollard [1984] Lemma VII.9, that $R_1^{(n)}$ is stochastic equicontinuous. Since the variance factor $\nu = (2CC')^2$ does not depend on $\theta \in \Theta$, we have stochastic equicontinuity uniformly over Θ by Corollary 2.B.19. This completes the second part of the proof and we are done. \square

Note that the second part of the proof above establishes stochastic equicontinuity by a bound on the probability that the increments of the process are large. This is a well known technique, see, e.g., Example 2.2.12 in van der Vaart and Wellner [1996], from which the same conclusion will follow if

$$\mathbb{E}(|R_{1,t}^{(n)} - R_{1,s}^{(n)}|^p) \leq K|t-s|^{1+r}$$

for $K, p, r > 0$.

Proposition 2.4.4 now follows from combining the Lemmas 2.A.7, 2.A.4, and 2.A.3. \square

2.A.6 Proof of Proposition 2.4.5

We separate the discussion of $D_1^{(n)}$ and $D_2^{(n)}$ into the Lemmas 2.A.8 and 2.A.10, respectively, which together amount to Proposition 2.4.5.

Lemma 2.A.8. *Suppose that Assumptions 2.4.1 and 2.4.2 hold. Then the stochastic process $\bar{D}^{(n)} := D_1^{(n)} - \sqrt{|J_n|} \cdot \gamma$ converges in distribution in $C[0, 1]$ uniformly over Θ .*

Proof. Let $\bar{D}^{(n)} := D_1^{(n)} - \sqrt{|J_n|} \cdot \gamma$ and note that

$$\bar{D}^{(n)} = |J_n|^{-\frac{1}{2}} \sum_{j \in J_n} W_j,$$

where W_j is given by $W_{j,t} := \int_0^t G_{j,s}(\boldsymbol{\lambda}_{j,s} - \lambda_{j,s})ds - \gamma_t$ for each $j \in J_n$. By assumption, the variables $\{W_j : j \in J_n\}$ are i.i.d. with the same distribution as the process W given by $W_t := \int_0^t G_s(\boldsymbol{\lambda}_s - \lambda_s)ds - \gamma_t$. For each $\theta \in \Theta$, let Γ^θ be a Gaussian process with mean zero and covariance function $(s, t) \mapsto \text{Cov}(W_s^\theta, W_t^\theta)$, which is well-defined by computations shown below.

We will show that $\bar{D}^{(n),\theta} \xrightarrow{\mathcal{D}/\Theta} \Gamma^\theta$ in $C[0, 1]$ by applying Lemma 2.B.14, which is an example of Prokhorov's method of "tightness + identification of limit". We first prove that for any given $k \in \mathbb{N}$ and $0 \leq t_1 < t_2 < \dots < t_k \leq 1$,

$$\mathbf{D}^{(n)} := (\bar{D}_{t_1}^{(n)}, \bar{D}_{t_2}^{(n)}, \dots, \bar{D}_{t_k}^{(n)}) \xrightarrow{\mathcal{D}/\Theta} (\Gamma_{t_1}^\theta, \Gamma_{t_2}^\theta, \dots, \Gamma_{t_k}^\theta).$$

To this end we will apply the uniform CLT of Lundborg et al. [2022b, Proposition 19] to the sequence of random vectors $\mathbf{D}^{(n)} \in \mathbb{R}^k$, i.e., the sequence of normalized sums of i.i.d. copies of $\mathbf{W} := (W_{t_1}, \dots, W_{t_k})$. The process (W_t) is mean zero and hence \mathbf{W} is also mean zero. For any $t \in [0, 1]$ we observe that

$$\text{Var}(W_t) = \text{Var}(W_t + \gamma_t) \leq \mathbb{E} \left[\left(\int_0^t |G_s| \cdot |\boldsymbol{\lambda}_s - \lambda_s| ds \right)^2 \right] \leq 2C^2(C')^2.$$

Therefore the trace of $\text{Var}(\mathbf{W})$ is uniformly bounded, which implies the trace condition in Proposition 19 of Lundborg et al. [2022b]. From Hölder's inequality and Minkowski's inequality, we note that for any $\mathbf{a}, \mathbf{b} \in \mathbb{R}^k$

$$\|\mathbf{a} + \mathbf{b}\|_2^3 \leq k^{3/2} \|\mathbf{a} + \mathbf{b}\|_3^3 \leq k^{3/2} (\|\mathbf{a}\|_3 + \|\mathbf{b}\|_3)^3 \leq 8k^{3/2} (\|\mathbf{a}\|_3^3 + \|\mathbf{b}\|_3^3).$$

Combining the above with Assumption 2.4.1 and Lemma 2.A.2 yields that

$$\mathbb{E}[\|\mathbf{W}\|_2^3] \leq C_k \mathbb{E} \left[\left(\int_0^1 |G_s| \cdot |\boldsymbol{\lambda}_s - \lambda_s| ds \right)^3 \right] + C_k \sup_{t \in [0,1]} |\gamma_t|^3 \leq 16C_k C^3 (C')^3,$$

where $C_k = 8k^{5/2}$. Hence Proposition 19 of Lundborg et al. [2022b] lets us conclude that $\mathbf{D}^{(n)} \xrightarrow{\mathcal{D}/\Theta} \mathcal{N}(0, \text{Var}(\mathbf{W}))$. By definition of Γ^θ , this is equivalent to $\mathbf{D}^{(n)} \xrightarrow{\mathcal{D}/\Theta} (\Gamma_{t_1}^\theta, \Gamma_{t_2}^\theta, \dots, \Gamma_{t_k}^\theta)$.

We now argue that $(\overline{D}^{(n)})$ and (Γ^θ) are stochastically equicontinuous uniformly over Θ . From the definition of Γ^θ and by Assumption 2.4.1, it follows that

$$\mathbb{E}[(\Gamma_t^\theta - \Gamma_s^\theta)^2] = \mathbb{E}[(W_t - W_s)^2] \leq (2CC'(t-s))^2. \quad (2.A.4)$$

Hence $\frac{1}{t-s}(\Gamma_t^\theta - \Gamma_s^\theta)$ is Gaussian with a variance bounded over Θ and $0 \leq s < t \leq 1$. In particular, it is sub-Gaussian with a uniform variance factor over Θ and $0 \leq s < t \leq 1$. Since W is uniformly bounded over Θ , an application of Hoeffding's Lemma yields that $A_j^{t,s} := \frac{1}{t-s}(W_{j,t} - W_{j,s})$ is also sub-Gaussian with a variance factor ν that is uniform over Θ , $0 \leq s < t \leq 1$, and $j \in J_n$. Letting $A_\bullet^{s,t} = \frac{1}{t-s}(\overline{D}_t^{(n)} - \overline{D}_s^{(n)})$, we have

$$\mathbb{E}e^{xA_\bullet^{s,t}} = \prod_{j \in J_n} \mathbb{E}[e^{x|J_n|^{-1/2}A_j^{s,t}}] \leq \prod_{j \in J_n} e^{\frac{x^2\nu}{2|J_n|}} = e^{x^2\nu/2}.$$

Hence $A_\bullet^{s,t}$ is also sub-Gaussian with a variance factor uniformly over Θ and $0 \leq s < t \leq 1$.

From the uniform chaining lemma, Corollary 2.B.19, we now conclude that both (Γ^θ) and $(\overline{D}^{(n)})$ are stochastically equicontinuous uniformly over Θ . By Proposition 2.B.20, this means that the collection $(\overline{D}^{(n),\theta})$ is sequentially tight and that (Γ^θ) , which is constant in n , is uniformly tight.

Now we have shown convergence of the finite-dimensional marginals and appropriate tightness conditions, so Lemma 2.B.14 lets us conclude that $\overline{D}^{(n)} \xrightarrow{\mathcal{D}/\Theta} \Gamma^\theta$ weakly in $C[0, 1]$. \square

Before moving on to the term $D_2^{(n)}$, we first note that Lemma 2.A.8 implies that stochastic boundedness, as we will use this result in the proof of Theorem 2.4.6.

Lemma 2.A.9. *Suppose that Assumptions 2.4.1 and 2.4.2 hold. Then $\overline{D}^{(n)} := D_1^{(n)} - \sqrt{|J_n|} \cdot \gamma$ is stochastically bounded uniformly over Θ , i.e., for every $\varepsilon > 0$ there exists $K > 0$ such that*

$$\limsup_{n \rightarrow \infty} \sup_{\theta \in \Theta} \mathbb{P}\left(\|\overline{D}^{(n),\theta}\|_\infty > K\right) < \varepsilon.$$

Proof. We have established in the proof of Lemma 2.A.8, under the same conditions, that $\overline{D}^{(n),\theta} \xrightarrow{\mathcal{D}/\Theta} \Gamma^\theta$ weakly in $C[0, 1]$. By the uniform continuous mapping theorem formulated in Proposition 2.B.6, it follows that $\|\overline{D}^{(n),\theta}\|_\infty \xrightarrow{\mathcal{D}/\Theta} \|\Gamma^\theta\|_\infty$. From Bengs and Holzmann [2019] Theorem 4.1 we then obtain that

$$\limsup_{n \rightarrow \infty} \sup_{\theta \in \Theta} \mathbb{P}(\|\overline{D}^{(n),\theta}\|_\infty > K) \leq \sup_{\theta \in \Theta} \mathbb{P}(\|\Gamma^\theta\|_\infty > K) \leq \frac{\mathbb{E}\|\Gamma^\theta\|_\infty}{K}.$$

Hence it suffices to argue that $\mathbb{E}\|\Gamma^\theta\|_\infty$ is uniformly bounded over Θ . To this end, we note that Equation (2.A.4) shows that square means of the increments of Γ^θ are smaller than those of a standard Brownian motion scaled by $2CC'$. Then the Sudakov–Fernique

2 Nonparametric conditional local independence testing

comparison inequality [Adler and Taylor, 2007, Theorem 2.2.3] allows us to leverage this relationship to the expected uniform norms, i.e., $\mathbb{E}\|\Gamma^\theta\|_\infty \leq 2CC'\mathbb{E}(\sup_{t \in [0,1]} |B_t|)$. It can be verified that $\mathbb{E}(\sup_{t \in [0,1]} |B_t|)$ is finite, and in fact, equal to $\sqrt{\pi/2}$ as shown in saz [2019]. \square

Lemma 2.A.10. *Suppose that Assumptions 2.4.1 and 2.4.2 hold, and that $G_t = X_t - \Pi_t$ is the additive residual process. Then $D_2^{(n)} \xrightarrow{P/\Theta} 0$ in $D[0,1]$ as $n \rightarrow \infty$.*

Proof. Note first that the terms in $D_2^{(n)}$ are i.i.d. conditionally on $\tilde{\mathcal{G}}_1^{n,c}$, with the same distribution as the process ξ given by

$$\xi_t = \frac{1}{\sqrt{|J_n|}} \int_0^t (\hat{G}_s^{(n)} - G_s)(\lambda_s - \lambda_s) ds.$$

Since λ_t is independent of $\tilde{\mathcal{G}}_1^{n,c}$, we have from the innovation theorem that

$$\mathbb{E}(\lambda_t | \mathcal{F}_{t-} \vee \tilde{\mathcal{G}}_1^{n,c}) = \mathbb{E}(\lambda_t | \mathcal{F}_{t-}) = \lambda_t.$$

For the additive residual process we also note that $G_t - \hat{G}_t^{(n)} = \hat{\Pi}_t^{(n)} - \Pi_t$ is \mathcal{F}_t -predictable conditionally on $\tilde{\mathcal{G}}_1^{n,c}$. It now follows that

$$\begin{aligned} \sqrt{|J_n|} \cdot \mathbb{E}[\xi_t | \tilde{\mathcal{G}}_1^{n,c}] &= \int_0^t \mathbb{E}[(\hat{G}_s^{(n)} - G_s)(\lambda_s - \lambda_s) | \tilde{\mathcal{G}}_1^{n,c}] ds \\ &= \int_0^t \mathbb{E}[(\hat{G}_s^{(n)} - G_s)(\mathbb{E}[\lambda_s | \mathcal{F}_{s-} \vee \tilde{\mathcal{G}}_1^{n,c}] - \lambda_s) | \tilde{\mathcal{G}}_1^{n,c}] ds = 0. \end{aligned}$$

We can therefore conclude that $D_2^{(n)}$ is mean zero conditionally on $\tilde{\mathcal{G}}_1^{n,c}$. Using that the terms of $D_2^{(n)}$ are i.i.d. conditionally on $\tilde{\mathcal{G}}_1^{n,c}$ once more, we now obtain that

$$\begin{aligned} \text{Var}(D_{2,t}^{(n)} | \tilde{\mathcal{G}}_1^{n,c}) &= |J_n| \cdot \text{Var}(\xi_t | \tilde{\mathcal{G}}_1^{n,c}) = \mathbb{E} \left[\left(\int_0^t (\hat{G}_s^{(n)} - G_s)(\lambda_s - \lambda_s) ds \right)^2 \middle| \tilde{\mathcal{G}}_1^{n,c} \right] \\ &\leq 4C^2 \cdot \mathbb{E} \left(\int_0^1 (\hat{G}_s^{(n)} - G_s)^2 ds \middle| \tilde{\mathcal{G}}_1^{n,c} \right). \end{aligned}$$

Taking expectation of the above we have $\text{Var}(D_{2,t}^{(n)}) = \mathbb{E}(\text{Var}(D_{2,t}^{(n)} | \tilde{\mathcal{G}}_1^{n,c})) \leq 4C^2 g(n)^2$. By Chebyshev's inequality we get for all $\epsilon > 0$

$$\mathbb{P} \left(|D_{2,t}^{(n)}| > \epsilon \right) \leq \frac{4C^2 g(n)^2}{\epsilon^2},$$

and by Assumption 2.4.2 we conclude that $D_{2,t}^{(n)} \xrightarrow{P/\Theta} 0$ for each $t \in [0,1]$.

We now apply the same chaining argument used in the proofs of Lemma 2.A.7 and Lemma 2.A.8. From Assumption 2.4.1, we have for $0 \leq s < t \leq 1$ that $|\xi_t - \xi_s| \leq 4\sqrt{|J_n|}CC'(t-s)$. Hence the conditional Hoeffding's lemma (Lemma 2.A.5) yields that

$$A_j^{s,t} = \frac{1}{t-s} \int_s^t (\hat{G}_{j,s}^{(n)} - G_{j,s})(\lambda_{j,s} - \lambda_{j,s}) ds$$

is sub-Gaussian conditionally on $\tilde{\mathcal{G}}_1^{n,c}$ with a variance factor ν that is uniform over Θ and $s < t$ (cf. the proof of Lemma 2.A.6). Letting $A_{\bullet}^{s,t} = \frac{1}{t-s}(D_{2,t}^{(n)} - D_{2,s}^{(n)})$, we have for any $x \in \mathbb{R}$

$$\begin{aligned} \mathbb{E}\left(e^{xA_{\bullet}^{s,t}}\right) &= \mathbb{E}\left(\mathbb{E}\left[e^{xA_{\bullet}^{s,t}} \mid \tilde{\mathcal{G}}_1^{n,c}\right]\right) \\ &= \mathbb{E}\left(\prod_{j \in J_n} \mathbb{E}\left[e^{x|J_n|^{-1/2}A_j^{s,t}} \mid \tilde{\mathcal{G}}_1^{n,c}\right]\right) \leq \prod_{j \in J_n} e^{\frac{x^2\nu}{2|J_n|}} = e^{x^2\nu/2}, \end{aligned}$$

so $A_{\bullet}^{s,t}$ is also sub-Gaussian uniformly over $s < t$ and Θ . In terms of $D_2^{(n)}$, this means that we can apply the uniform chaining lemma, Corollary 2.B.19, and conclude that it is stochastically equicontinuous uniformly over Θ .

Since $D_{2,t}^{(n)} \xrightarrow{P/\Theta} 0$ for each $t \in [0, 1]$ and $(D_2^{(n)})$ is stochastically equicontinuous uniformly over Θ , Lemma 2.B.16 now lets us conclude that $\sup_{t \in [0,1]} |D_{2,t}^{(n)}| \xrightarrow{P/\Theta} 0$ and we are done. \square

2.A.7 Proof of Theorem 2.4.6

Before proving Theorem 2.4.6, we first prove that the collection of Gaussian martingales from Proposition 2.4.3 is tight in $C[0, 1]$ (see Definition 2.B.7).

Lemma 2.A.11. *Let $(U^\theta)_{\theta \in \Theta}$ be the collection of Gaussian martingales from Proposition 2.4.3, i.e., U^θ is a mean zero continuous Gaussian martingale with variance function \mathcal{V}^θ . Under Assumption 2.4.1, $(U^\theta)_{\theta \in \Theta}$ is uniformly tight in $C[0, 1]$.*

Proof. We will use Theorem 7.3 in Billingsley [2013], which characterizes tightness of measures in $C[0, 1]$. The first condition of the theorem is trivially satisfied for $(U^\theta)_{\theta \in \Theta}$ since $\mathbb{P}(U_0^\theta = 0) = 1$ for all $\theta \in \Theta$.

By Proposition 2.C.2, U^θ has a distributional representation as a time-transformed Brownian motion such that $(U_t^\theta)_{t \in [0,1]} \stackrel{\mathcal{D}}{=} (B_{\mathcal{V}^\theta(t)})_{t \in [0,1]}$, where B is a Brownian motion. Recall that Brownian motion is α -Hölder continuous for $\alpha \in (0, \frac{1}{2})$, which means that

$$K(\alpha) = \sup_{s \neq t} \frac{|B_t - B_s|}{|t - s|^\alpha} < \infty.$$

Note also that the collection of variance functions is uniformly Lipschitz by Lemma 2.A.2 with uniform Lipschitz constant C_0 , say. It follows that for every $\epsilon > 0$,

$$\begin{aligned} \lim_{\delta \rightarrow 0^+} \sup_{\theta \in \Theta} \mathbb{P}\left(\sup_{|t-s| < \delta} |U_t^\theta - U_s^\theta| > \epsilon\right) &= \lim_{\delta \rightarrow 0^+} \sup_{\theta \in \Theta} \mathbb{P}\left(\sup_{|t-s| < \delta} |B_{\mathcal{V}^\theta(t)} - B_{\mathcal{V}^\theta(s)}| > \epsilon\right) \\ &\leq \lim_{\delta \rightarrow 0^+} \sup_{\theta \in \Theta} \mathbb{P}\left(K(\alpha) \sup_{|t-s| < \delta} |\mathcal{V}^\theta(t) - \mathcal{V}^\theta(s)|^\alpha > \epsilon\right) \\ &= \lim_{\delta \rightarrow 0^+} \mathbb{P}\left(K(\alpha)C_0^\alpha \delta^\alpha > \epsilon\right) = 0. \end{aligned}$$

This establishes the second condition of Theorem 7.3 in Billingsley [2013], and we thus conclude that $(U^\theta)_{\theta \in \Theta}$ is uniformly tight in $C[0, 1]$. \square

2 Nonparametric conditional local independence testing

We now return to the proof of Theorem 2.4.6.

For part i), we first note that under H_0 we can take $\lambda_t = \lambda_t$, which implies that both $D_1^{(n)}$ and $D_2^{(n)}$ equal the zero-process.

Combining Propositions 2.4.3 and 2.4.4 with the uniform version of Slutsky's theorem formulated in Lemma 2.B.5, we conclude that

$$\sqrt{|J_n|}\hat{\gamma}^{(n)} = \underbrace{U^{(n)}}_{\xrightarrow{P/\Theta_0} U^\theta} + \underbrace{R_1^{(n)} + R_2^{(n)} + R_3^{(n)}}_{\xrightarrow{P/\Theta} 0} + \underbrace{D_1^{(n)} + D_2^{(n)}}_{=0 \text{ under } H_0} \xrightarrow{P/\Theta_0} U^\theta,$$

in $D[0, 1]$ as $n \rightarrow \infty$, where U^θ is the Gaussian martingale from Proposition 2.4.3.

For part ii) we can, in addition to Propositions 2.4.3 and 2.4.4, apply Proposition 2.4.5 and Lemma 2.A.9. Using the triangle inequality on the decomposition (2.4.17) yields that

$$\begin{aligned} \sqrt{|J_n|} \cdot \|\hat{\gamma}^{(n)} - \gamma\|_\infty &\leq \|U^{(n)}\|_\infty + \|D_1^{(n)} - \sqrt{|J_n|}\gamma\|_\infty \\ &\quad + \|R_1^{(n)}\|_\infty + \|R_2^{(n)}\|_\infty + \|R_3^{(n)}\|_\infty + \|D_2^{(n)}\|_\infty. \end{aligned}$$

All the terms in the second line converge in probability to zero uniformly over Θ . Combined with the convergences established in Proposition 2.4.3 and Lemma 2.A.8, we obtain that

$$\begin{aligned} &\limsup_{n \rightarrow \infty} \sup_{\theta \in \Theta} \mathbb{P} \left(\sqrt{|J_n|} \cdot \|\hat{\gamma}^{(n), \theta} - \gamma^\theta\|_\infty > K \right) \\ &\leq \sup_{\theta \in \Theta} \mathbb{P} \left(\|U^\theta\|_\infty > K/6 \right) + \sup_{\theta \in \Theta} \mathbb{P} \left(\|\Gamma^\theta\|_\infty > K/6 \right), \end{aligned} \quad (2.A.5)$$

where Γ^θ is the limiting Gaussian process from (the proof of) Lemma 2.A.8. The last term in (2.A.5) can be made arbitrarily small for K sufficiently large by Lemma 2.A.9. Lemma 2.A.11 states that the family $(U^\theta)_{\theta \in \Theta}$ is tight in $C[0, 1]$, and hence the family $(\|U^\theta\|_\infty)_{\theta \in \Theta}$ is tight in $\mathbb{R}_{\geq 0}$. This implies that the first term in (2.A.5) can also be made arbitrarily small for K sufficiently large. This establishes (2.4.26) and we are done. \square

2.A.8 Proof of Proposition 2.4.7

Consider the decomposition of the variance function estimator given by

$$\hat{\mathcal{V}}_n(t) = A_t^{(n)} + B_t^{(n)} + 2C_t^{(n)}$$

where

$$\begin{aligned} A_t^{(n)} &= \frac{1}{|J_n|} \sum_{j \in J_n} \int_0^t G_{j,s}^2 dN_{j,s}, \\ B_t^{(n)} &= \frac{1}{|J_n|} \sum_{j \in J_n} \int_0^t \left(G_{j,s} - \hat{G}_{j,s}^{(n)} \right)^2 dN_{j,s}, \\ C_t^{(n)} &= \frac{1}{|J_n|} \sum_{j \in J_n} \int_0^t G_{j,s} \left(G_{j,s} - \hat{G}_{j,s}^{(n)} \right) dN_{j,s}. \end{aligned}$$

We first consider the asymptotic limit of $A^{(n)}$, which is the empirical mean of $|J_n|$ i.i.d. samples of the process $\int_0^t G_s^2 dN_s$. Under Assumption 2.4.1, we can apply the first part of Lemma 2.A.1 which states $\mathbf{M}_t^2 - \mathbf{\Lambda}_t$ is a martingale. We use this fact to note that

$$\mathbb{E}(N_1^2) = \mathbb{E}((\mathbf{M}_1 + \mathbf{\Lambda}_1)^2) \leq 2(\mathbb{E}(\mathbf{M}_1^2) + \mathbb{E}(\mathbf{\Lambda}_1^2)) = 4\mathbb{E}\left(\left(\int_0^1 \lambda_s ds\right)^2\right) \leq 4C^2.$$

Now, another use of Assumption 2.4.1 shows that $\int_0^t G_s^2 dN_s$ has a second moment bounded by $4(CC')^2$. Thus we can apply the uniform law of large numbers [Shah and Peters, 2020, Lemma 19] to conclude for each $t \in [0, 1]$,

$$A_t^{(n)} = \frac{1}{|J_n|} \sum_{j \in J_n} \int_0^t G_{j,s}^2 dN_{j,s} \xrightarrow{P/\Theta} \mathbb{E}\left(\int_0^t G_s^2 dN_s\right) = \mathcal{V}(t).$$

Note also that $A^{(n)}$ and \mathcal{V} are non-decreasing and that the collection $(\mathcal{V}^\theta)_{\theta \in \Theta}$ is uniformly equicontinuous by Lemma 2.A.2. These are exactly the conditions for Lemma 2.B.13, so we can automatically conclude that $\sup_{t \in [0,1]} |A_t^{(n)} - \mathcal{V}(t)| \xrightarrow{P/\Theta} 0$.

Next we show that the remainder terms $B^{(n)}$ and $C^{(n)}$ converge uniformly to zero in expectation. Similarly to the proof of Lemma 2.A.4, we have under Assumptions 2.4.1 and 2.4.2,

$$\begin{aligned} \mathbb{E}\left(\sup_{0 \leq t \leq 1} B_t^{(n)}\right) &= \mathbb{E}(B_1^{(n)}) = \mathbb{E}\left(\frac{1}{|J_n|} \sum_{j \in J_n} \int_0^1 (G_{j,s} - \hat{G}_{j,s}^{(n)})^2 \lambda_{j,s} ds\right) \\ &= \mathbb{E}\left(\mathbb{E}\left(\frac{1}{|J_n|} \sum_{j \in J_n} \int_0^1 (G_{j,s} - \hat{G}_{j,s}^{(n)})^2 \lambda_{j,s} ds \mid \tilde{G}_1^c\right)\right) \\ &= \mathbb{E}\left(\int_0^1 (G_s - \hat{G}_s^{(n)})^2 \lambda_s ds\right) \\ &\leq C \cdot g(n)^2 \longrightarrow 0 \end{aligned}$$

2 Nonparametric conditional local independence testing

as $n \rightarrow \infty$ uniformly over Θ . Lastly, we see that

$$\begin{aligned}
\mathbb{E} \left| \sup_{0 \leq t \leq 1} C_t^{(n)} \right| &\leq \mathbb{E} \left(\sup_{0 \leq t \leq 1} |C_t^{(n)}| \right) \\
&\leq \mathbb{E} \left(\frac{1}{|J_n|} \sum_{j \in J_n} \sup_{0 \leq t \leq 1} \int_0^t |G_{j,s}| |G_{j,s} - \widehat{G}_{j,s}^{(n)}| \lambda_{j,s} ds \right) \\
&= \mathbb{E} \left(\frac{1}{|J_n|} \sum_{j \in J_n} \int_0^1 |G_{j,s}| |G_{j,s} - \widehat{G}_{j,s}^{(n)}| \lambda_{j,s} ds \right) \\
&= \mathbb{E} \left(\mathbb{E} \left(\frac{1}{|J_n|} \sum_{j \in J_n} \int_0^1 |G_{j,s}| |G_{j,s} - \widehat{G}_{j,s}^{(n)}| \lambda_{j,s} ds \mid \tilde{\mathcal{G}}_1^c \right) \right) \\
&= \mathbb{E} \left(\int_0^1 |G_s| |G_s - \widehat{G}_s^{(n)}| \lambda_s ds \right) \\
&\leq CC' \mathbb{E} \left(\int_0^1 |G_s - \widehat{G}_s^{(n)}| ds \right) \\
&\leq CC' \cdot g(n) \rightarrow 0
\end{aligned}$$

as $n \rightarrow \infty$ uniformly over Θ by Assumption 2.4.2. Combining the convergences established for $A^{(n)}$, $B^{(n)}$, and $C^{(n)}$, we get by a generalized Slutsky (Lemma 2.B.11) that

$$\sup_{t \in [0,1]} |\widehat{\mathcal{V}}_n(t) - \mathcal{V}(t)| \xrightarrow{P/\Theta} 0.$$

□

2.A.9 Proof of Corollary 2.4.8

Under Assumptions 2.4.1 and 2.4.2 we know by Theorem 2.4.6 and Proposition 2.4.7 that

$$\sqrt{|J_n|} \widehat{\gamma}^{(n),\theta} \xrightarrow{D/\Theta_0} U^\theta \quad \text{and} \quad \widehat{\mathcal{V}}_n^\theta \xrightarrow{P/\Theta_0} \mathcal{V}^\theta \tag{2.A.6}$$

in $D[0,1]$ as $n \rightarrow \infty$. If we were to show pointwise convergence of the test statistic, this would now be a straightforward consequence of the continuous mapping theorem. However, to show uniform convergence, we will need an additional tightness argument.

Let $(\theta_n)_{n \in \mathbb{N}} \subset \Theta_0$ be an arbitrary sequence. Proposition 2.B.3 then states that it suffices to show that there exists a subsequence $(\theta_{k(n)})_{n \in \mathbb{N}} \subseteq (\theta_n)_{n \in \mathbb{N}}$, with $k: \mathbb{N} \rightarrow \mathbb{N}$ strictly increasing, such that

$$\lim_{n \rightarrow \infty} d_{BL}(\widehat{D}_{k(n)}^{\theta_{k(n)}}, \mathcal{J}(U^{\theta_{k(n)}}, \mathcal{V}^{\theta_{k(n)}})) = 0. \tag{2.A.7}$$

Here d_{BL} denotes the bounded Lipschitz metric defined in Section 2.B. By Lemma 2.A.11, the collection $(U^\theta)_{\theta \in \Theta}$ is tight in $C[0,1]$ under Assumption 2.4.1. Therefore, Prokhorov's

theorem [Kallenberg, 2021, Theorem 23.2] asserts that there exists a subsequence $(\theta_{a(n)}) \subset (\theta_n)$, and a $C[0, 1]$ -valued random variable \tilde{U} such that $U^{\theta_{a(n)}} \xrightarrow{\mathcal{D}} \tilde{U}$ in $C[0, 1]$.

Likewise, Lemma 2.A.2 states that the collection $(V^\theta)_{\theta \in \Theta}$ is uniformly bounded and uniformly equicontinuous under Assumption 2.4.1. Thus the Arzelà-Ascoli theorem yields that there exists a further subsequence $(\theta_{b(n)}) \subset (\theta_{a(n)})$ and a function $\tilde{V} \in C[0, 1]$ such that $\|\mathcal{V}^{\theta_{b(n)}} - \tilde{V}\|_\infty \rightarrow 0$.

Combining the convergences of $U^{\theta_{b(n)}}$ and $\mathcal{V}^{\theta_{b(n)}}$ with those in Equation (2.A.6), it follows from the triangle inequality of the metric d_{BL} that also

$$\sqrt{|J_{b(n)}|} \hat{\gamma}^{(b(n)), \theta_{b(n)}} \xrightarrow{\mathcal{D}} \tilde{U} \quad \text{and} \quad \hat{\mathcal{V}}_{b(n)}^{\theta_{b(n)}} \xrightarrow{P} \tilde{V},$$

in $D[0, 1]$ as $n \rightarrow \infty$. Now we may use that convergence in Skorokhod topology is equivalent to convergence in uniform topology whenever the limit variable continuous, see e.g. Kallenberg [2021, Theorem 23.9]. Hence the convergences above also hold in $(D[0, 1], \|\cdot\|_\infty)$.

Since \mathcal{V} is deterministic, this implies the joint convergences

$$(U^{\theta_{b(n)}}, \mathcal{V}^{\theta_{b(n)}}) \xrightarrow{\mathcal{D}} (\tilde{U}, \tilde{V}) \quad \text{and} \quad \left(\sqrt{|J_{b(n)}|} \hat{\gamma}^{(b(n)), \theta_{b(n)}}, \hat{\mathcal{V}}_{b(n)}^{\theta_{b(n)}} \right) \xrightarrow{\mathcal{D}} (\tilde{U}, \tilde{V})$$

in the product space $D[0, 1] \times D[0, 1]$ endowed with the uniform topology. Since $(\tilde{U}, \tilde{V}) \in C[0, 1] \times \{\mathcal{V}^\theta : \theta \in \Theta_0\}$ takes values in the continuity set of \mathcal{J} by assumption, the classical continuous mapping theorem lets us conclude that

$$\mathcal{J}(U^{\theta_{b(n)}}, \mathcal{V}^{\theta_{b(n)}}) \xrightarrow{\mathcal{D}} \mathcal{J}(\tilde{U}, \tilde{V}) \quad \text{and} \quad \hat{D}_{b(n)}^{\theta_{b(n)}} = \mathcal{J}\left(\sqrt{|J_{b(n)}|} \hat{\gamma}^{(b(n)), \theta_{b(n)}}, \hat{\mathcal{V}}_{b(n)}^{\theta_{b(n)}}\right) \xrightarrow{\mathcal{D}} \mathcal{J}(\tilde{U}, \tilde{V})$$

as $n \rightarrow \infty$. Now another application of the triangle inequality with $\mathcal{J}(\tilde{U}, \tilde{V})$ as intermediate value shows that (2.A.7) holds with $k(n) = b(n)$, so we are done. \square

2.A.10 Proof of Theorem 2.5.1

We will apply Corollary 2.4.8 with the functional \mathcal{J} given by

$$\mathcal{J}(f_1, f_2) = \mathbb{1}(f_2 \neq 0) \frac{\|f_1\|_\infty}{\sqrt{|f_2(1)|}}, \quad f_1, f_2 \in D[0, 1].$$

Under Assumption 2.5.1, it suffices to check continuity of \mathcal{J} on the set Υ given by

$$\Upsilon := C[0, 1] \times \{f \in C[0, 1] \mid \delta_1 \leq |f(1)|\} \supset C[0, 1] \times \overline{\{\mathcal{V}^\theta : \theta \in \Theta_0\}}.$$

To see that \mathcal{J} is continuous on Υ in the uniform topology, we note that it can be written as a composition of the continuous maps

$$\begin{aligned} \Upsilon &\longrightarrow [0, \infty) \times [\delta_1, \infty), & (f_1, f_2) &\mapsto (\|f_1\|_\infty, |f_2(1)|), \\ [0, \infty) \times [\delta_1, \infty) &\longrightarrow \mathbb{R}, & (x_1, x_2) &\mapsto \frac{x_1}{\sqrt{x_2}}. \end{aligned}$$

2 Nonparametric conditional local independence testing

Thus it follows from Corollary 2.4.8 that

$$\hat{T}_n = \frac{\sqrt{|J_n|} \sup_{t \in [0,1]} |\hat{\gamma}_t^{(n)}|}{\sqrt{\hat{\mathcal{V}}_n(1)}} = \mathcal{J} \left(\sqrt{|J_n|} \hat{\gamma}^{(n)}, \hat{\mathcal{V}}_n \right) \xrightarrow{\mathcal{D}/\Theta_0} \mathcal{J}(U, \mathcal{V}) = \frac{\|U\|_\infty}{\mathcal{V}(1)}.$$

With (B_u) a Brownian motion it follows by Proposition 2.C.2 that

$$\frac{\|U\|_\infty}{\sqrt{\mathcal{V}(1)}} \stackrel{\mathcal{D}}{=} \frac{\sup_{0 \leq t \leq 1} |B_{\mathcal{V}(t)}|}{\sqrt{\mathcal{V}(1)}} = \frac{\sup_{0 \leq u \leq \mathcal{V}(1)} |B_u|}{\sqrt{\mathcal{V}(1)}} \stackrel{\mathcal{D}}{=} \sup_{0 \leq t \leq 1} |B_t| \stackrel{\mathcal{D}}{=} S, \quad (2.A.8)$$

where we have used that \mathcal{V} is continuous and that Brownian motion is scale invariant. This establishes the first part of the theorem.

For the second part, we first note that the distribution of S is absolutely continuous with respect to Lebesgue measure, which follows from Equation (2.5.31). Then we can use Theorem 4.1 of Bengs and Holzmann [2019] to conclude that

$$\limsup_{n \rightarrow \infty} \sup_{\theta \in \Theta} |\mathbb{P}(\hat{T}_n \leq z_{1-\alpha}) - (1 - \alpha)| = 0.$$

It follows from the triangle inequality that

$$\limsup_{n \rightarrow \infty} \sup_{\theta \in \Theta} \mathbb{P}(\Psi_n^\alpha = 1) = \limsup_{n \rightarrow \infty} \sup_{\theta \in \Theta} \mathbb{P}(\hat{T}_n > z_{1-\alpha}) \leq \alpha.$$

□

2.A.11 Proof of Theorem 2.5.2

Let $0 < \alpha < \beta < 1$ be given. The second part of Theorem 2.4.6 permits us to choose $K > 0$ sufficiently large such that

$$\limsup_{n \rightarrow \infty} \sup_{\theta \in \Theta} \mathbb{P} \left((\sqrt{|J_n|} \|\hat{\gamma}^{(n),\theta} - \gamma^\theta\|_\infty) > K \right) < 1 - \beta. \quad (2.A.9)$$

We then choose $c > K + z_{1-\alpha} \sqrt{1 + C(C')^2}$ such that for all $\theta \in \mathcal{A}_{c,n}$, it holds that

$$\sqrt{|J_n|} \|\gamma^\theta\|_\infty - z_{1-\alpha} \sqrt{1 + \mathcal{V}^\theta(1)} \geq c - z_{1-\alpha} \sqrt{1 + C(C')^2} > K,$$

where we have used Lemma 2.A.2 in the first inequality. The (reverse) triangle inequality now yields that for any $\theta \in \mathcal{A}_{c,n}$

$$\begin{aligned} (\Psi_n^\theta = 0) &= (\hat{T}_n^\theta \leq z_{1-\alpha}) = \left(\|\hat{\gamma}^{(n),\theta}\|_\infty \leq \sqrt{\hat{\mathcal{V}}_n^\theta(1)} \frac{z_{1-\alpha}}{\sqrt{|J_n|}} \right) \\ &\subseteq \left(\|\gamma^\theta\|_\infty - \|\hat{\gamma}^{(n),\theta} - \gamma^\theta\|_\infty \leq \sqrt{\hat{\mathcal{V}}_n^\theta(1)} \frac{z_{1-\alpha}}{\sqrt{|J_n|}} \right) \\ &\subseteq E_1^{(n),\theta} \cup E_2^{(n),\theta}, \end{aligned}$$

where

$$\begin{aligned} E_1^{(n),\theta} &= \left(\sqrt{|J_n|} \left\| \hat{\gamma}^{(n),\theta} - \gamma^\theta \right\|_\infty > K \right), \\ E_2^{(n),\theta} &= \left(\hat{\nu}_n^\theta(1) > 1 + \nu^\theta(1) \right) \subseteq \left(|\hat{\nu}_n^\theta(1) - \nu^\theta(1)| > 1 \right). \end{aligned}$$

From Proposition 2.4.7 we know that $\limsup_{n \rightarrow \infty} \sup_{\theta \in \Theta} \mathbb{P}(E_2^{(n),\theta}) = 0$, so from the choice of K we conclude that

$$\limsup_{n \rightarrow \infty} \sup_{\theta \in \Theta} \mathbb{P}(\Psi_n = 0) \leq \limsup_{n \rightarrow \infty} \sup_{\theta \in \Theta} \mathbb{P}(E_1^{(n),\theta}) < 1 - \beta.$$

The desired statement follows from substituting $\mathbb{P}(\Psi_n = 0) = 1 - \mathbb{P}(\Psi_n = 1)$ into the above equation and simplifying. \square

2.A.12 Proof of Theorem 2.5.4

Assume that H_0 holds and note that Assumptions 2.4.1 and 2.4.2 are satisfied for every sample split $J_n^k \cup (J_n^k)^c$, $k = 1, \dots, K$.

We consider the decomposition in Equation (2.4.17) for each sample split $J_n^k \cup (J_n^k)^c$, and denote the corresponding processes by $U^{k,(n)}$, $R_1^{k,(n)}$, $R_2^{k,(n)}$, $R_3^{k,(n)}$, $D_1^{k,(n)}$, and $D_2^{k,(n)}$. For each fold $k \in \{1, \dots, K\}$, we can then apply the results in Section 2.4 for a single data split:

- By Proposition 2.4.3, we have that $U^{k,(n)} \xrightarrow{\mathcal{D}/\Theta} U$ in $D[0, 1]$, where U is a mean zero continuous Gaussian martingale with variance function \mathcal{V} .
- By Proposition 2.4.4, $R_\ell^{k,(n)} \xrightarrow{P/\Theta} 0$ in $(D[0, 1], \|\cdot\|_\infty)$ as $n \rightarrow \infty$.
- Under H_0 , the processes $D_1^{k,(n)}$ and $D_2^{k,(n)}$ are equal to the zero process almost surely.

Recall that the folds are assumed to have uniform asymptotic density, which is equivalent to $\frac{\sqrt{n}}{\sqrt{K|J_n^k|}} \rightarrow 1$ as $n \rightarrow \infty$. Thus we may also conclude that for each fixed k and ℓ ,

$$\frac{\sqrt{n}}{\sqrt{K|J_n^k|}} U^{k,(n)} \xrightarrow{\mathcal{D}/\Theta} U \quad \text{and} \quad \frac{\sqrt{n}}{K\sqrt{|J_n^k|}} R_\ell^{k,(n)} \xrightarrow{P/\Theta} 0,$$

where the convergences hold in the Skorokhod and uniform topology, respectively. Now the key observation is that

$$U^{1,(n)} \perp\!\!\!\perp \dots \perp\!\!\!\perp U^{K,(n)}.$$

To see this, note that $U^{k,(n)}$ is constructed from $(G_j, M_j)_{j \in J_n^k}$ only, and by the i.i.d. assumption of the data, the collections $(G_j, M_j)_{j \in J_n^1}, \dots, (G_j, M_j)_{j \in J_n^K}$ are jointly independent. We can therefore apply Lemma 2.B.12 iteratively to the sequences

$$\frac{\sqrt{n}}{\sqrt{K|J_n^1|}} U^{1,(n)}, \dots, \frac{\sqrt{n}}{\sqrt{K|J_n^K|}} U^{K,(n)}$$

2 Nonparametric conditional local independence testing

to conclude that their sum is uniformly convergent to the sum of K independent copies of U . Using the convolution property of the Gaussian distribution, it therefore follows that

$$\check{U}^{K,(n)} := \frac{1}{\sqrt{K}} \sum_{k=1}^K \frac{\sqrt{n}}{\sqrt{K|J_n^k|}} U^{k,(n)} \xrightarrow{\mathcal{D}/\Theta} U$$

in $D[0,1]$ as $n \rightarrow \infty$. By the uniform Slutsky theorem formulated in Lemma 2.B.11, we can therefore conclude that

$$\sqrt{n}\check{\gamma}^{K,(n)} = \check{U}^{K,(n)} + \sum_{k=1}^K \frac{\sqrt{n}}{K\sqrt{|J_n^k|}} \left(R_1^{k,(n)} + R_2^{k,(n)} + R_3^{k,(n)} + D_1^{k,(n)} + D_2^{k,(n)} \right) \xrightarrow{\mathcal{D}/\Theta_0} U$$

in $D[0,1]$ as $n \rightarrow \infty$. Note that we use $\theta \in \Theta_0$ to ensure that $D_1^{k,(n)} + D_2^{k,(n)}$ is equal to the zero process almost surely. Since the limit $(U^\theta)_{\theta \in \Theta_0}$ is tight in $C[0,1]$ by Lemma 2.A.11, Proposition 2.B.9 lets us conclude that $\sqrt{n}\|\check{\gamma}^{K,(n)}\|_\infty \xrightarrow{\mathcal{D}/\Theta_0} \|U\|_\infty$.

Consider now the cross-fitted variance estimator at its endpoint

$$\check{V}_{K,n}(1) = \frac{1}{K} \sum_{k=1}^K \frac{1}{|J_n^k|} \sum_{j \in J_n^k} \int_0^1 \left(\hat{G}_{j,s}^{k,(n)} \right)^2 dN_{j,s}.$$

From Proposition 2.4.7, we see that $\check{V}_{K,n}(1)$ is an average of K variables converging uniformly in probability to $\mathcal{V}(1)$ in the uniform topology. Hence $\check{V}_{K,n}(1)$ also converges uniformly in probability to $\mathcal{V}(1)$ in the uniform topology. We can then apply Theorem 6.3 of Bengs and Holzmann [2019], which is a uniform version of Slutsky's theorem, to conclude that

$$\check{T}_n^K = \frac{\sqrt{n}\|\check{\gamma}^{K,(n)}\|_\infty}{\check{V}_{K,n}(1)} \xrightarrow{\mathcal{D}/\Theta_0} \frac{\|U\|_\infty}{\mathcal{V}(1)} \stackrel{\mathcal{D}}{=} S,$$

as $n \rightarrow \infty$, where last equality in distribution was established in (2.A.8).

Following the second part of the proof of Theorem 2.5.1, we conclude in the X-LCT has uniform asymptotic level. \square

2.B. Uniform stochastic convergence

In this section, we discuss weak convergence of random variables with values in a metric space uniformly over a parameter set Θ . The uniformity over the parameter set can be used, for example, to establish uniform asymptotic level as well as power under local alternatives.

The content of this section extends the works of Bengs and Holzmann [2019] and Kasy [2019], and we especially build upon Appendix B of Lundborg et al. [2022b], in which uniform stochastic convergence is considered in *separable* Banach spaces and Hilbert spaces. The space space $(D[0,1], \|\cdot\|_\infty)$ of càdlàg functions endowed with the uniform

norm is a Banach space, but it is unfortunately not separable. Therefore we extend the notion of uniform stochastic convergence to random variables in metric spaces, with the condition that the limit is supported on a separable set. This allows to consider uniform weak convergence in two important special cases: i) convergence in $(D[0, 1], \|\cdot\|_\infty)$ towards variables in $(C[0, 1], \|\cdot\|_\infty)$, and ii) convergence in $D[0, 1]$ endowed with the Skorokhod metric.

The Skorokhod space $D[0, 1]$ is, if not specified otherwise, equipped with the complete Skorokhod metric d° , which makes it a Polish space, i.e., a complete and separable metric space. See for example Section 12 in Billingsley [2013] for a discussion of the Skorokhod space and in particular Equation (12.16) for a definition of d° .

2.B.1 Uniform stochastic convergence in metric spaces

Throughout this section we consider a background probability space $(\Omega, \mathbb{F}, \mathbb{P})$ and let $(\mathbb{D}, d_{\mathbb{D}})$ denote a generic metric space. We define $BL_1(\mathbb{D})$ as the set of real-valued functions on \mathbb{D} with Lipschitz norm bounded by 1, that is, functions $f: \mathbb{D} \rightarrow \mathbb{R}$ with $\|f\|_\infty \leq 1$ and $|f(x) - f(y)| \leq d_{\mathbb{D}}(x, y)$ for every $x, y \in \mathbb{D}$. Let $\mathcal{M}_1(\mathbb{D})$ denote the set of Borel probability measures on \mathbb{D} . We then define the *bounded Lipschitz metric* on $\mathcal{M}_1(\mathbb{D})$ by

$$d_{BL}(\mu, \nu) := \sup_{f \in BL_1(\mathbb{D})} \left| \int f d\mu - \int f d\nu \right|, \quad \mu, \nu \in \mathcal{M}_1(\mathbb{D}).$$

For any pair (X, Y) of \mathbb{D} -valued random variables we use the shorthand notation

$$d_{BL}(X, Y) = d_{BL}(X(\mathbb{P}), Y(\mathbb{P})) = \sup_{f \in BL_1(\mathbb{D})} |\mathbb{E}(f(X) - f(Y))|.$$

If the underlying metric space is ambiguous for d_{BL} , we will specify that it is the bounded Lipschitz metric on $\mathcal{M}_1(\mathbb{D})$ by writing $d_{BL(\mathbb{D})}$. Our interest in the bounded Lipschitz metric is due to its characterization of weak convergence.

Proposition 2.B.1. *Let X, X_1, X_2, \dots be a sequence of \mathbb{D} -valued random variables. Assume that there exists a separable subset $\mathbb{D}_0 \subseteq \mathbb{D}$ such that $\mathbb{P}(X \in \mathbb{D}_0) = 1$. Then the following are equivalent:*

- *The sequence $(X_n)_{n \geq 1}$ converges in distribution to X , i.e., for all $f \in C_b(\mathbb{D})$ it holds that $\mathbb{E}[f(X_n)] \rightarrow \mathbb{E}[f(X)]$ as $n \rightarrow \infty$.*
- *It holds that $d_{BL}(X_n, X) \rightarrow 0$ as $n \rightarrow \infty$.*

Proof. See Theorem 1.12.2, Addendum 1.12.3, and the following discussion in van der Vaart and Wellner [1996]. \square

To discuss uniform stochastic convergence, we will for the remaining part of this section let Θ be fixed set, which is used as a (possible) parameter set for every random variable. We say that a collection $(X^\theta)_{\theta \in \Theta}$ of \mathbb{D} -valued random variables is *separable*

2 Nonparametric conditional local independence testing

if there exists a separable subset $\mathbb{D}_0 \subseteq \mathbb{D}$ such that $\mathbb{P}(X^\theta \in \mathbb{D}_0) = 1$ for all $\theta \in \Theta$. If \mathbb{D} is a separable metric space, then any collection of \mathbb{D} -valued random variables is automatically separable.

Now Lemma 2.B.1 justifies the following generalization of weak convergence uniformly over Θ :

Definition 2.B.2. Let $(X_n^\theta)_{n \in \mathbb{N}, \theta \in \Theta}$ and $(X^\theta)_{\theta \in \Theta}$ be collections of \mathbb{D} -valued random variables and assume that $(X^\theta)_{\theta \in \Theta}$ is separable. We say that:

- (i) X_n^θ converges uniformly in distribution over Θ to X^θ in \mathbb{D} , and write $X_n^\theta \xrightarrow{D/\Theta} X^\theta$, if

$$\lim_{n \rightarrow \infty} \sup_{\theta \in \Theta} d_{BL(\mathbb{D})}(X_n^\theta, X^\theta) = 0.$$

- (ii) X_n^θ converges uniformly in probability over Θ to X^θ in \mathbb{D} , and write $X_n^\theta \xrightarrow{P/\Theta} X^\theta$, if

$$\lim_{n \rightarrow \infty} \sup_{\theta \in \Theta} \mathbb{P}(d_{\mathbb{D}}(X_n^\theta, X^\theta) > \epsilon) = 0$$

for every $\epsilon > 0$.

♣

If for some $\mu \in \mathcal{M}_1(\mathbb{D})$, it holds that $X_n^\theta \xrightarrow{D/\Theta} X^\theta$ with $X^\theta(\mathbb{P}) = \mu$ for all $\theta \in \Theta$, we also write $X_n^\theta \xrightarrow{D/\Theta} \mu$. Similarly, we may replace the limit random variable X^θ by a point $x \in \mathbb{D}$ by interpreting x as the constant map $(\theta, \omega) \mapsto x$ for $\theta \in \Theta$ and $\omega \in \Omega$.

Note that if the parameter set $\Theta = \{\theta_0\}$ is a singleton, then each type of uniform convergence reduces to the corresponding classical definition of convergence in distribution or probability. If \mathbb{D} is a separable Banach space, we note that Definition 2.B.2 coincides with Definition 3 in Lundborg et al. [2022b].

Proposition 2.B.3. Let $(X_n^\theta)_{n \in \mathbb{N}, \theta \in \Theta}$ and $(X^\theta)_{\theta \in \Theta}$ be collections of \mathbb{D} -valued random variables and assume $(X^\theta)_{\theta \in \Theta}$ is separable. Then the following are equivalent:

- $X_n^\theta \xrightarrow{D/\Theta} X^\theta$ as $n \rightarrow \infty$.
- For any sequence $(\theta_n)_{n \in \mathbb{N}} \subseteq \Theta$ it holds that $d_{BL}(X_n^{\theta_n}, X^{\theta_n}) \rightarrow 0$ as $n \rightarrow \infty$.
- For any sequence $(\theta_n)_{n \in \mathbb{N}} \subseteq \Theta$ there exists a subsequence $(\theta_{k(n)})_{n \in \mathbb{N}}$, with $k: \mathbb{N} \rightarrow \mathbb{N}$ strictly increasing, such that

$$\lim_{k \rightarrow \infty} d_{BL}(X_{k(n)}^{\theta_{k(n)}}, X^{\theta_{k(n)}}) = 0.$$

Moreover, $X_n^\theta \xrightarrow{P/\Theta} X^\theta$ if and only if for any sequence $(\theta_n)_{n \in \mathbb{N}} \subseteq \Theta$ and any $\epsilon > 0$ it holds that

$$\lim_{n \rightarrow \infty} \mathbb{P}(d_{\mathbb{D}}(X_n^{\theta_n}, X^{\theta_n}) > \epsilon) = 0.$$

Proof. This is essentially Lemma 1 in Kasy [2019] for \mathbb{D} -valued random variables, except that we have added the equivalent condition c). The proof for the characterization of uniform convergence in probability is identical to the one given by Kasy [2019], so we focus on the equivalence between a), b), and c). To this end, we to prove that $a) \implies b) \implies c) \implies a)$.

The fact that a) implies b) follows directly from applying the bound

$$d_{BL}(X_n^{\theta_n}, X^{\theta_n}) \leq \sup_{\theta \in \Theta} d_{BL}(X_n^\theta, X^\theta)$$

and taking the limit as $n \rightarrow \infty$. We also see that b) implies c) since any sequence is a subsequence of itself.

We show that c) implies a) by contraposition. Assume the negation of a), that is, there exists an $\epsilon > 0$ and a sequence $(\theta_n)_{n \in \mathbb{N}} \subseteq \Theta$ such that

$$d_{BL}(X_n^{\theta_n}, X^{\theta_n}) > \epsilon$$

for all $n \in \mathbb{N}$. Then, for all subsequences $(\theta_{k(n)})$ of (θ_n) , it holds that $d_{BL}(X_{k(n)}^{\theta_{k(n)}}, X^{\theta_{k(n)}})$ does not converge to zero. This implies the negation of c). \square

Proposition 2.B.3 will allow us to extend many results for classical stochastic convergence to uniform stochastic convergence.

Corollary 2.B.4. *Let $(X_n^\theta)_{n \in \mathbb{N}, \theta \in \Theta}$ be a collection of \mathbb{D} -valued random variables and let $x \in \mathbb{D}$. Then $X_n^\theta \xrightarrow{\mathcal{D}/\Theta} x$ if and only if $X_n^\theta \xrightarrow{P/\Theta} x$.*

Proof. For any sequence $(\theta_n)_{n \in \mathbb{N}} \subseteq \Theta$, recall that $X_n^{\theta_n} \xrightarrow{\mathcal{D}} x$ if and only if $X_n^{\theta_n} \xrightarrow{P} x$, see e.g. Lemma 5.1 in Kallenberg [2021]. Hence the statement follows directly from Proposition 2.B.3 (combined with Proposition 2.B.1). \square

Our goal is to prove uniform versions of Slutsky's theorem for $D[0, 1]$, Rebolledo's central limit theorem, and the chaining lemma for stochastic processes. To prove Slutsky's lemma for $D[0, 1]$, we first prove a general result for metric spaces.

Lemma 2.B.5. *Let $(X^\theta, X_n^\theta, Y_n^\theta)_{n \in \mathbb{N}, \theta \in \Theta}$ be a collection of \mathbb{D} -valued random variables and assume that $(X^\theta)_{\theta \in \Theta}$ is separable. If $X_n^\theta \xrightarrow{\mathcal{D}/\Theta} X^\theta$ and $d_{\mathbb{D}}(X_n^\theta, Y_n^\theta) \xrightarrow{P/\Theta} 0$, then it also holds that $Y_n^\theta \xrightarrow{\mathcal{D}/\Theta} X^\theta$.*

Proof. By the triangle inequality of the bounded Lipschitz metric, we observe that

$$\sup_{\theta \in \Theta} d_{BL}(Y_n^\theta, X^\theta) \leq \sup_{\theta \in \Theta} d_{BL}(Y_n^\theta, X_n^\theta) + \sup_{\theta \in \Theta} d_{BL}(X_n^\theta, X^\theta).$$

The last term converges to zero by the assumption of $X_n^\theta \xrightarrow{\mathcal{D}/\Theta} X^\theta$. For the other term, let $\epsilon > 0$ and use the partition

$$(d_{\mathbb{D}}(X_n^\theta, Y_n^\theta) \leq \epsilon) \cup (d_{\mathbb{D}}(X_n^\theta, Y_n^\theta) > \epsilon)$$

2 Nonparametric conditional local independence testing

to obtain that

$$\begin{aligned} d_{BL}(X_n^\theta, Y_n^\theta) &= \sup_{f \in BL_1(\mathbb{D})} |\mathbb{E}[f(X_n^\theta) - f(Y_n^\theta)]| \\ &\leq \epsilon + \sup_{f \in BL_1(\mathbb{D})} \mathbb{E}[|f(X_n^\theta) - f(Y_n^\theta)| \mathbb{1}(d_{\mathbb{D}}(X_n^\theta, Y_n^\theta) > \epsilon)] \\ &\leq \epsilon + \mathbb{P}(d_{\mathbb{D}}(X_n^\theta, Y_n^\theta) > \epsilon). \end{aligned}$$

Taking the supremum over Θ and the limit superior for $n \rightarrow \infty$ finishes the proof. \square

The following formulation of the continuous mapping theorem is analogous to Theorem 1 in Kasy [2019]. The proof is almost identical, but we repeat it here for completeness.

Proposition 2.B.6. *Let (\mathbb{D}_1, d_1) and (\mathbb{D}_2, d_2) be metric spaces, and let $\Phi: \mathbb{D}_1 \rightarrow \mathbb{D}_2$ be a Lipschitz continuous map. Let $(X_n^\theta)_{n \in \mathbb{N}, \theta \in \Theta}$ and $(X^\theta)_{\theta \in \Theta}$ be collections of \mathbb{D}_1 -valued random variables, and assume $(X^\theta)_{\theta \in \Theta}$ is separable.*

If $X_n^\theta \xrightarrow{D/\Theta} X^\theta$ in \mathbb{D}_1 , then $\Phi(X_n^\theta) \xrightarrow{D/\Theta} \Phi(X^\theta)$ in \mathbb{D}_2 .

Proof. Note first that if X^θ is in a separable subset $\mathbb{D}_0 \subseteq \mathbb{D}_1$, then the variables $\Phi(X_\theta)$ for $\theta \in \Theta$ are all in the separable subset $\Phi(\mathbb{D}_0) \subseteq \mathbb{D}_2$. Hence it is well-defined to consider uniform convergence in distribution towards $(\Phi(X^\theta))_{\theta \in \Theta}$. Let $f \in BL_1(\mathbb{D}_2)$ and let K be the Lipschitz constant of Φ . Consider the map

$$g: \mathbb{D}_1 \rightarrow \mathbb{R}, \quad g(x) = \min(1, K^{-1})f(\Phi(x)).$$

Then $\|g\|_\infty \leq \|f\|_\infty \leq 1$ and for all $x, y \in \mathbb{D}_1$,

$$\begin{aligned} |g(x) - g(y)| &\leq \min(1, K^{-1})d_2(\Phi(x), \Phi(y)) \\ &\leq \min(1, K^{-1})Kd_1(x, y) \leq d_1(x, y) \end{aligned}$$

Hence $g \in BL_1(\mathbb{D}_1)$. It follows that

$$\begin{aligned} d_{BL_1(\mathbb{D}_2)}(\Phi(X_n^\theta), \Phi(X^\theta)) &= \sup_{f \in BL_1(\mathbb{D}_2)} |\mathbb{E}[f(\Phi(X_n^\theta)) - f(\Phi(X^\theta))]| \\ &\leq \frac{1}{\min(1, K^{-1})} \sup_{g \in BL_1(\mathbb{D}_1)} |\mathbb{E}[g(X_n^\theta) - g(X^\theta)]| \\ &\leq \max(1, K) \cdot d_{BL_1(\mathbb{D}_1)}(X_n^\theta, X^\theta) \end{aligned}$$

Taking the supremum over Θ and the limit superior as $n \rightarrow \infty$ finish the proof. \square

We will also need the following two notions of tightness.

Definition 2.B.7. Let $(\mu^\theta)_{\theta \in \Theta}$ be a family of probability measures on \mathbb{D} , and let $(X^\theta)_{\theta \in \Theta}$ and $(X_n^\theta)_{n \in \mathbb{N}, \theta \in \Theta}$ be collections of \mathbb{D} -valued random variables.

- i) We say that $(\mu^\theta)_{\theta \in \Theta}$ is *tight* if for any $\varepsilon > 0$, there exists a compact set $K \subseteq \mathbb{D}$ such that $\sup_{\theta \in \Theta} \mu^\theta(K^c) < \varepsilon$. We say that $(X^\theta)_{\theta \in \Theta}$ is *uniformly tight* if the collection of distributions $(X^\theta(\mathbb{P}))_{\theta \in \Theta}$ is tight.
- ii) The sequence $((X_n^\theta)_{\theta \in \Theta})_{n \in \mathbb{N}}$ of collections is said to be *sequentially tight* if for any sequence $(\theta_n)_{n \in \mathbb{N}} \subset \Theta$, the sequence of distributions $(X_n^{\theta_n}(\mathbb{P}))_{n \in \mathbb{N}}$ is tight.

♣

Definition 2.B.7 i) is a classical concept, whereas sequential tightness was introduced by Lundborg et al. [2022b] and relaxes uniform tightness for sequences of variables parametrized over an infinite set.

The importance of tightness is mainly due to Prokhorov's theorem [Kallenberg, 2021, Theorem 23.2], which states that if \mathbb{D} is a Polish space², then $(\mu^\theta)_{\theta \in \Theta}$ is tight if and only if all sequences in $(\mu^\theta)_{\theta \in \Theta}$ have a weakly convergent subsequence.

The continuous mapping theorem in Proposition 2.B.6 is more restrictive than the classical theorem as it requires Lipschitz continuity. However, we also have an alternative version of uniform continuous mapping when the limit variable is tight.

Proposition 2.B.8. *Let (\mathbb{D}_1, d_1) and (\mathbb{D}_2, d_2) be Polish spaces, and let $(X_n^\theta)_{n \in \mathbb{N}, \theta \in \Theta}$ and $(X^\theta)_{\theta \in \Theta}$ be collections of \mathbb{D}_1 -valued random variables. Assume $(X^\theta)_{\theta \in \Theta}$ is uniformly tight, and let $\Phi: \mathbb{D}_1 \rightarrow \mathbb{D}_2$ be a map that is continuous on the support of $(X^\theta)_{\theta \in \Theta}$.*

If $X_n^\theta \xrightarrow{\mathcal{D}/\Theta} X^\theta$ in \mathbb{D}_1 , then $\Phi(X_n^\theta) \xrightarrow{\mathcal{D}/\Theta} \Phi(X^\theta)$ in \mathbb{D}_2 .

Proof. Same as the proof of Proposition 10 in Lundborg et al. [2022b], but with norms of differences replaced by metric distances. \square

2.B.2 Uniform stochastic convergence in Skorokhod space

In this section we consider the special case where $(\mathbb{D}, d_{\mathbb{D}})$ is the Skorokhod space $(D[0, 1], d^\circ)$. We can also equip $D[0, 1]$ with the uniform norm, $\|x\|_\infty = \sup_{t \in [0, 1]} |x_t|$, and it known that weak convergence based on either $\|\cdot\|_\infty$ or d° are equivalent when the limit is continuous. We now extend this result to stochastic convergence uniformly over Θ .

Proposition 2.B.9 (Skorokhod equivalence). *Let $(X_n^\theta)_{n \in \mathbb{N}, \theta \in \Theta}$ be a collection of $D[0, 1]$ -valued random variables and let $(X^\theta)_{\theta \in \Theta}$ be a uniformly tight collection of $C[0, 1]$ -valued random variables. Then $X_n^\theta \xrightarrow{\mathcal{D}/\Theta} X^\theta$ in $(D[0, 1], d^\circ)$ if and only if $X_n^\theta \xrightarrow{\mathcal{D}/\Theta} X^\theta$ in $(D[0, 1], \|\cdot\|_\infty)$. In the affirmative, $\|X_n^\theta\|_\infty \xrightarrow{\mathcal{D}/\Theta} \|X^\theta\|_\infty$.*

Proof. To avoid ambiguity in the topology on $D[0, 1]$, we will throughout this proof use \mathbb{D}° to denote the metric space $(D[0, 1], d^\circ)$ and we use \mathbb{D}_∞ to denote the Banach space $(D[0, 1], \|\cdot\|_\infty)$. Note also that $C[0, 1]$ is separable within \mathbb{D}_∞ , so $(X^\theta)_{\theta \in \Theta}$ is separable, and hence the convergence $X_n^\theta \xrightarrow{\mathcal{D}/\Theta} X^\theta$ is well-defined in the non-separable space \mathbb{D}_∞ .

The 'if' part is clear since $d^\circ(x, y) \leq \|x - y\|_\infty$ for all $x, y \in D[0, 1]$.

²The 'only if' part does not require separability nor completeness.

2 Nonparametric conditional local independence testing

For the ‘only if’ part, assume that $X_n^\theta \xrightarrow{\mathcal{D}/\Theta} X^\theta$ in \mathbb{D}° and let $(\theta_n) \subseteq \Theta$ be an arbitrary sequence. Since $(X^{\theta_n}(\mathbb{P}))$ is tight, Prokhorov’s Theorem asserts that there exists a subsequence $(\theta_{k(n)})$ and a probability distribution μ on $C[0, 1]$ such that $X^{\theta_{k(n)}}(\mathbb{P}) \xrightarrow{wk} \mu$ in \mathbb{D}° . By the triangle inequality

$$d_{BL(\mathbb{D}^\circ)}(X_{k(n)}^{\theta_{k(n)}}, \mu) \leq d_{BL(\mathbb{D}^\circ)}(X_{k(n)}^{\theta_{k(n)}}, X^{\theta_{k(n)}}) + d_{BL(\mathbb{D}^\circ)}(X^{\theta_{k(n)}}, \mu) \rightarrow 0, \quad n \rightarrow 0.$$

This shows that also $X_{k(n)}^{\theta_{k(n)}}(\mathbb{P}) \xrightarrow{wk} \mu$ in \mathbb{D}° . Now we can use that weak convergence in the Skorokhod topology and the uniform topology are equivalent when the limit is continuous [Kallenberg, 2021, Theorem 23.9 (iii)]. We therefore conclude that the convergences $X^{\theta_{k(n)}}(\mathbb{P}) \xrightarrow{wk} \mu$ and $X_{k(n)}^{\theta_{k(n)}}(\mathbb{P}) \xrightarrow{wk} \mu$ also hold in \mathbb{D}_∞ . But then another use of the triangle inequality shows that

$$d_{BL(\mathbb{D}_\infty)}(X_{k(n)}^{\theta_{k(n)}}, X^{\theta_{k(n)}}) \leq d_{BL(\mathbb{D}_\infty)}(X_{k(n)}^{\theta_{k(n)}}, \mu) + d_{BL(\mathbb{D}_\infty)}(\mu, X^{\theta_{k(n)}}) \rightarrow 0.$$

Since $(\theta_{k(n)})$ is a subsequence of the arbitrarily chosen sequence (θ_n) , we conclude that $X_n^\theta \xrightarrow{\mathcal{D}/\Theta} X^\theta$ in \mathbb{D}_∞ by Proposition 2.B.3.

Finally, as the uniform norm is Lipschitz continuous as a map from \mathbb{D}_∞ to \mathbb{R} , the continuous mapping theorem formulated in Proposition 2.B.6 yields that

$$X_n^\theta \xrightarrow{\mathcal{D}/\Theta} X^\theta \text{ in } \mathbb{D}_\infty \quad \implies \quad \|X_n^\theta\|_\infty \xrightarrow{\mathcal{D}/\Theta} \|X\|_\infty.$$

This establishes the last part of the lemma. \square

Using $\|\mu\|_\infty$ to denote the pushforward measure for any $\mu \in \mathcal{M}_1(D([0, 1]))$ we restate the result above for a fixed limit distribution.

Corollary 2.B.10. *Let $(X_n^\theta)_{n \in \mathbb{N}, \theta \in \Theta}$ be a collection of $D[0, 1]$ -valued random variables and let μ be a probability measure on $C[0, 1]$. Then $X_n^\theta \xrightarrow{\mathcal{D}/\Theta} \mu$ in $(D[0, 1], d^\circ)$ if and only if $X_n^\theta \xrightarrow{\mathcal{D}/\Theta} \mu$ in $(D[0, 1], \|\cdot\|_\infty)$. In the affirmative, $\|X_n^\theta\|_\infty \xrightarrow{\mathcal{D}/\Theta} \|\mu\|_\infty$.*

Proof. Since μ is a probability measure on the Polish space $C[0, 1]$, it is, in particular, tight [Billingsley, 2013, Theorem 1.3]. Hence the statement is a special case of Proposition 2.B.9. \square

Now we are ready to prove a uniform version of Slutsky’s theorem in the Skorokhod space.

Lemma 2.B.11 (Uniform Slutsky in Skorokhod space). *Let $(X^\theta, X_n^\theta, Y_n^\theta)_{n \in \mathbb{N}, \theta \in \Theta}$ be a collection of $D[0, 1]$ -valued random variables such that $Y_n^\theta \xrightarrow{P/\Theta} 0$ and $X_n^\theta \xrightarrow{\mathcal{D}/\Theta} X^\theta$ in $D[0, 1]$. Then it holds that $X_n^\theta + Y_n^\theta \xrightarrow{\mathcal{D}/\Theta} X^\theta$.*

Proof. Since $Y_n^\theta \xrightarrow{P/\Theta} 0$, Corollary 2.B.10 implies that $\|Y_n^\theta\|_\infty \xrightarrow{\mathcal{D}/\Theta} 0$, and Corollary 2.B.4 implies that $\|Y_n^\theta\|_\infty \xrightarrow{P/\Theta} 0$. Using the trivial estimate $d^\circ(x+y, x) \leq \|(x+y) - x\|_\infty = \|y\|_\infty$ for $x, y \in D[0, 1]$, it follows that $d^\circ(X_n^\theta + Y_n^\theta, X_n^\theta) \xrightarrow{P/\Theta} 0$. Combining the latter with $X_n^\theta \xrightarrow{\mathcal{D}/\Theta} X^\theta$, the desired conclusion now follows from Lemma 2.B.5. \square

We also have a related result for sums of independent sequences.

Lemma 2.B.12. *Let $(X_n^\theta, Y_n^\theta)_{n \in \mathbb{N}, \theta \in \Theta}$ be a collection of $D[0, 1]$ -valued random variables and let $(X^\theta)_{\theta \in \Theta}$ and $(Y^\theta)_{\theta \in \Theta}$ be uniformly tight collections of $C[0, 1]$ -valued random variables. Assume that $X_n^\theta \xrightarrow{\mathcal{D}/\Theta} X^\theta$ and $Y_n^\theta \xrightarrow{\mathcal{D}/\Theta} Y^\theta$ in $D[0, 1]$, and that for each $\theta \in \Theta$ and $n \in \mathbb{N}$, it holds that $X_n^\theta \perp\!\!\!\perp Y_n^\theta$. Let Z^θ have distribution $X^\theta(\mathbb{P}) * Y^\theta(\mathbb{P})$, that is, the same distribution as the sum of two independent copies of each of X^θ and Y^θ .*

Then it also holds that $X_n^\theta + Y_n^\theta \xrightarrow{\mathcal{D}/\Theta} Z^\theta$ in $(D[0, 1], \|\cdot\|_\infty)$.

Proof. We may assume without loss of generality that $X^\theta \perp\!\!\!\perp Y^\theta$ and that $Z^\theta = X^\theta + Y^\theta$. Let $(\theta_n) \subseteq \Theta$ be an arbitrary sequence. By tightness of $(X^\theta)_{\theta \in \Theta}$ and $(Y^\theta)_{\theta \in \Theta}$, we can apply Prokhorov's theorem twice to obtain probability measures μ and ν on $C[0, 1]$, and a subsequence $(\theta_{k(n)})$, such that $X^{\theta_{k(n)}}(\mathbb{P}) \xrightarrow{wk} \mu$ and $Y^{\theta_{k(n)}}(\mathbb{P}) \xrightarrow{wk} \nu$. Hence the product measures converge,

$$X^{\theta_{k(n)}}(\mathbb{P}) \otimes Y^{\theta_{k(n)}}(\mathbb{P}) \xrightarrow{wk} \mu \otimes \nu,$$

in $C[0, 1]$ as $n \rightarrow \infty$, see, for example, Theorem 2.8 (ii) in Billingsley [2013].

Since $X_n^\theta \xrightarrow{\mathcal{D}/\Theta} X^\theta$ in $D[0, 1]$ by assumption and (X^θ) is uniformly tight in $C[0, 1]$, Proposition 2.B.9 implies that the convergence also holds in $(D[0, 1], \|\cdot\|_\infty)$. The triangle inequality now yields

$$d_{BL}(X_{k(n)}^{\theta_{k(n)}}, \mu) \leq d_{BL}(X_{k(n)}^{\theta_{k(n)}}, X^{\theta_{k(n)}}) + d_{BL}(X^{\theta_{k(n)}}, \mu) \rightarrow 0,$$

so also $X_{k(n)}^{\theta_{k(n)}}(\mathbb{P}) \xrightarrow{wk} \mu$ in $(D[0, 1], \|\cdot\|_\infty)$. An analogous computation shows that $Y_{k(n)}^{\theta_{k(n)}}(\mathbb{P}) \xrightarrow{wk} \nu$, and hence also

$$X_{k(n)}^{\theta_{k(n)}}(\mathbb{P}) \otimes Y_{k(n)}^{\theta_{k(n)}}(\mathbb{P}) \xrightarrow{wk} \mu \otimes \nu$$

in the product space $D[0, 1] \times D[0, 1]$ endowed with the uniform product topology. From the independence statements $X^\theta \perp\!\!\!\perp Y^\theta$ and $X_n^\theta \perp\!\!\!\perp Y_n^\theta$, we have thus shown that

$$(X^{\theta_{k(n)}}, Y^{\theta_{k(n)}}) \xrightarrow{\mathcal{D}} \mu \otimes \nu \quad \text{and} \quad (X_{k(n)}^{\theta_{k(n)}}, Y_{k(n)}^{\theta_{k(n)}}) \xrightarrow{\mathcal{D}} \mu \otimes \nu$$

in the uniform product topology. Since addition $+: D[0, 1] \times D[0, 1] \rightarrow D[0, 1]$ is continuous with respect to this topology, we conclude by the classical continuous mapping theorem that

$$Z^{\theta_{k(n)}} = X^{\theta_{k(n)}} + Y^{\theta_{k(n)}} \xrightarrow{\mathcal{D}} \mu * \nu \quad \text{and} \quad X_{k(n)}^{\theta_{k(n)}} + Y_{k(n)}^{\theta_{k(n)}} \xrightarrow{\mathcal{D}} \mu * \nu.$$

It now follows that

$$\begin{aligned} & d_{BL}(X_{k(n)}^{\theta_{k(n)}} + Y_{k(n)}^{\theta_{k(n)}}, Z^{\theta_{k(n)}}) \\ & \leq d_{BL}(X_{k(n)}^{\theta_{k(n)}} + Y_{k(n)}^{\theta_{k(n)}}, \mu * \nu) + d_{BL}(\mu * \nu, Z^{\theta_{k(n)}}) \rightarrow 0. \end{aligned}$$

Since $(\theta_{k(n)})$ is a subsequence of the arbitrarily chosen sequence (θ_n) , we conclude that $X_n^\theta + Y_n^\theta \xrightarrow{\mathcal{D}/\Theta} Z^\theta$ in $(D[0, 1], \|\cdot\|_\infty)$ by Proposition 2.B.3. \square

2 Nonparametric conditional local independence testing

We also need the following lemma, which is a generalization of the classical result: pointwise convergence of a sequence of monotone functions towards a continuous limit is in fact uniform over compact intervals.

Lemma 2.B.13. *Let $(X_n^\theta)_{n \in \mathbb{N}, \theta \in \Theta}$ be a collection of $D[0, 1]$ -valued random variables with non-decreasing sample paths. Let $(f^\theta)_{\theta \in \Theta} \subset C[0, 1]$ be a uniformly equicontinuous collection of non-decreasing functions. If $X_n^\theta(t) \xrightarrow{P/\Theta} f^\theta(t)$ for each $t \in [0, 1]$, then it also holds that*

$$\sup_{t \in [0, 1]} |X_n^\theta(t) - f^\theta(t)| \xrightarrow{P/\Theta} 0.$$

Proof. Let $\epsilon > 0$. By uniform equicontinuity we can find $0 = t_1 < \dots < t_k = 1$ such that $f^\theta(t_i) - f^\theta(t_{i-1}) < \epsilon/2$ for all θ and i . Using that X_n^θ and f^θ are non-decreasing, we observe that for $t_{i-1} \leq t \leq t_i$:

$$\begin{aligned} X_n^\theta(t) - f^\theta(t) &\leq X_n^\theta(t_i) - f^\theta(t_i) + \epsilon/2, \\ X_n^\theta(t) - f^\theta(t) &\geq X_n^\theta(t_{i-1}) - f^\theta(t_{i-1}) - \epsilon/2. \end{aligned}$$

Combining the inequalities over the entire grid we have

$$\sup_{t \in [0, 1]} |X_n^\theta(t) - f^\theta(t)| \leq \max_{i=0, \dots, k} |X_n^\theta(t_i) - f^\theta(t_i)| + \epsilon/2.$$

By assumption, $X_n^\theta(t) \xrightarrow{P/\Theta} f^\theta(t)$ for each t , and in particular

$$\max_{i=0, \dots, k} |X_n^\theta(t_i) - f^\theta(t_i)| \xrightarrow{P/\Theta} 0$$

as $n \rightarrow \infty$. We therefore conclude that

$$\sup_{\theta \in \Theta} \mathbb{P} \left(\sup_{t \in [0, 1]} |X_n^\theta(t) - f^\theta(t)| > \epsilon \right) \leq \sup_{\theta \in \Theta} \mathbb{P} \left(\max_{i=0, \dots, k} |X_n^\theta(t_i) - f^\theta(t_i)| > \epsilon/2 \right) \rightarrow 0$$

as $n \rightarrow \infty$. □

The last auxiliary result of this section is an example of Prokhorov's method of "tightness + identification of limit".

Lemma 2.B.14. *Let $(\mathbb{D}, d_{\mathbb{D}})$ be either $(C[0, 1], \|\cdot\|_\infty)$ or $(D[0, 1], d^\circ)$, and let $(X^\theta, X_n^\theta)_{n \in \mathbb{N}, \theta \in \Theta}$ be a collection of \mathbb{D} -valued random variables with $(X^\theta)_{\theta \in \Theta}$ separable. Suppose that*

- *The finite dimensional marginals converge uniformly: for any $0 \leq t_1 < \dots < t_k \leq 1$*

$$\pi_{t_1, \dots, t_k}(X_n^\theta) \xrightarrow{D/\Theta} \pi_{t_1, \dots, t_k}(X^\theta), \quad n \rightarrow \infty,$$

where $\pi_{t_1, \dots, t_k} : \mathbb{D} \rightarrow \mathbb{R}^k$ is the projection given by $\pi_{t_1, \dots, t_k}(x) = (x(t_1), \dots, x(t_k))$.

- *$(X_n^\theta)_{n \in \mathbb{N}, \theta \in \Theta}$ is sequentially tight.*

- $(X^\theta)_{n \in \mathbb{N}, \theta \in \Theta}$ is uniformly tight.

Then $X_n^\theta \xrightarrow{D/\Theta} X^\theta$ as $n \rightarrow \infty$.

Proof. The statement is analogous to Proposition 18 in Lundborg et al. [2022b], the difference being that the functionals $\langle \cdot, h \rangle$ in Lundborg et al. [2022b] have been replaced by the functionals π_{t_1, \dots, t_k} .

The proof of Lundborg et al. [2022b] also works in our case, given that the finite dimensional marginals form a separating class for the both the Borel algebra on $C[0, 1]$ and the Borel algebra on $D[0, 1]$. This is established in Billingsley [2013], Example 1.3 and Theorem 12.5 (iii). \square

2.B.3 Chaining in time uniformly over a parameter

We extend the basic chaining arguments to hold uniformly over Θ . Our arguments closely follow those of Pollard [1984, Chapter VII.2.] and Newey [1991]. The results are formulated for processes indexed over a general metric space T , but we will only apply the results in the case $T = [0, 1]$. We have the following extension of stochastic equicontinuity to the uniform setting.

Definition 2.B.15. A collection of sequences

$$(Z^{(n), \theta})_{n \in \mathbb{N}, \theta \in \Theta} = \left(Z_t^{(n), \theta} \right)_{t \in T, n \in \mathbb{N}, \theta \in \Theta}$$

of stochastic processes indexed over a metric space (T, d) is called *stochastically equicontinuous uniformly over Θ* if for all $\epsilon, \eta > 0$ there exists $\delta > 0$ such that

$$\limsup_{n \rightarrow \infty} \sup_{\theta \in \Theta} \mathbb{P} \left(\sup_{s, t \in T: d(s, t) \leq \delta} |Z_s^{(n), \theta} - Z_t^{(n), \theta}| > \epsilon \right) < \eta.$$

♣

In Section 2.8.2 of van der Vaart and Wellner [1996], the same definition is given in the context of empirical processes. Recall that we write, e.g., $Z^{(n)}$ as a shorthand for $Z^{(n), \theta}$ and let the dependency on θ be implicit for notational ease. We also write $\sup_{d(s, t) \leq \delta}$ as a shorthand for $\sup_{s, t \in T: d(s, t) \leq \delta}$. Definition 2.B.15 is a direct extension of pointwise stochastic equicontinuity. Accordingly, Theorem 2.1 from Newey [1991] generalizes as follows:

Lemma 2.B.16. Let $(Z_t^{(n)})_{t \in T, n \in \mathbb{N}}$ be a sequence of stochastic processes indexed by a compact metric space T . Assume that $(Z_t^{(n)})$ is stochastically equicontinuous uniformly over Θ and that for each $t \in T$ it holds that $Z_t^{(n)} \xrightarrow{P/\Theta} 0$. Then $\sup_{t \in T} |Z_t^{(n)}| \xrightarrow{P/\Theta} 0$ as $n \rightarrow \infty$.

2 Nonparametric conditional local independence testing

Proof. Let $\varepsilon, \eta > 0$ be given, and let $\delta > 0$ be the corresponding distance obtained from the uniform stochastic equicontinuity of $(Z^{(n)})$. By compactness of T there exists a finite set $T^* \subseteq T$ such that $T = \bigcup_{t \in T^*} B(t, \delta)$. By the triangle inequality we get that

$$\sup_{t \in T} |Z_t^{(n)}| = \sup_{t \in T^*} \sup_{s \in B(t, \delta)} |Z_s^{(n)}| \leq \sup_{t \in T^*} |Z_t^{(n)}| + \sup_{t \in T^*} \sup_{s \in B(t, \delta)} |Z_s^{(n)} - Z_t^{(n)}|.$$

Since T^* is finite, it follows that $\sup_{t \in T^*} |Z_t^{(n)}| \xrightarrow{P/\Theta} 0$, which combined with the inequality implies that

$$\begin{aligned} & \limsup_{n \rightarrow \infty} \sup_{\theta \in \Theta} \mathbb{P} \left(\sup_{t \in T} |Z_t^{(n)}| > 2\varepsilon \right) \\ & \leq 0 + \limsup_{n \rightarrow \infty} \sup_{\theta \in \Theta} \mathbb{P} \left(\sup_{t \in T^*} \sup_{s \in B(t, \delta)} |Z_t^{(n)} - Z_s^{(n)}| > \varepsilon \right) \leq \eta. \end{aligned}$$

As $\varepsilon, \eta > 0$ were chosen arbitrarily, we conclude that $\sup_{t \in T} |Z_t^{(n)}| \xrightarrow{P/\Theta} 0$. \square

To establish uniform stochastic equicontinuity we extend the chaining lemma to a uniform setting. To formulate the theorem we first need some classical definitions related to chaining.

Definition 2.B.17. Let T be a compact metric space. A subset $T^* \subseteq T$ is called a δ -net if $\bigcup_{t \in T^*} B(t, \delta) = T$. The covering number

$$N(\delta) = N(\delta, T) := \min\{|T^*| : T^* \subseteq T, T^* \text{ is a } \delta\text{-net}\}$$

is the smallest possible cardinality of a δ -net, which is finite by compactness. The associated covering integral is

$$J(\delta) = \int_0^\delta (2 \log(N(\epsilon)/\epsilon))^{\frac{1}{2}} d\epsilon, \quad 0 \leq \delta \leq 1.$$

♣

Lemma 2.B.18. Let (T, d) be a metric space with finite covering integral $J(\cdot)$ and let $(Z_t^\theta)_{t \in T, \theta \in \Theta}$ be a collection of stochastic processes indexed by T with continuous sample paths. Assume there is a uniform constant $\varsigma > 0$ such that, for all $s, t \in T$ and $\eta > 0$,

$$\sup_{\theta \in \Theta} \mathbb{P} \left(|Z_s^\theta - Z_t^\theta| > \eta \cdot d(s, t) \right) \leq 2e^{-\frac{\eta^2}{2\varsigma^2}}.$$

Then, for all $0 < \epsilon < 1$,

$$\sup_{\theta \in \Theta} \mathbb{P} \left(\sup_{d(s, t) \leq \epsilon} |Z_s^\theta - Z_t^\theta| > 26\varsigma J(\epsilon) \right) \leq 2\epsilon.$$

Proof. The lemma is a direct consequence of classical chaining lemma [Pollard, 1984, page 144]. For each $\theta \in \Theta$, the conditions of the chaining lemma are met for $(Z_t^\theta)_{t \in T}$ with sub-exponential factor ς . This implies, in particular, that for any $\theta \in \Theta$ and $0 < \epsilon < 1$,

$$\mathbb{P}\left(\sup_{d(s,t) \leq \epsilon} |Z_s^\theta - Z_t^\theta| > 26\varsigma J(\epsilon)\right) \leq 2\epsilon,$$

which is equivalent to the conclusion of the lemma. \square

This immediately implies the following corollary.

Corollary 2.B.19. *Let (T, d) be a metric space with finite covering integral $J(\cdot)$ and let $(Z^{(n),\theta})$ be a sequence of stochastic processes on T with continuous sample paths. Assume there exists a constant $\varsigma > 0$ such that, for all $s, t \in T$ and $\eta > 0$ and $n \in \mathbb{N}$,*

$$\sup_{\theta \in \Theta} \mathbb{P}\left(|Z_s^{(n),\theta} - Z_t^{(n),\theta}| > \eta \cdot d(s, t)\right) \leq 2e^{-\frac{\eta^2}{2\varsigma^2}}.$$

Then $(Z^{(n)})$ is stochastically equicontinuous uniformly over Θ .

For stochastic processes with continuous sample paths, stochastic equicontinuity turns out to be equivalent to sequential tightness (Definition 2.B.7 ii).

Proposition 2.B.20. *Let $(Z^{(n),\theta})_{n \in \mathbb{N}, \theta \in \Theta}$ be a collection of $C[0, 1]$ -valued random variables such that $\mathbb{P}(Z_0^{(n),\theta} = 0) = 1$ all $n \in \mathbb{N}$ and $\theta \in \Theta$. The following are equivalent:*

1. $(Z^{(n),\theta})$ is stochastically equicontinuous uniformly over Θ .
2. $(Z^{(n),\theta})$ is sequentially tight.

Proof. The equivalence is a straightforward application of Theorem 7.3 in Billingsley [2013]. Condition (i) of the aforementioned theorem is satisfied for any sequence of measures from the collection $(Z^{(n),\theta}(\mathbb{P}))_{n \in \mathbb{N}, \theta \in \Theta}$, since $Z_0^{(n),\theta} = 0$ almost surely for all n and θ . For any sequence $(\theta_n) \subseteq \Theta$, stochastic equicontinuity uniformly over Θ implies condition (ii) of Theorem 7.3 in Billingsley [2013] for the measures $((Z^{(n),\theta_n})(\mathbb{P}))$. We therefore conclude that stochastic equicontinuity uniformly over Θ implies sequential tightness.

On the contrary, assume that $(Z^{(n),\theta})$ is sequentially tight and let $\epsilon, \eta > 0$ be given. For each n , choose θ_n such that

$$\sup_{\theta \in \Theta} \mathbb{P}\left(\sup_{|s-t| \leq \delta} |Z_s^{(n),\theta} - Z_t^{(n),\theta}| \geq \epsilon\right) \leq \mathbb{P}\left(\sup_{|s-t| \leq \delta} |Z_s^{(n),\theta_n} - Z_t^{(n),\theta_n}| \geq \epsilon\right) + \frac{1}{n}.$$

Since $((Z^{(n),\theta_n})(\mathbb{P}))$ is tight by assumption, condition (ii) of Theorem 7.3 asserts that there exists $\delta, N > 0$ such that

$$\mathbb{P}\left(\sup_{|s-t| \leq \delta} |Z_s^{(n),\theta_n} - Z_t^{(n),\theta_n}| \geq \epsilon\right) < \eta$$

for $n \geq N$. Combining both inequalities and taking the limit superior finish the proof. \square

2.C. The Functional Martingale CLT

In this section we state Rebolledo's martingale CLT [Rebolledo, 1980] based on its formulation in Andersen et al. [1993], and then we extend the result to a uniform version without fixed variance functions. The one-dimensional case suffices for our purpose, so for simplicity, every local martingale in the following is a real-valued stochastic process. For a local square integrable martingale (M_t) , we let $\langle M \rangle(t)$ denote its quadratic characteristic. The theorem requires a condition on the jumps of the local martingales, for which we will need the following definition.

Definition 2.C.1. Let M_t be a local square integrable \mathcal{F}_t -martingale. For any $\varepsilon > 0$, we define $\langle M_\varepsilon \rangle(t)$ to be the quadratic characteristic of the pure jump-process given by

$$t \mapsto \sum_{0 \leq s \leq t} M_s \mathbb{1}(|\Delta M_s| > \varepsilon).$$

♣

We also need a representation of Gaussian martingales, which ensures their continuity.

Proposition 2.C.2. Let $(B_t)_{t \in [0, \infty)}$ be a Brownian motion on $[0, \infty)$ with continuous sample paths. For every non-decreasing $f \in C[0, 1]$, the process $(B_{f(t)})_{t \in [0, 1]}$ is a continuous mean zero Gaussian martingale on $[0, 1]$ with variance function f .

Consequently, if $U = (U_t)_{t \in [0, 1]}$ is a mean zero Gaussian martingale with a continuous variance function V , then U has the distributional representation

$$(U_t)_{t \in [0, 1]} \stackrel{\mathcal{D}}{=} (B_{V(t)})_{t \in [0, 1]}. \quad (2.C.1)$$

Proof. Let $f \in C[0, 1]$ be non-decreasing. From the properties of Brownian motion, it follows directly that the time-transformed process $(B_{f(t)})_{t \in [0, 1]}$ is a mean zero Gaussian process with variance function f . Since f is continuous, each sample path $t \mapsto B_{f(t)}$ is a composition of continuous functions and thus continuous itself. Since f is non-decreasing, the time-transformation also preserves the martingale property. This establishes the first part.

For the second part, recall that the covariance function of a martingale is determined by its variance function. Hence the first part implies that the right-hand side in (2.C.1) is a Gaussian process with the same mean and covariance structure as the left-hand side. Since the distribution of a Gaussian processes is uniquely determined by its mean and covariance structure, the equality in distribution follows. \square

Proposition 2.C.2 is a simple, distributional variant of the Dubins-Schwarz theorem, see Revuz and Yor [2013], Chapter V, Theorems 1.6 and 1.7. The Dubins-Schwarz theorem implies that, in fact, $U_t = B_{V(t)}$ for $t \in [0, 1]$, where B is a Brownian motion on $[0, V(1)]$. For the purpose of this work we only need the simpler, distributional equality (2.C.1).

We can now formulate Rebolledo's CLT for local martingales. To this end, note that Proposition 2.C.2 ensures the existence of the continuous Gaussian limit martingale U when the variance function V is continuous.

Theorem 2.C.3 (Rebolledo's CLT). *Let $(U^{(n)})_{n \in \mathbb{N}}$ be a sequence a local square integrable martingales in $D[0, 1]$, possibly defined on different sample spaces and with different filtrations for each $n \in \mathbb{N}$. Let U be a continuous Gaussian martingale with continuous variance function $V: [0, 1] \rightarrow [0, \infty)$, and assume that $U_0^{(n)} = U_0 = 0$. Suppose that for every $t \in [0, 1]$ and $\varepsilon > 0$,*

$$\langle U^{(n)} \rangle(t) \xrightarrow{P} V(t) \quad \text{and} \quad \langle U_\varepsilon^{(n)} \rangle(t) \xrightarrow{P} 0,$$

as $n \rightarrow \infty$. Then it holds that $U^{(n)} \xrightarrow{\mathcal{D}} U$ in $D[0, 1]$ as $n \rightarrow \infty$.

Proof. This is a special case of Theorem II.5.2 in Andersen et al. [1993]. \square

The general formulation of Rebolledo's CLT above, which allows for n -dependent sample spaces and filtrations, can now be leveraged to obtain a uniform version via the sequential characterization of uniform stochastic convergence.

Theorem 2.C.4 (Uniform Rebolledo CLT). *For each $n \in \mathbb{N}$ and $\theta \in \Theta$:*

- *Let $\mathcal{F}^{(n), \theta} = (\mathcal{F}_t^{(n), \theta})_{t \in [0, 1]}$ be a filtration satisfying the usual conditions.*
- *Let $U_t^{(n), \theta}$ be a local square integrable $\mathcal{F}_t^{(n), \theta}$ -martingale in $D[0, 1]$ with $U_0^{(n), \theta} = 0$.*
- *Let $V^\theta: [0, 1] \rightarrow [0, \infty)$ be a non-decreasing function with $V^\theta(0) = 0$.*

Assume that $(V^\theta)_{\theta \in \Theta}$ is uniformly equicontinuous and that $\sup_{\theta \in \Theta} V^\theta(1) < \infty$. Assume further that for every $\varepsilon > 0$ and $t \in [0, 1]$,

$$\langle U^{(n), \theta} \rangle(t) \xrightarrow{P/\theta} V^\theta(t) \quad \text{and} \quad \langle U_\varepsilon^{(n), \theta} \rangle(t) \xrightarrow{P/\theta} 0, \quad (2.C.2)$$

as $n \rightarrow \infty$. Then it holds that

$$U^{(n), \theta} \xrightarrow{\mathcal{D}/\theta} U^\theta, \quad n \rightarrow \infty,$$

in $D[0, 1]$ uniformly over Θ , where for each $\theta \in \Theta$, U^θ is a mean zero continuous Gaussian martingale on $[0, 1]$ with variance function V^θ .

Proof. We will use the characterization of uniform convergence as stated in Proposition 2.B.3 c). To this end, let $(\theta_n) \subseteq \Theta$ be an arbitrary sequence. By assumption $(V_{\theta_n})_{n \in \mathbb{N}}$ is a uniformly equicontinuous and bounded sequence of functions on a compact interval, so the Arzelà–Ascoli theorem states that there exists a subsequence $\theta_{k(n)}$, with $k: \mathbb{N} \rightarrow \mathbb{N}$ strictly increasing, and a function $\tilde{V} \in C[0, 1]$ such that

$$\sup_{t \in [0, 1]} |V^{\theta_{k(n)}}(t) - \tilde{V}(t)| \longrightarrow 0, \quad n \rightarrow \infty.$$

Since each function $V^{\theta_{k(n)}}$ is non-decreasing, it follows that \tilde{V} is non-decreasing. It also holds that $\tilde{V}(0) = \lim_{n \rightarrow \infty} V^{\theta_{k(n)}}(0) = 0$, and therefore \tilde{V} is the variance function of a continuous Gaussian martingale \tilde{U} with $\tilde{U}_0 = 0$.

2 Nonparametric conditional local independence testing

By assumption of the convergences in (2.C.2), we may conclude that

$$|\langle U^{(k(n)), \theta_{k(n)}} \rangle(t) - \tilde{V}(t)| \leq \underbrace{|\langle U^{(k(n)), \theta_{k(n)}} \rangle(t) - V^{\theta_{k(n)}}(t)|}_{\xrightarrow{P} 0} + \underbrace{|V^{\theta_{k(n)}}(t) - \tilde{V}(t)|}_{\rightarrow 0} \xrightarrow{P} 0$$

and that $\langle U_\epsilon^{(k(n)), \theta_{k(n)}} \rangle(t) \rightarrow 0$ as $n \rightarrow \infty$. Thus we have established the conditions of the classical Rebolledo CLT – Theorem 2.C.3 – for the sequence $U^{(k(n)), \theta_{k(n)}}$ and the Gaussian martingale \tilde{U} with variance function \tilde{V} . We therefore conclude that

$$U^{(k(n)), \theta_{k(n)}} \xrightarrow{\mathcal{D}} \tilde{U}$$

in $D[0, 1]$ as $n \rightarrow \infty$.

We now establish that the sequence $(U^{\theta_{k(n)}})$ also converges in distribution to \tilde{U} in $C[0, 1]$, and in particular also in $D[0, 1]$. To this end, we use the characterization of convergence in distribution in $C[0, 1]$ from Theorem 7.5 in Billingsley [2013], which states that we need to show that

1. For all $0 \leq t_1 < \dots < t_m \leq 1$, it holds that

$$(U_{t_1}^{\theta_{k(n)}}, \dots, U_{t_m}^{\theta_{k(n)}}) \xrightarrow{\mathcal{D}} (\tilde{U}_{t_1}, \dots, \tilde{U}_{t_m}), \quad n \rightarrow \infty.$$

2. For all $\epsilon > 0$

$$\lim_{\delta \rightarrow 0^+} \limsup_{n \rightarrow \infty} \mathbb{P} \left(\sup_{|t-s| < \delta} |U_t^{\theta_{k(n)}} - U_s^{\theta_{k(n)}}| > \epsilon \right) = 0.$$

The first condition is clear since all the marginals are multivariate Gaussian, and the mean and variance of the sequence converges to the mean and variance of the limit distribution. The second condition follows from the same computation as in the proof of Lemma 2.A.11. By Theorem 7.5 in Billingsley [2013] we therefore conclude that

$$U^{\theta_{k(n)}} \xrightarrow{\mathcal{D}} \tilde{U}, \quad \text{for } n \rightarrow \infty,$$

in $C[0, 1]$, and hence also in $D[0, 1]$.

We can now apply the triangle inequality for the bounded Lipschitz metric to conclude that

$$d_{BL}(U^{(k(n)), \theta_{k(n)}}, U^{\theta_{k(n)}}) \leq d_{BL}(U^{(k(n)), \theta_{k(n)}}, \tilde{U}) + d_{BL}(\tilde{U}, U^{\theta_{k(n)}}) \longrightarrow 0.$$

Since $(\theta_n) \subseteq \Theta$ was an arbitrary sequence, we conclude that $U^{(n), \theta} \xrightarrow{\mathcal{D}/\Theta} U^\theta$ by Proposition 2.B.3. \square

The following proposition gives explicit expressions for the quadratic characteristics that appear in Rebolledo's CLT in the special case where the local martingales are given as stochastic integrals with respect to a compensated counting processes.

Proposition 2.C.5. *Let N_1, \dots, N_n be counting processes and assume that for each $j = 1, \dots, n$, N_j has an absolutely continuous $\mathcal{F}_t^{(n)}$ -compensator $\Lambda_{j,t}$ such that $M_{j,t} = N_{j,t} - \Lambda_{j,t}$ is a locally square integrable $\mathcal{F}_t^{(n)}$ -martingale. Let H_1, \dots, H_n be locally bounded $\mathcal{F}_t^{(n)}$ -predictable processes, and define the process $U_t^{(n)} = \sum_{j=1}^n \int_0^t H_{j,s} dM_{j,s}$. Then $U_t^{(n)}$ is a local square integrable $\mathcal{F}_t^{(n)}$ -martingale, and for any $t, \epsilon > 0$ it holds that*

$$\begin{aligned} \langle U^{(n)} \rangle(t) &= \sum_{j=1}^n \int_0^t H_{j,s}^2 d\Lambda_{j,s}, \\ \langle U_\epsilon^{(n)} \rangle(t) &= \sum_{j=1}^n \int_0^t H_{j,s}^2 \mathbb{1}(|H_{j,s}| \geq \epsilon) d\Lambda_{j,s}. \end{aligned}$$

Proof. See the discussion following Theorem II.5.2 in Andersen et al. [1993], in particular equations (2.5.6) and (2.5.8). \square

2.D. Estimation of λ and G

The asymptotic theory for estimation of the LCM crucially relies on $\hat{\lambda}^{(n)}$ and $\hat{G}^{(n)}$ being consistent, and more importantly, having a product error decaying at an $n^{-1/2}$ -rate. Therefore, a central question when applying the test, is how to model λ and G .

In principle, we could use parametric models to learn $\hat{\lambda}^{(n)}$ and $\hat{G}^{(n)}$, and under such models it should be possible to achieve $n^{-1/2}$ -rates. For example, if we consider a parametrization $(t, \theta) \mapsto \lambda_t(\theta)$ which is $\kappa(t)$ -Lipschitz in $\theta \in \Theta \subseteq \mathbb{R}^p$ for each t , then

$$h(n)^2 = E \left(\int_0^1 (\lambda_t(\theta_0) - \lambda_t(\hat{\theta}^{(n)}))^2 dt \right) \leq \|\kappa\|_{L_2([0,1])}^2 E \|\theta_0 - \hat{\theta}^{(n)}\|_{\mathbb{R}^p}^2.$$

Thus the rates from parametric asymptotic theory can be converted to rates for g and h .

However, it is of greater interest if sufficient rates can be achieved with nonparametric estimators. Below we give concrete examples of nonparametric models and discuss which rates are achievable. For simplicity, we focus on the case where $\mathcal{F}_t = \mathcal{F}_t^{N,Z}$ and where $G_t = X_t - \Pi_t$ as in the introductory example.

2.D.1 Nonparametric functional estimation of Π

As seen in Section 2.6.1, assumptions on the form of Π turn the general estimation problem into a concrete problem of estimating a function.

If the system is Markovian, it can be reasonable to assume a *functional concurrent model*. The model asserts that $\Pi_t = \mu(t, Z_t)$ for a bivariate function μ , and a survey of methods for estimating μ is given by Maity [2017]. Notably, Jiang and Wang [2011] achieve an $n^{-1/3}$ -rate of $g(n)$ under certain regularity and moment assumptions, see their Theorem 3.3. That result also holds if Z is replaced by a linear predictor $\beta^T \mathbf{Z}$ of several covariates.

2 Nonparametric conditional local independence testing

Consider again the historical linear regression model from Section 2.6.1, and assume that the effect of Z on X is homogeneous over time. That is, $\rho_X(s, t) = \tilde{\rho}_X(t - s)$ for some function $\tilde{\rho}_X$. This submodel is known as the *functional convolution model*, since Π can be written as the convolution of Z and $\tilde{\rho}_X$. Applying the Fourier transform converts it into a (complex) linear concurrent model, so by Plancherel's theorem one can leverage the convergence rates from the concurrent model. Manrique [2016, Theorem 16] uses this idea to transfer the $n^{-1/4}$ -rate of the *functional ridge regression estimator* [Manrique et al., 2018] to the convolution model, which holds under modest moment conditions on the data. With additional distributional assumptions, we conjecture that faster rate results for the linear concurrent model can also be leveraged to the convolution model. Şentürk and Müller [2010] consider a similar model under the assumption that

$$\rho_X(s, t) = \mathbb{1}(t - \Delta \leq s \leq t) \tilde{\rho}_X^1(t) \tilde{\rho}_X^2(t - s),$$

for two functions $\tilde{\rho}_X^1$ and $\tilde{\rho}_X^2$ and a lag $\Delta > 0$. They establish a pointwise rate result for the response curve, but it is not obvious how to cast their result as a polynomial rate for $g(n)$.

For the full historical functional linear model we are not aware of any published rate results. Yuan and Cai [2010], Cai and Yuan [2012] establish rates on the prediction error for scalar-on-function regression, and Yao et al. [2005] establish various rates for function-on-function regression, but in a non-historical setting. Based on the former, we give a heuristic for which rates are achievable for $g(n)$ in this model. If $\hat{\Pi}$ is based on a kernel estimate $\hat{\rho}_X^{(n)}$ of ρ_X , then Tonelli's theorem yields

$$\begin{aligned} g(n)^2 &= \left\| \Pi - \hat{\Pi}^{(n)} \right\|_2^2 = \mathbb{E} \left(\int_0^1 \left(\int_0^t (\rho_X(s, t) - \hat{\rho}_X^{(n)}(s, t)) Z_s ds \right)^2 dt \right) \\ &= \int_0^1 \mathbb{E} \left(\left(\int_0^t (\rho_X(s, t) - \hat{\rho}_X^{(n)}(s, t)) Z_s ds \right)^2 \right) dt. \end{aligned}$$

Theorem 4 in Cai and Yuan [2012] asserts that we, under certain regularity conditions, can estimate $\rho_X(\cdot, t)$ such that

$$\mathbb{E} \left(\left(\int_0^t (\rho_X(s, t) - \hat{\rho}_X^{(n)}(s, t)) Z_s ds \right)^2 \right)$$

decays at a $n^{-2r_t/(2r_t+1)}$ -rate for a fixed t . Here r_t is a constant describing the eigenvalue decay of a certain operator related to the autocovariance of Z and the regularity of ρ_X . As a concrete example, if Z is a Wiener process and $\rho_X(\cdot, t) \in \mathcal{W}_2^m([0, t])$ is in the m -th Sobolev space for each $t > 0$, then $r_t = 1 + m$ and $g(n)$ will converge at an $n^{-(1+m)/(2m+3)}$ -rate, see the discussion after Corollary 8 in Yuan and Cai [2010]. Based on these arguments, we believe that the desired $n^{-(1/4+\varepsilon)}$ -rate for $g(n)$ is achievable with suitable regularity assumptions on Z and ρ_X .

2.D.2 Estimation of λ

Within the framework of the Cox model, Wells [1994] demonstrate that the baseline intensity can be estimated with rate $n^{-2/5}$ using a standard kernel smoothing technique. With the parametric $n^{-1/2}$ -rate on the remaining parameters, this translates readily into $h(n) = O(n^{-2/5})$.

We suspect that the same rate should also be attainable in a sparse setting with high-dimensional covariates, for example by applying the smoothing approach of Wells [1994] to the baseline hazard estimators of e.g. Fang et al. [2017] and Hou et al. [2023].

Hiabu et al. [2021] consider the more general multiplicative intensity model with $\lambda_t = \mathbb{1}(T \geq t)f(t, Z_t)$, where f has a multiplicative structure over its arguments. They introduce an estimator with optimal rate $h(n) = n^{-2/(5+d)}$, where d is the dimension of Z . For $d > 3$, we therefore need faster rates on $g(n)$ in order for the LCT to maintain type I error control.

Omitting the multiplicative structure on f , Bender et al. [2020] propose a general framework for nonparametric estimation of Markovian intensities. They survey existing methods such as gradient boosted trees and neural networks and relate them to this setting. Based on real and synthetic data, they find that both gradient boosted trees and neural networks outperform the Cox model in terms of predictive performance as measured by the Brier score. In essence, the framework relies on discretizing time and approximating the intensity with successive Poisson regressions. Using the same idea, Rytgaard et al. [2021] argue that $h(n) = o(n^{-1/4})$ can be achieved for time-independent covariates.

Similarly, Rytgaard et al. [2022] mention that $h(n) = o(n^{-1/4})$ can be achieved for estimation of intensities in a multivariate point process with a uniformly bounded number of events, which we place into a general modeling framework below.

2.D.3 Estimation of λ and Π for counting processes

In Sections 2.2 and 2.4 we considered the setup where N was a counting process adapted to a filtration \mathcal{F}_t , which could contain information on baseline covariates and covariate processes that were not necessarily counting processes. In this section we explore how our testing framework can be applied when all stochastic processes of interest are counting processes.

More specifically, let $(N_t^d)_{d \in [p]}$ be a p -dimensional counting process. For $a, b \in [p]$ and $C \subset [p] \setminus \{b\}$ with $a \neq b$ and $a \in C$ we are interested in testing the hypothesis that N^a is conditionally locally independent of N^b given the filtration, \mathcal{F}_t^C , generated by $N^C = (N_t^d)_{d \in C}$.

We can cast this setup in the framework of Section 2.2 as follows. Naturally, we let $N = N^a$ and $\mathcal{F}_t = \mathcal{F}_t^C$. The auxiliary process X is chosen to be càglàd and predictable with respect to the filtration, \mathcal{F}_t^b , generated by N^b . For example, we could choose $X_t = N_{t-}^b$. But X_t could be any functional of N^b such as $X_t = f(N_{t-}^b)$ for a suitable

2 Nonparametric conditional local independence testing

function f or a linear filter of N^b ,

$$X_t = \int_0^{t-} \kappa(t-s) dN_s^b,$$

where κ is a suitable kernel function, see also Section 2.3.1. In principle, the process X could also depend on the process N^C , but it is important that the filtration, \mathcal{G}_t , generated by \mathcal{F}_t and X_t is strictly larger than \mathcal{F}_t , i.e., X_t should depend on N^b , in order to get a non-trivial test as explained in Section 2.2.

In the framework of counting processes, we can approach the estimation of both λ and Π in a unified and general way as follows: Let $(\tau_j, z_j)_{j \geq 1}$ be the marked point process associated with the counting process N^C , i.e., $(\tau_j)_{j \geq 1}$ is a sequence of almost surely strictly increasing event times located at the jumps of N^C , and $(z_j)_{j \geq 1}$ for $z_j \in C$ are the corresponding event types.

Since both λ_t and Π_t are real-valued and \mathcal{F}_{t-}^C -measurable for each fixed $t \geq 0$, they can be represented as measurable functions of $\{(\tau_j, z_j) \mid \tau_j < t, z_j \in C\}$. Hence, we can model both λ and Π using any sequence-to-number model. For the intensity process, Rytgaard et al. [2022] propose a sequence of HAL estimators when the total event count is uniformly bounded. As an alternative, Xiao et al. [2019] propose using a recurrent neural network (LSTM). Unless there is a uniform bound on the total number of events, as assumed by Rytgaard et al. [2022], there are currently no published results available on the rates of convergence for nonparametric estimation of sequence-to-number functions.

2.E. Relation to semiparametric survival models

In this section, we relate the LCM to existing work on treatment effects in survival analysis. We resume to the setting of Section 2.3.2, that is, the case where $N_t = \mathbb{1}(T \geq t)$ is the counting process of a survival time, X is a baseline treatment variable, and where $\mathcal{F}_t = \sigma(N_s, Z; s \leq t)$ for additional baseline covariates Z . Supposing that X is also non-negative, we may consider two different models for the intensity:

$$\lambda_t^\times = \mathbb{1}(T \geq t) \underline{\lambda}(t) \exp(\theta X + \varphi(Z)), \quad (2.E.1)$$

$$\lambda_t^+ = \mathbb{1}(T \geq t) (\vartheta X + \phi(t, Z)), \quad (2.E.2)$$

where $\theta, \vartheta \in \mathbb{R}$ are treatment parameters of interest, and where $\underline{\lambda}, \varphi, \phi$ deterministic nuisance functions. The model in (2.E.1) is known as the partially linear Cox model (PLCM, Sasieni [1992]), and the additive model in (2.E.2) was considered by Dukes et al. [2019] among others.

While the parameters θ and ϑ are difficult to compare directly, the hypothesis of conditional local independence corresponds to the hypothesis of zero treatment effect within each of the models, and testing this hypothesis can be done using a score test.

Sasieni [1992] shows that within the PLCM, the efficient score for θ is given by

$$S^\times(\theta; \underline{\lambda}, \varphi) = \int_0^1 (X - \alpha^*(t) - h^*(Z)) (dN_t - \mathbb{1}(T \geq t) \underline{\lambda}(t) e^{\theta X + \varphi(Z)} dt), \quad (2.E.3)$$

where (α^*, h^*) are defined as the minimizers $\mathbb{E}(X - \alpha(T) - h(Z))^2$. Recall that, when X and Z are time-independent, the null hypothesis H_0 of conditional local independence reduces to the conditional independence statement $X \perp\!\!\!\perp T \mid Z$. Consequently, it holds that $\alpha^*(T) = 0$ and $h^*(Z) = \mathbb{E}[X \mid Z] = \Pi_0$ under H_0 . Evaluating the efficient score at $\theta = 0$ under H_0 therefore gives

$$S^\times(0; \underline{\lambda}, \varphi) = \int_0^1 (X - \Pi_0)(dN_t - \mathbb{1}(T \geq t)\underline{\lambda}(t)e^{\varphi(Z)}dt) = (X - \Pi_0)(1 - \Lambda_T(\underline{\lambda}, \varphi)).$$

We see that the empirical version of S^\times is exactly the endpoint of the LCM estimator with additive residual process (cf. Equation 2.3.16). This means that our test, the X-LCT, can under the PLCM be interpreted as a score test based on the efficient score.

A similar connection can be made for the additive model. Dukes et al. [2019] show that the efficient score for ϑ is given by

$$S^+(\vartheta; \phi) = \int_0^1 \left(X - \frac{\mathbb{E}[(\vartheta X + \phi(t, Z))^{-1} X e^{-\vartheta X t} \mid Z]}{\mathbb{E}[(\vartheta X + \phi(t, Z))^{-1} e^{-\vartheta X t} \mid Z]} \right) \frac{(dN_t - \lambda_t(\vartheta, \phi)dt)}{(\vartheta X + \phi(t, Z))}.$$

Plugging in $\vartheta = 0$ and simplifying under H_0 yields

$$S^+(0; \phi) = \int_0^1 \frac{(X - \mathbb{E}[X \mid Z])}{\phi(t, Z)} (dN_t - \mathbb{1}(T \geq t)\phi(t, Z)dt) = \int_0^1 \frac{(X - \Pi_0)}{\lambda_t} (dN_t - \lambda_t dt).$$

We recognize the empirical version of S^+ as the endpoint of the LCM estimator with the time-constant X replaced by the hazard weighted process X/λ_t .

Other works that consider effect estimation based on orthogonal scores include Huang [1999], Fang et al. [2017], Niu et al. [2022], Zhong et al. [2022] for the PLCM and Hou et al. [2023] for the additive model in (2.E.2).

We also suspect, as the derivations in Section 2.3.1 likewise suggest, that the LCM is still an efficient score for certain semiparametric survival models even when the covariates vary with time, but we are not aware of existing results on such a connection.

2.F. Details on Neyman orthogonality

In this section, we first show by direct computation that the LCM is Neyman orthogonal with respect to both general residual processes and intensities. We then show that the LCM with an additive residual process can be viewed as a concentrated-out score in the sense of Newey [1994].

2.F.1 General Neyman orthogonality

The definition of Neyman orthogonality by Chernozhukov et al. [2018, Def. 2.1.] requires that we formally define function spaces for the collections of nuisance parameters. However, to avoid extensive technical specifications (and redundant model assumptions), we prove a simpler – but more general – condition, from which Neyman orthogonality can be derived within specific semiparametric models.

2 Nonparametric conditional local independence testing

First, we generalize the integral I_t from Definition 2.2.4 to a function of pairs (x, y) of càglàd functions, given by

$$I_t(x, y) = \int_0^t x_s(dN_s - y_s ds).$$

With this notation, the LCM is given by $\gamma_t = \mathbb{E}[I_t(G, \lambda)]$, where G is a residual process and λ is the \mathcal{F}_t -intensity of N . We assume for simplicity that λ and G are bounded such that the expectation is well-defined.

Now let \tilde{G}_t be an arbitrary bounded \mathcal{G}_t -predictable càglàd process, and let $\tilde{\lambda}_t$ be an arbitrary bounded \mathcal{F}_t -predictable càglàd process. We establish the following orthogonality condition: under H_0 it holds that

$$\partial_r \mathbb{E}[I_t(G + r(\tilde{G} - G), \lambda + r(\tilde{\lambda} - \lambda))]_{|r=0} = 0. \quad (2.F.1)$$

Indeed, observe that

$$\begin{aligned} I_t(G + r(\tilde{G} - G), \lambda + r(\tilde{\lambda} - \lambda)) \\ = (1-r)^2 I_t(G, \lambda) + r(1-r)(I_t(G, \tilde{\lambda}) + I_t(\tilde{G}, \lambda)) + r^2 I_t(\tilde{G}, \tilde{\lambda}), \end{aligned}$$

from which it follows that

$$\partial_r \mathbb{E}[I_t(G + r(\tilde{G} - G), \lambda + r(\tilde{\lambda} - \lambda))]_{|r=0} = 2 \cdot \mathbb{E}[I_t(G, \lambda)] + \mathbb{E}[I_t(G, \tilde{\lambda})] + \mathbb{E}[I_t(\tilde{G}, \lambda)].$$

The first term is zero under H_0 by Proposition 2.2.5, and the third term vanishes under H_0 by the same argument. For the second term, we note that

$$\mathbb{E}[I_t(G, \tilde{\lambda})] = \gamma_t + \mathbb{E} \left[\int_0^t G_s(\lambda_s - \tilde{\lambda}_s) ds \right] = \gamma_t + \int_0^t \mathbb{E} \left[\mathbb{E}[G_s | \mathcal{F}_{s-}] (\lambda_s - \tilde{\lambda}_s) \right] ds = \gamma_t,$$

which also vanishes under H_0 . This lets us conclude that (2.F.1) holds under H_0 .

2.F.2 Concentrating-out

To derive the concentrated-out score, we first need to formalize the nuisance parameters and the collection thereof. We consider the case where $\mathcal{F}_t = \mathcal{F}_t^{N, Z}$ for a process $Z = (Z_t)$, and let D_X and D_Z denote the respective sample spaces of the X and Z . We posit the following semiparametric model for the intensity:

$$\lambda_t(\beta, h) = \mathbb{1}(T \geq t) e^{\beta X_t} h(t, Z),$$

where $h: [0, 1] \times D_Z \rightarrow [0, \infty)$ is a function such that $t \mapsto h(t, Z)$ is an \mathcal{F}_t^Z -predictable càglàd process. Denote the collection of such functions by \mathcal{T}_1 . To make the space not dependent on the particular instantiation of the process Z , we could also take it to be the set:

$$\left\{ (h_t)_{t \in [0, 1]} \in \int_{[0, 1]}^{\oplus} L_2(D_Z |_{[0, t]}, \mathbb{R}) dt \mid t \mapsto h_t((z_s)_{s < t}) \text{ is a bounded nonnegative} \right. \\ \left. \text{càglàd function for all } (z_s)_{0 \leq s \leq 1} \in D_Z \right\},$$

where \int^\oplus denotes the direct integral, and where $D_Z|_{[0,t]}$ is the path space D_Z restricted to the domain $[0,t]$.

Recall that the likelihood at $t = 1$ is given by

$$\ell_1(\beta, h) = \int_0^1 \log(\lambda_s(\beta, h)) dN_s - \int_0^1 \lambda_s(\beta, h) ds.$$

We show that concentrating out the nuisance parameter h of $\ell_1(\beta, h)$ yields to the LCM with additive residual process under H_0 .

Let P_0 be fixed distribution satisfying H_0 with ground truth $\beta^0 = 0$ and $h^0 \in \mathcal{T}_1$. Suppose that for each $\beta \in (-\epsilon, \epsilon)$ in a neighborhood of zero, there is a function $\varpi_\beta \in \mathcal{T}_1$ such that

$$\varpi_\beta(t, Z) = \mathbb{E}_{P_0}[e^{\beta X_t} \mid T \geq t, \mathcal{F}_{t-}^Z].$$

The first step of concentrating-out is to maximize the expected likelihood over h . We claim that for each fixed β , the function $h_\beta^* := h^0 / \varpi_\beta \in \mathcal{T}_1$ maximizes the objective $u(h) := \mathbb{E}_{P_0}[\ell_1(\beta, h)]$ over $h \in \mathcal{T}_1$.

Since the logarithm is concave and $\lambda(\beta, h)$ is linear in h , it follows that u is a concave objective. Moreover, one can show that h_β^* is the unique critical point of u by equating its Gateaux derivative to zero and invoking the fundamental lemma of the calculus of variations. For simplicity, we settle with verifying (see below) that the Gateaux derivative is zero at h_β^* . From these properties we may conclude that the global maximum of u is attained at h_β^* .

A straightforward differentiation shows that for any $\tilde{h} \in \mathcal{T}_1$,

$$\partial_r \ell_1(\beta, h_\beta^* + r \cdot \tilde{h}) \big|_{r=0} = \int_0^1 \frac{\tilde{h}(s, Z)}{h_\beta^*(s, Z)} dN_s - \int_0^1 \lambda_s(\beta, \tilde{h}) ds.$$

Because $\frac{\tilde{h}(t, Z)}{h_\beta^*(t, Z)}$ is \mathcal{G}_t -predictable and $N_t - \int_0^t \lambda_s(0, h^0) ds$ is a \mathcal{G}_t -martingale, taking expectation under P_0 yields that

$$\begin{aligned} \partial_r u(h_\beta^* + r \cdot \tilde{h}) \big|_{r=0} &= \mathbb{E}[\partial_r \ell_1(\beta, h_\beta^* + r \cdot \tilde{h}) \big|_{r=0}] \\ &= \mathbb{E} \left[\int_0^1 \frac{\tilde{h}(s, Z)}{h_\beta^*(s, Z)} \lambda_s(0, h^0) ds - \int_0^1 \lambda_s(\beta, \tilde{h}) ds \right] \\ &= \mathbb{E} \left[\int_0^1 \mathbb{1}(T \geq s) \left(\varpi_\beta(t, Z) - e^{\beta X_s} \right) \tilde{h}(s, Z) ds \right] \\ &= \int_0^1 \mathbb{E} \left[\mathbb{1}(T \geq s) \left(\varpi_\beta(t, Z) - \mathbb{E}[e^{\beta X_s} \mid T \geq t, \mathcal{F}_{s-}^Z] \right) \tilde{h}(s, Z) \right] ds = 0, \end{aligned}$$

where we have used that $\tilde{h}(s, Z)$ is \mathcal{F}_s -predictable.

Now, by the method of concentrating-out, we are led to consider nuisance functions of the form

$$\begin{aligned} \eta_h &: (-\epsilon, \epsilon) \longrightarrow \mathcal{T}_1, \\ \eta_h(\beta) &= \left((t, z) \mapsto \frac{h(t, z)}{\varpi_\beta(t, z)} \right), \end{aligned}$$

2 Nonparametric conditional local independence testing

for any $h \in \mathcal{T}_1$, which in particular includes $\eta_{h_0}(\beta) = h_\beta^*$. The resulting concentrating-out score is therefore

$$\begin{aligned} \psi(\beta, \eta_h) &= \frac{\partial \ell_1(\beta, \eta_h(\beta))}{\partial \beta} \\ &= \int_0^1 \left(X_s - \frac{\partial_\beta \varpi_\beta(t, Z)}{\varpi_\beta(t, Z)} \right) dN_s - \int_0^1 \left(\frac{X_s}{\varpi_\beta(t, Z)} - \frac{\partial_\beta \varpi_\beta(t, Z)}{\varpi_\beta(t, Z)^2} \right) \lambda(\beta, h) ds \\ &= \int_0^1 \left(X_s - \frac{\mathbb{E}_{P_0}[X_s e^{\beta X_s} \mid T \geq s, \mathcal{F}_{s-}^Z]}{\mathbb{E}[e^{\beta X_s} \mid T \geq s, \mathcal{F}_{s-}^Z]} \right) dN_s \\ &\quad - \int_0^1 \left(\frac{X_s}{\mathbb{E}_{P_0}[e^{\beta X_s} \mid T \geq s, \mathcal{F}_{s-}^Z]} - \frac{\mathbb{E}_{P_0}[X_s e^{\beta X_s} \mid T \geq s, \mathcal{F}_{s-}^Z]}{(\mathbb{E}_{P_0}[e^{\beta X_s} \mid T \geq s, \mathcal{F}_{s-}^Z])^2} \right) \lambda(\beta, h) ds. \end{aligned}$$

After plugging in $\beta = 0$ and simplifying we see that

$$\psi(0, \eta_h) = \int_0^1 (X_s - \mathbb{E}[X_s \mid \mathcal{F}_{s-}]) (dN_s - \lambda(0, h) ds).$$

This shows that the concentrating-out score evaluated at $\beta = 0$ is exactly the score for the endpoint of the LCM estimator.

2.G. Additional details of simulation study

This section contains additional details, numerical results, and figures related to the simulations of Section 2.6.

2.G.1 Implementation of estimators and tests

For our proof-of-concept implementation we used two simple off-the-shelf estimators.

To estimate λ we used the `BoXHED2.0` estimator from Pakbin et al. [2021], based on the works of Wang et al. [2020] and Lee et al. [2021]. In essence, the estimator is a gradient boosted forest adapted to the setting of hazard estimation with time-dependent covariates. The maximum depth and number of trees were tuned by 5-fold cross-validation over the same grid as in Pakbin et al. [2021]. For computational ease, the hyperparameters were tuned once on the entire dataset instead of tuning them on each fold J_n^k . In principle, this may invalidate the asymptotic properties of $\check{\Psi}_n^K$ since it breaks the independence between $\hat{\lambda}^{k,(n)}$ and $(T_j, X_j, Z_j)_{j \in J_n^k}$, but we believe that this dependency is negligible.

To estimate the predictable projection $\Pi_t = \mathbb{E}(X_t \mid \mathcal{F}_{t-})$, we fitted a series of linear least squares estimators by regressing X_t on $(Z_s)_{s \in \mathbb{T}: s < t}$ for each $t \in \mathbb{T}$. To stabilize the estimation error $g(n)$, we added a small L_2 -penalty with coefficient 0.001 fixed across all experiments for simplicity. Since X_t was sampled from a discretized historical linear model, the error $g(n)$ should in principle converge with a classical $n^{-1/2}$ -rate. The historical linear regression estimator from the `scikit-fda` library was also considered initially, but we found that fitting this model was too computationally expensive for




Figure 2.G.1: Empirical distribution functions of p -values for the three different conditional local independence tests considered, simulated under the sampling scheme described in Section 2.6. The dotted line shows $y = x$ corresponding to a uniform distribution.

2 Nonparametric conditional local independence testing

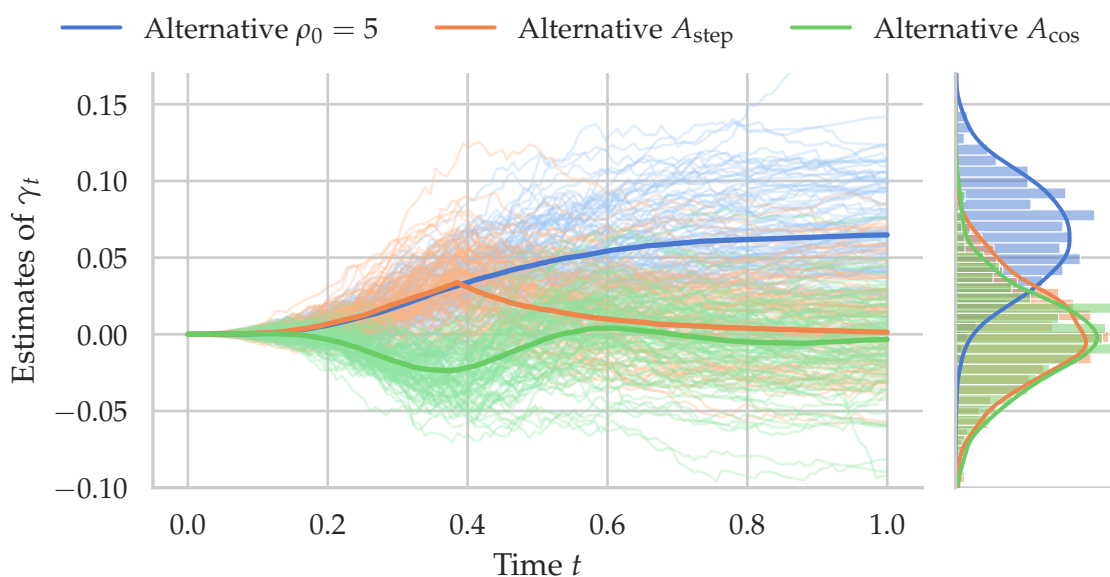


Figure 2.G.2: Sample paths of $\gamma_t^{K,(500)}$ fitted on data sampled from three different alternatives as described in Section 2.6.4. Here (X, Y, Z) are sampled from the scheme described in Section 2.6, with both ρ_X and ρ_Y being the constant kernel and with $\beta = -1$. For each alternative, 100 paths are shown. The empirical mean functions and the endpoint distributions are highlighted and computed based on 500 samples.

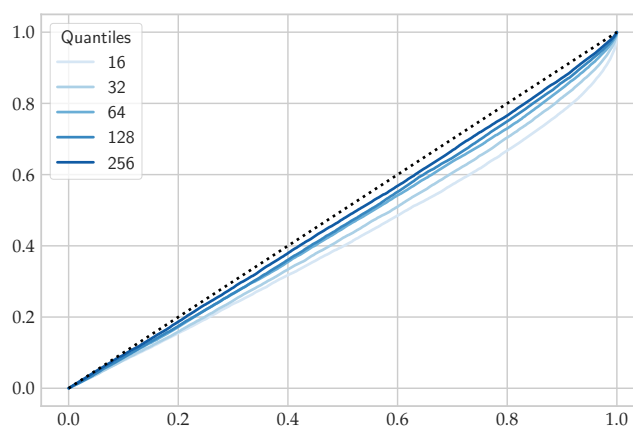


Figure 2.G.3: Empirical distribution functions of $p^{(q)} = 1 - F_S(M^{(q)})$, where $M^{(q)} = (M_t^{(q)})_{t=1,\dots,q}$ is a random walk with Gaussian increments such that $M_q^{(q)}$ has unit variance for each $q \in \{2^\ell: \ell = 4, \dots, 8\}$. Each empirical distribution function is based on $N = 20\,000$ samples.

a simulation study with cross-fitting. In principle, in our time-continuous setting, we would like to use a functional estimator of Π that would utilize the regularity along s and t . Initial experiments, however, suggested that the simpler historical regression described above gave similar results as using the `scikit-fda` library, and we went with the less time consuming implementation.

Based on these estimators, the X-LCT was implemented based on Algorithm 2. Following the recommendation by Chernozhukov et al. [2018, Remark 3.1.], we computed the X-LCT with $K = 5$ folds. The associated p -value was computed with the series representation of F_S truncated to the first 1000 terms.

We compared our results for X-LCT with a hazard ratio test in the possibly misspecified marginal Cox model given by (2.2.11). This test was computed using the `lifelines` library [Davidson-Pilon, 2019], specifically the `CoxTimeVaryingFitter` model. The model was fitted with an L_2 -penalty with a coefficient set to 0.1 (the default), and as a consequence the hazard ratio test is expected to be conservative.

2.G.2 Comparison with endpoint statistic

We compare the X-LCT, which is based on the uniform norm of the X-LCM, with its endpoint counterpart. More precisely, we consider the test statistic

$$(\check{\mathcal{V}}_{K,n}(1))^{-\frac{1}{2}} \sqrt{n} \check{\gamma}_1^{K,(n)},$$

which is asymptotically standard normal under H_0 . With the simulation settings in Section 2.6.4, the X-LCT turns out to be more or less indistinguishable from the corresponding endpoint test. This is because the alternatives considered have corresponding LCMs, which are most extreme towards $t = 1$. Therefore, the supremum and the endpoint behave similarly in these cases.

For this reason we consider local alternatives that result in a non-monotonic LCM. Using the same expression for the intensity (2.6.39), but with a time-varying ρ_0 , we consider the alternatives

$$\begin{aligned} A_{\text{step}}: \rho_0(t) &= 5 \cdot \mathbf{1}(t \leq 0.4) - 5 \cdot \mathbf{1}(t > 0.4), \\ A_{\text{cos}}: \rho_0(t) &= 7 \cdot \cos(4\pi \cdot t). \end{aligned}$$

The idea behind the alternative A_{step} is that the LCM should be increasing on $[0, 0.4]$ and decreasing on $(0.4, 1]$. Figure 2.G.2 shows sample paths of $\check{\gamma}^{K,(n)}$ for data simulated under each of the alternatives $\rho_0 = 5$, A_{step} and A_{cos} . The figure illustrates that $t \mapsto |\check{\gamma}_t^{K,(n)}|$ is, indeed, mostly maximal towards $t = 1$ for the alternative $\rho_0 = 5$, but not for the time-varying alternatives A_{step} and A_{cos} .

With the same sampling scheme for (X, Y, Z) as in Section 2.6.2, we conducted an analogous experiment with 400 runs for each setting. Figure 2.G.4 shows the rejection rates for the two tests.

Under the hypothesis of conditional local independence, the left plot in Figure 2.G.4 shows that the endpoint test behaves similarly to $\check{\Psi}_n^K$ as expected. Both tests have power against the local alternatives, but for A_{step} the power does not seem to stabilize before

2 Nonparametric conditional local independence testing

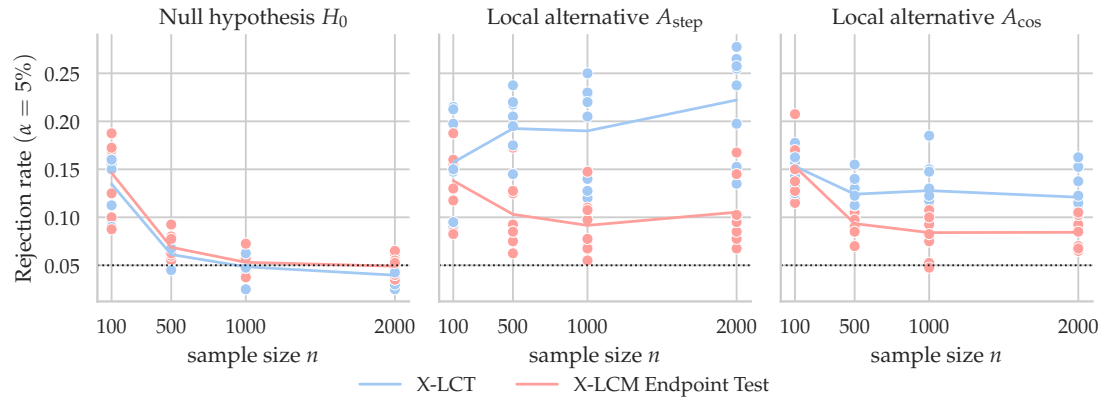


Figure 2.G.4: The plots show the average rejection rate of the double machine learning tests based on the supremum statistic (blue) and the endpoint statistic (red).

$n = 2000$. This is different from the previous settings, and can be explained by a slower convergence of the intensity estimator due to the more complex dependency on X . For both of the local alternatives, we observe that $\check{\Psi}_n^K$ is more powerful than the endpoint test, with the difference being largest for A_{step} . In conclusion, these results show that the supremum test dominates the endpoint test in certain situations.

3 Efficient adjustment for complex covariates: Gaining efficiency with DOPE

ALEXANDER MANGULAD CHRISTGAU AND NIELS RICHARD HANSEN

Abstract

Covariate adjustment is a ubiquitous method used to estimate the average treatment effect (ATE) from observational data. Assuming a known graphical structure of the data generating model, recent results give graphical criteria for optimal adjustment, which enables efficient estimation of the ATE. However, graphical approaches are challenging for high-dimensional and complex data, and it is not straightforward to specify a meaningful graphical model of non-Euclidean data such as texts. We propose a general framework that accommodates adjustment for any subset of information expressed by the covariates. We generalize prior works and leverage these results to identify the optimal covariate information for efficient adjustment. This information is minimally sufficient for prediction of the outcome conditionally on treatment.

Based on our theoretical results, we propose the Debiased Outcome-adapted Propensity Estimator (DOPE) for efficient estimation of the ATE, and we provide asymptotic results for the DOPE under general conditions. Compared to the augmented inverse propensity weighted (AIPW) estimator, the DOPE can retain its efficiency even when the covariates are highly predictive of treatment. We illustrate this with a single-index model, and with an implementation of the DOPE based on neural networks, we demonstrate its performance on simulated and real data. Our results show that the DOPE provides an efficient and robust methodology for ATE estimation in various observational settings.

3.1 Introduction

Estimating the population average treatment effect (ATE) of a treatment on an outcome variable is a fundamental statistical task. A naive approach is to contrast the mean outcome of a treated population with the mean outcome of an untreated population. Using observational data this is, however, generally a flawed approach due to

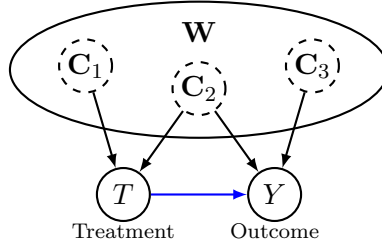


Figure 3.1.1: The covariate \mathbf{W} can have a complex data structure, even if the information it represents is structured and can be categorized into components that influence treatment and outcome separately.

confounding. If the underlying confounding mechanisms are captured by a set of pre-treatment covariates \mathbf{W} , it is possible to adjust for confounding by conditioning on \mathbf{W} in a certain manner. Given that multiple subsets of \mathbf{W} may be valid for this adjustment, it is natural to ask if there is an ‘optimal adjustment subset’ that enables the most efficient estimation of the ATE.

Assuming a causal linear graphical model, Henckel et al. [2022] established the existence of – and gave graphical criteria for – an *optimal adjustment set* for the OLS estimator. Rotnitzky and Smucler [2020] extended the results of Henckel et al. [2022], by proving that the optimality was valid within general causal graphical models and for all regular and asymptotically linear estimators. Critically, this line of research assumes knowledge of the underlying graphical structure.

To accommodate the assumption of *no unmeasured confounding*, observational data is often collected with as many covariates as possible, which means that \mathbf{W} can be high-dimensional. In such cases, assumptions of a known graph are unrealistic, and graphical estimation methods are statistically unreliable [Uhler et al., 2013, Shah and Peters, 2020, Chickering et al., 2004]. Furthermore, for non-Euclidean data such as images or texts, it is not clear how to impose any graphical structure pertaining to causal relations. Nevertheless, we can in these cases still imagine that the information that \mathbf{W} represents can be separated into distinct components that affect treatment and outcome directly, as illustrated in Figure 3.1.1.

In this paper, we formalize this idea and formulate a novel and general adjustment theory with a focus on efficiency bounds for the estimation of the average treatment effect. Based on the adjustment theory, we propose a general estimation procedure and analyze its asymptotic behavior.

3.1.1 Setup

Throughout we consider a discrete treatment variable $T \in \mathbb{T}$, a square-integrable outcome variable $Y \in \mathbb{R}$, and pre-treatment covariates $\mathbf{W} \in \mathbb{W}$. For now we only require that \mathbb{T} is a finite set and that \mathbb{W} is a measurable space. The joint distribution of (T, \mathbf{W}, Y) is denoted by P and it is assumed to belong to a collection of probability measures \mathcal{P} . If we need to make the joint distribution explicit, we denote expectations and probabilities

with \mathbb{E}_P and \mathbb{P}_P , respectively, but usually the P is omitted for ease of notation.

Our model-free target parameters of interest are of the form

$$\chi_t := \mathbb{E}[\mathbb{E}[Y | T = t, \mathbf{W}]] = \mathbb{E} \left[\frac{\mathbb{1}(T = t)Y}{\mathbb{P}(T = t | \mathbf{W})} \right], \quad t \in \mathbb{T}. \quad (3.1.1)$$

In words, these are treatment specific means of the outcome when adjusting for the covariate \mathbf{W} . To ensure that this quantity is well-defined, we assume the following condition, commonly known as *positivity*.

Assumption 3.1.1 (Positivity). *It holds that $0 < \mathbb{P}(T = t | \mathbf{W}) < 1$ almost surely for each $t \in \mathbb{T}$.*

Under additional assumptions common in causal inference literature – which informally entail that \mathbf{W} captures all confounding – the target parameter χ_t has the interpretation as the *interventional mean*, which is expressed by $\mathbb{E}[Y | \text{do}(T = t)]$ in *do-notation* or by $\mathbb{E}[Y^t]$ using *potential outcome* notation [Peters et al., 2017, van der Laan and Rose, 2011]. Under such causal assumptions, the average treatment effect is identified, and when $\mathbb{T} = \{0, 1\}$ it is typically expressed as the contrast $\chi_1 - \chi_0$. The theory in this paper is agnostic with regards to whether or not χ_t has this causal interpretation, although it is the primary motivation for considering χ_t as a target parameter.

Given n i.i.d. observations of (T, \mathbf{W}, Y) , one may proceed to estimate χ_t by estimating either of the equivalent expressions for χ_t in Equation (3.1.1). Within parametric models, the outcome regression function

$$g(t, \mathbf{w}) = \mathbb{E}[Y | T = t, \mathbf{W} = \mathbf{w}]$$

can typically be estimated with a \sqrt{n} -rate. In this case, the sample mean of the estimated regression function yields a \sqrt{n} -consistent estimator of χ_t under Donsker class conditions or sample splitting. However, many contemporary datasets indicate that parametric model-based regression methods get outperformed by nonparametric methods such as boosting and neural networks [Bojer and Meldgaard, 2021].

Nonparametric estimators of the regression function typically converge at rates slower than \sqrt{n} , and likewise for estimators of the propensity score

$$m(t | \mathbf{w}) = \mathbb{P}(T = t | \mathbf{W} = \mathbf{w}).$$

Even if both nonparametric estimators have rates slower than \sqrt{n} , it is in some cases possible to achieve a \sqrt{n} -rate of χ_t by modeling both m and g , and then combining their estimates in a way that achieves ‘rate double robustness’ [Smucler et al., 2019]. That is, an estimation error of the same order as the product of the errors for m and g . Two prominent estimators that have this property are the Augmented Inverse Probability Weighted estimator (AIPW) and the Targeted Minimum Loss-based Estimator (TMLE) [Robins and Rotnitzky, 1995, Chernozhukov et al., 2018, van der Laan and Rose, 2011].

In what follows, the premise is that even with a \sqrt{n} -rate estimator of χ_t , it might still be intractable to model m – and possibly also g – as a function of \mathbf{W} directly. This can

happen for primarily two reasons: (i) the sample space \mathbb{W} is high-dimensional or has a complex structure, or (ii) the covariate is highly predictive of treatment, leading to unstable predictions of the inverse propensity score $\mathbb{P}(T = t | \mathbf{W})^{-1}$.

In either of these cases, which are not exclusive, we can try to manage these difficulties by instead working with a *representation* $\mathbf{Z} = \varphi(\mathbf{W})$, given by a measurable mapping φ from \mathbb{W} into a more tractable space such as \mathbb{R}^d . In the first case above, such a representation might be a pre-trained word embedding, e.g., the celebrated BERT and its offsprings [Devlin et al., 2018]. The second case has been well-studied in the special case where $\mathbb{W} = \mathbb{R}^k$ and where \mathcal{P} contains the distributions that are consistent with respect to a fixed DAG (or CPDAG). We develop a general theory that subsumes both cases, and we discuss how to represent the original covariates to efficiently estimate the adjusted mean χ_t .

3.1.2 Relations to existing literature

Various studies have explored the adjustment for complex data structures by utilizing a (deep) representation of the covariates, as demonstrated in works such as Shi et al. [2019], Veitch et al. [2020]. In a different research direction, the *Collaborative TMLE* [van der Laan and Gruber, 2010] has emerged as a robust method for estimating average treatment effects by collaboratively learning the outcome regression and propensity score, particularly in scenarios where covariates are highly predictive of treatment [Ju et al., 2019]. Our overall estimation approach shares similarities with the mentioned strategies; for instance, our proof-of-concept estimator in the experimental section employs neural networks with shared layers. However, unlike the cited works, it incorporates the concept of efficiently tuning the representation specifically for predicting outcomes, rather than treatment. Related to this idea is another interesting line of research, which builds upon the *outcome adapted lasso* proposed by Shortreed and Ertefaie [2017]. Such works include Ju et al. [2020], Benkeser et al. [2020], Greenewald et al. [2021], Baldé et al. [2023]. These works all share the common theme of proposing estimation procedures that select covariates based on L_1 -penalized regression onto the outcome, and then subsequently estimate the propensity score based on the selected covariates adapted to the outcome. The theory of this paper generalizes the particular estimators proposed in the previous works, and also allows for other feature selection methods than L_1 -penalization. Moreover, our generalization of (parts of) the efficient adjustment theory from Rotnitzky and Smucler [2020] allows us to theoretically quantify the efficiency gains from these estimation methods. Finally, our asymptotic theory considers a novel regime, which, according to the simulations, seems more adequate for describing the finite sample behavior than the asymptotic results of Benkeser et al. [2020] and Ju et al. [2020].

Our general adjustment results in Section 3.3 draw on the vast literature on classical adjustment and confounder selection, for example Rosenbaum and Rubin [1983], Hahn [1998], Henckel et al. [2022], Rotnitzky and Smucler [2020], Guo et al. [2022], Perković et al. [2018], Peters et al. [2017], Forré and Mooij [2023]. In particular, two of our results are direct extensions of results from Rotnitzky and Smucler [2020], Henckel et al. [2022].

3.1.3 Organization of the paper

In Section 3.2 we discuss generalizations of classical adjustment concepts to abstract conditioning on information. In Section 3.3 we discuss information bounds in the framework of Section 3.2. In Section 3.4 we propose a novel method, the DOPE, for efficient estimation of adjusted means, and we discuss the asymptotic behavior of the resulting estimator. In Section 3.5 we implement the DOPE and demonstrate its performance on synthetic and real data. The paper is concluded by a discussion in Section 3.6.

3.2 Generalized adjustment concepts

In this section we discuss generalizations of classical adjustment concepts. These generalizations are motivated by the premise from the introduction: it might be intractable to model the propensity score directly as a function of \mathbf{W} , so instead we consider adjusting for a representation $\mathbf{Z} = \varphi(\mathbf{W})$. This is, in theory, confined to statements about conditioning on \mathbf{Z} , and is therefore equivalent to adjusting for any representation of the form $\tilde{\mathbf{Z}} = \psi \circ \varphi(\mathbf{W})$, where ψ is a bijective and bimeasurable mapping. The equivalence class of such representations is characterized by the σ -algebra generated by \mathbf{Z} , denoted by $\sigma(\mathbf{Z})$, which informally describes the information contained in \mathbf{Z} . In view of this, we define adjustment with respect to sub- σ -algebras contained in $\sigma(\mathbf{W})$.

Remark 3.2.1. Conditional expectations and probabilities are, unless otherwise indicated, defined conditionally on σ -algebras as in Kolmogoroff [1933], see also Kallenberg [2021, Ch. 8]. Equalities between conditional expectations are understood to hold almost surely. When conditioning on the random variable T and a σ -algebra \mathcal{Z} we write ‘ $|T, \mathcal{Z}$ ’ as a shorthand for ‘ $|\sigma(T) \vee \mathcal{Z}$ ’. Finally, we define conditioning on both the event $(T = t)$ and on a σ -algebra $\mathcal{Z} \subseteq \sigma(\mathbf{W})$ by $\mathbb{E}[Y | T = t, \mathcal{Z}] := \frac{\mathbb{E}[Y \mathbf{1}(T=t) | \mathcal{Z}]}{\mathbb{P}(T=t | \mathcal{Z})}$, which is well-defined under Assumption 3.1.1. \diamond

Definition 3.2.2. A sub- σ -algebra $\mathcal{Z} \subseteq \sigma(\mathbf{W})$ is called a *description* of \mathbf{W} . For each $t \in \mathbb{T}$ and $P \in \mathcal{P}$, and with \mathcal{Z} a description of \mathbf{W} , we define

$$\begin{aligned} \pi_t(\mathcal{Z}; P) &:= \mathbb{P}_P(T = t | \mathcal{Z}), \\ b_t(\mathcal{Z}; P) &:= \mathbb{E}_P[Y | T = t, \mathcal{Z}], \\ \chi_t(\mathcal{Z}; P) &:= \mathbb{E}_P[b_t(\mathcal{Z}; P)] = \mathbb{E}_P[\pi_t(\mathcal{Z}; P)^{-1} \mathbf{1}(T = t) Y]. \end{aligned}$$

If a description \mathcal{Z} of \mathbf{W} is given as $\mathcal{Z} = \sigma(\mathbf{Z})$ for a representation $\mathbf{Z} = \varphi(\mathbf{W})$, we may write $\chi_t(\mathbf{Z}; P)$ instead of $\chi_t(\sigma(\mathbf{Z}); P)$ etc.

We say that

$$\begin{aligned} \mathcal{Z} \text{ is } P\text{-valid if:} & \quad \chi_t(\mathcal{Z}; P) = \chi_t(\mathbf{W}; P), & \quad \text{for all } t \in \mathbb{T}, \\ \mathcal{Z} \text{ is } P\text{-OMS if:} & \quad b_t(\mathcal{Z}; P) = b_t(\mathbf{W}; P), & \quad \text{for all } t \in \mathbb{T}, \\ \mathcal{Z} \text{ is } P\text{-ODS if:} & \quad Y \perp\!\!\!\perp_P \mathbf{W} | T, \mathcal{Z}. \end{aligned}$$

Here OMS means *Outcome Mean Sufficient* and ODS means *Outcome Distribution Sufficient*. If \mathcal{Z} is P -valid for all $P \in \mathcal{P}$, we say that it is \mathcal{P} -valid. We define \mathcal{P} -OMS and \mathcal{P} -ODS analogously. \clubsuit

3 Efficient adjustment for complex covariates: Gaining efficiency with DOPE

A few remarks are in order.

- We have implicitly extended the shorthand notation described in Remark 3.2.1 to the quantities in Definition 3.2.2.
- The quantity $\chi_t(\mathcal{Z}; P)$ is deterministic, whereas $\pi_t(\mathcal{Z}; P)$ and $b_t(\mathcal{Z}; P)$ are \mathcal{Z} -measurable real valued random variables. Thus, if \mathcal{Z} is generated by a representation \mathbf{Z} , then by the Doob-Dynkin lemma [Kallenberg, 2021, Lemma 1.14] these random variables can be expressed as functions of \mathbf{Z} . This fact does not play a role until the discussion of estimation in Section 3.4.
- There are examples of descriptions $\mathcal{Z} \subseteq \sigma(\mathbf{W})$ that are not given by representations. Such descriptions might be of little practical importance, but our results do not require \mathcal{Z} to be given by a representation. We view the σ -algebraic framework as a convenient abstraction that generalizes equivalence classes of representations.
- We have the following hierarchy of the properties in Definition 3.2.2:

$$P\text{-ODS} \implies P\text{-OMS} \implies P\text{-valid},$$

and the relations also hold if P is replaced by \mathcal{P} .

- The P -ODS condition relates to existing concepts in statistics. The condition can be viewed as a σ -algebraic analogue of *prognostic scores* [Hansen, 2008], and it holds that any prognostic score generates a P -ODS description. Moreover, if a description $\mathcal{Z} = \sigma(\mathbf{Z})$ is P -ODS, then its generator \mathbf{Z} is *c-equivalent* to \mathbf{W} in the sense of Pearl [2009], see in particular the claim following his Equation (11.8). In Remark 3.3.12 we discuss the relations between \mathcal{P} -ODS descriptions and classical statistical sufficiency.

The notion of \mathcal{P} -valid descriptions can be viewed as a generalization of valid adjustment sets, where subsets are replaced with sub- σ -algebras.

Example 3.2.3 (Comparison with adjustment sets in causal DAGs). Suppose $\mathbf{W} \in \mathbb{R}^k$ and let \mathcal{D} be a DAG on the nodes $\mathbf{V} = (T, \mathbf{W}, Y)$. Let $\mathcal{P} = \mathcal{M}(\mathcal{D})$ be the collection of continuous distributions (on \mathbb{R}^{k+2}) that are Markovian with respect to \mathcal{D} and with $\mathbb{E}|Y| < \infty$.

Any subset $\mathbf{Z} \subseteq \mathbf{W}$ is a representation of \mathbf{W} given by a coordinate projection, and the corresponding σ -algebra $\sigma(\mathbf{Z})$ is a description of \mathbf{W} . In this framework, a subset $\mathbf{Z} \subseteq \mathbf{W}$ is called a *valid adjustment set* for (T, Y) if for all $P \in \mathcal{P}$, $t \in \mathbb{T}$, and $y \in \mathbb{R}$

$$\mathbb{E}_P \left[\frac{\mathbb{1}(T = t)\mathbb{1}(Y \leq y)}{\mathbb{P}_P(T = t | \text{pa}_{\mathcal{D}}(T))} \right] = \mathbb{E}_P[\mathbb{P}_P(Y \leq y | T = t, \mathbf{Z})],$$

where $\text{pa}_{\mathcal{D}}$ denotes the parents T in \mathcal{D} , see for example Definition 2 in Rotnitzky and Smucler [2020]. It turns out that \mathbf{Z} is a valid adjustment set if and only if

$$\mathbb{E}_P \left[\frac{\mathbb{1}(T = t)Y}{\mathbb{P}_P(T = t | \text{pa}_{\mathcal{D}}(T))} \right] = \mathbb{E}_P[\mathbb{E}_P[Y | T = t, \mathbf{Z}]] = \chi_t(\mathcal{Z}; P)$$

for all $P \in \mathcal{P}$ and $t \in \mathbb{T}$. This follows¹ from results of Perković et al. [2018], which we discuss for completeness in Proposition 3.A.1 in the supplement. Thus, if we assume that \mathbf{W} is a valid adjustment set, then any subset $\mathbf{Z} \subseteq \mathbf{W}$ is a valid adjustment set if and only if the corresponding description $\sigma(\mathbf{Z})$ is \mathcal{P} -valid. ♠

In general, \mathcal{P} -valid descriptions are not necessarily generated from valid adjustment sets. This can happen if structural assumptions are imposed on the conditional mean, which is illustrated in the following example.

Example 3.2.4. Suppose that the outcome regression is known to be invariant under rotations of $\mathbf{W} \in \mathbb{R}^k$ such that for any $P \in \mathcal{P}$,

$$\mathbb{E}_P[Y | T, \mathbf{W}] = \mathbb{E}_P[Y | T, \|\mathbf{W}\|] =: g_P(T, \|\mathbf{W}\|).$$

Without further graphical assumptions, we cannot deduce that any proper subset of \mathbf{W} is a valid adjustment set. In contrast, the magnitude $\|\mathbf{W}\|$ generates a \mathcal{P} -OMS – and hence also \mathcal{P} -valid – description of \mathbf{W} by definition.

Suppose that there is a distribution $\tilde{P} \in \mathcal{P}$ for which $g_{\tilde{P}}(t, \cdot)$ is bijective. Then $\sigma(\|\mathbf{W}\|)$ is also the smallest \mathcal{P} -OMS description up to $\mathbb{P}_{\tilde{P}}$ -negligible sets: if \mathcal{Z} is another \mathcal{P} -OMS description, we see that $b_t(\mathcal{Z}; \tilde{P}) = b_t(\mathbf{W}; \tilde{P}) = g_{\tilde{P}}(t, \|\mathbf{W}\|)$ almost surely. Hence

$$\sigma(\|\mathbf{W}\|) = \sigma(g_{\tilde{P}}(t, \|\mathbf{W}\|)) \subseteq \overline{\sigma(b_t(\mathcal{Z}; \tilde{P}))} \subseteq \overline{\mathcal{Z}},$$

where overline denotes the $\mathbb{P}_{\tilde{P}}$ -completion of a σ -algebra. That is, $\overline{\mathcal{Z}}$ is the smallest σ -algebra containing \mathcal{Z} and all its $\mathbb{P}_{\tilde{P}}$ -negligible sets. ♠

Even in the above example, where the regression function is known to depend on a one-dimensional function of \mathbf{W} , it is generally not possible to estimate the regression function at a \sqrt{n} -rate without restrictive assumptions. Thus, modeling of the propensity score is required in order to obtain a \sqrt{n} -rate estimator of the adjusted mean. If \mathbf{W} is highly predictive of treatment, then naively applying a doubly robust estimator (AIPW, TMLE) can be statistically unstable due to large inverse propensity weights. Alternatively, since $\sigma(\|\mathbf{W}\|)$ is a \mathcal{P} -valid description, we could also base our estimator on the *pruned propensity* $\mathbb{P}(T = t | \|\mathbf{W}\|)$. This approach should intuitively provide more stable weights, as we expect $\|\mathbf{W}\|$ to be less predictive of treatment. We proceed to analyze the difference between the asymptotic efficiencies of the two approaches and we show that, under reasonable conditions, the latter approach is never worse asymptotically.

3.3 Efficiency bounds for adjusted means

We now discuss efficiency bounds based on the concepts introduced in Section 3.2. To this end, let $\mathbf{Z} = \varphi(\mathbf{W})$ be a representation of \mathbf{W} . Under a sufficiently dense model \mathcal{P} , the *influence function* for $\chi_t(\mathbf{Z}; P)$ is given by

$$\psi_t(\mathbf{Z}; P) = b_t(\mathbf{Z}; P) + \frac{\mathbb{1}(T = t)}{\pi_t(\mathbf{Z}; P)}(Y - b_t(\mathbf{Z}; P)) - \chi_t(\mathbf{Z}; P). \quad (3.3.2)$$

¹We thank Leonard Henckel for pointing this out.

3 Efficient adjustment for complex covariates: Gaining efficiency with DOPE

The condition that \mathcal{P} is sufficiently dense entails that no ‘structural assumptions’ are imposed on the functional forms of π_t and b_t , see also Robins et al. [1994], Robins and Rotnitzky [1995], Hahn [1998]. Structural assumptions do not include smoothness conditions, but do include, for example, parametric model assumptions on the outcome regression.

Using the formalism of Section 3.2, we define $\psi_t(\mathcal{Z}; P)$ analogously to (3.3.2) for any description \mathcal{Z} of \mathbf{W} . We denote the variance of the influence function by

$$\begin{aligned} \mathbb{V}_t(\mathcal{Z}; P) &:= \text{Var}_P[\psi_t(\mathcal{Z}; P)] \\ &= \mathbb{E}_P \left[\frac{\text{Var}_P(Y | T = t, \mathcal{Z})}{\pi_t(\mathcal{Z}; P)} \right] + \text{Var}_P(b_t(\mathcal{Z}; P)). \end{aligned} \quad (3.3.3)$$

The importance of this variance was highlighted by Hahn [1998], who computed $\mathbb{V}_t(\mathbf{W}; P)$ as the semiparametric efficiency bound for regular asymptotically linear (RAL) estimators of $\chi_t(\mathbf{W}; P)$. That is, for any \mathcal{P} -consistent RAL estimator $\hat{\chi}_t(\mathbf{W}; P)$ based on n i.i.d. observations, the asymptotic variance of $\sqrt{n}(\hat{\chi}_t(\mathbf{W}; P) - \chi_t(\mathbf{W}; P))$ is bounded below by $\mathbb{V}_t(\mathbf{W}; P)$, provided that \mathcal{P} does not impose structural assumptions on $\pi_t(\mathbf{W})$ and $b_t(\mathbf{W})$. Moreover, the asymptotic variance of both the TMLE and AIPW estimators achieve this lower bound when the propensity score and outcome regression can be estimated with sufficiently fast rates [Chernozhukov et al., 2018, van der Laan and Rose, 2011].

Since the same result can be applied for the representation $\mathbf{Z} = \varphi(\mathbf{W})$, we see that estimation of $\chi_t(\mathbf{Z}; P)$ has a semiparametric efficiency bound of $\mathbb{V}_t(\mathbf{Z}; P)$. Now suppose that $\sigma(\mathbf{Z})$ is a \mathcal{P} -valid description of \mathbf{W} , which means that $\chi_t := \chi_t(\mathbf{W}; P) = \chi_t(\mathbf{Z}; P)$ for all P . It is then natural to ask if we should estimate χ_t based on \mathbf{W} or the representation \mathbf{Z} ? We proceed to investigate this question based on which of the efficiency bounds $\mathbb{V}_t(\mathbf{Z}; P)$ or $\mathbb{V}_t(\mathbf{W}; P)$ is smallest.

If $\mathbb{V}_t(\mathbf{Z}; P)$ is smaller than $\mathbb{V}_t(\mathbf{W}; P)$, then $\mathbb{V}_t(\mathbf{W}; P)$ is not an actual efficiency bound for $\chi_t(\mathbf{W}; P)$. This does not contradict the result of Hahn [1998], as we have assumed the existence of a non-trivial \mathcal{P} -valid description of \mathbf{W} , which implicitly imposes structural assumptions on the functional form of $b_t(\mathbf{W})$. Nevertheless, it is sensible to compare the variances $\mathbb{V}_t(\mathbf{Z}; P)$ and $\mathbb{V}_t(\mathbf{W}; P)$, as these are the asymptotic variances of the AIPW when using either \mathbf{Z} or \mathbf{W} , respectively, as a basis for the nuisance functions.

We formulate our efficiency bounds in terms of the more general contrast parameter

$$\Delta = \Delta(\mathbf{W}; P) = \sum_{t \in \mathbb{T}} c_t \chi_t(\mathbf{W}; P), \quad (3.3.4)$$

where $\mathbf{c} := (c_t)_{t \in \mathbb{T}}$ are fixed real-valued coefficients. The prototypical example, when $\mathbb{T} = \{0, 1\}$, is $\Delta = \chi_1 - \chi_0$, which is the average treatment effect under causal assumptions, cf. the discussion following Assumption 3.1.1. Note that the family of Δ -parameters includes the adjusted mean χ_t as a special case.

To estimate Δ , we consider estimators of the form

$$\hat{\Delta}(\mathcal{Z}; P) = \sum_{t \in \mathbb{T}} c_t \hat{\chi}_t(\mathcal{Z}; P), \quad (3.3.5)$$

where $\hat{\chi}_t(\mathcal{Z}; P)$ denotes a consistent RAL estimator of $\chi_t(\mathcal{Z}; P)$. Correspondingly, the efficiency bound for such an estimator is

$$\mathbb{V}_\Delta(\mathcal{Z}; P) := \text{Var}_P \left(\sum_{t \in \mathbb{T}} c_t \psi_t(\mathcal{Z}; P) \right). \quad (3.3.6)$$

It turns out that two central results by Rotnitzky and Smucler [2020], specifically their Lemmas 4 and 5, can be generalized from covariate subsets to descriptions. One conceptual difference is that there is *a priori* no natural generalization of *precision variables* and *overadjustment (instrumental) variables*. To wit, if \mathcal{Z}_1 and \mathcal{Z}_2 are descriptions, there is no canonical way² to subtract \mathcal{Z}_2 from \mathcal{Z}_1 in a way that maintains their join $\mathcal{Z}_1 \vee \mathcal{Z}_2 := \sigma(\mathcal{Z}_1, \mathcal{Z}_2)$. Apart from this technical detail, the proofs translate more or less directly. The following lemma is a direct extension of Rotnitzky and Smucler [2020, Lemma 4].

Lemma 3.3.1 (Deletion of overadjustment). *Fix a distribution $P \in \mathcal{P}$ and let $\mathcal{Z}_1 \subseteq \mathcal{Z}_2$ be σ -algebras such that $Y \perp\!\!\!\perp_P \mathcal{Z}_2 | T, \mathcal{Z}_1$. Then it always holds that*

$$\mathbb{V}_\Delta(\mathcal{Z}_2; P) - \mathbb{V}_\Delta(\mathcal{Z}_1; P) = \sum_{t \in \mathbb{T}} c_t^2 D_t(\mathcal{Z}_1, \mathcal{Z}_2; P) \geq 0,$$

where for each $t \in \mathbb{T}$,

$$\begin{aligned} D_t(\mathcal{Z}_1, \mathcal{Z}_2; P) &:= \mathbb{V}_t(\mathcal{Z}_2; P) - \mathbb{V}_t(\mathcal{Z}_1; P) \\ &= \mathbb{E}_P \left[\pi_t(\mathcal{Z}_1; P) \text{Var}(Y | T = t, \mathcal{Z}_1) \text{Var}(\pi_t(\mathcal{Z}_2; P)^{-1} | T = t, \mathcal{Z}_1) \right]. \end{aligned}$$

Moreover, if \mathcal{Z}_2 is a description of \mathbf{W} then \mathcal{Z}_1 is P -valid if and only if \mathcal{Z}_2 is P -valid.

The lemma quantifies the efficiency lost in adjustment when adding information that is irrelevant for the outcome.

We proceed to apply this lemma to the minimal information in \mathbf{W} that is predictive of Y conditionally on T . To define this information, we use the regular conditional distribution function of Y given $T = t, \mathbf{W} = \mathbf{w}$, which we denote by

$$F(y | t, \mathbf{w}; P) := \mathbb{P}_P(Y \leq y | T = t, \mathbf{W} = \mathbf{w}), \quad y \in \mathbb{R}, t \in \mathbb{T}, \mathbf{w} \in \mathbb{W}.$$

See Kallenberg [2021, Sec. 8] for a rigorous treatment of regular conditional distributions. We will in the following, by convention, take

$$b_t(\mathbf{w}; P) = \int y \, dF(y | t, \mathbf{w}; P), \quad (3.3.7)$$

so that $b_t(\mathbf{W}; P)$ is given in terms of the regular conditional distribution.

²equivalent conditions to the existence of an independent complement are given in Proposition 4 in Emery and Schachermayer [2001].

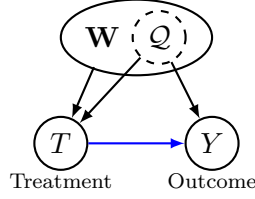


Figure 3.3.2: The σ -algebra \mathcal{Q} given in Definition 3.3.2 as a description of \mathbf{W} .

Definition 3.3.2. Define the σ -algebras

$$\begin{aligned}\mathcal{Q} &:= \bigvee_{P \in \mathcal{P}} \mathcal{Q}_P, & \mathcal{Q}_P &:= \sigma(F(y|t, \mathbf{W}; P); y \in \mathbb{R}, t \in \mathbb{T}), \\ \mathcal{R} &:= \bigvee_{P \in \mathcal{P}} \mathcal{R}_P, & \mathcal{R}_P &:= \sigma(b_t(\mathbf{W}; P); t \in \mathbb{T}).\end{aligned}$$

♣

Note that \mathcal{Q}_P , \mathcal{Q} , \mathcal{R}_P and \mathcal{R} are all descriptions of \mathbf{W} , see Figure 3.3.2 for a depiction of the information contained in \mathcal{Q} . Note also that $\mathcal{R}_P \subseteq \mathcal{Q}_P$ by the convention (3.3.7).

We now state one of the main results of this section.

Theorem 3.3.3. Fix $P \in \mathcal{P}$. Then, under Assumption 3.1.1, it holds that

- (i) If \mathcal{Z} is a description of \mathbf{W} and $\mathcal{Q}_P \subseteq \overline{\mathcal{Z}}$, then \mathcal{Z} is P -ODS. In particular, \mathcal{Q}_P is a P -ODS description of \mathbf{W} .
- (ii) If \mathcal{Z} is P -ODS description, then $\mathcal{Q}_P \subseteq \overline{\mathcal{Z}}$.
- (iii) For any P -ODS description \mathcal{Z} it holds that

$$\mathbb{V}_\Delta(\mathcal{Z}; P) - \mathbb{V}_\Delta(\mathcal{Q}_P; P) = \sum_{t \in \mathbb{T}} c_t^2 D_t(\mathcal{Q}_P, \mathcal{Z}; P) \geq 0, \quad (3.3.8)$$

where D_t is given as in Lemma 3.3.1.

Together parts (i) and (ii) state that a description of \mathbf{W} is P -ODS if and only if its \mathbb{P}_P -completion contains \mathcal{Q}_P . Part (iii) states that, under $P \in \mathcal{P}$, \mathcal{Q}_P leads to the optimal efficiency bound among all P -ODS descriptions.

In the following corollary we use $\overline{\mathcal{Z}}^P$ to denote the \mathbb{P}_P -completion of a σ -algebra \mathcal{Z} .

Corollary 3.3.4. Let \mathcal{Z} be a description of \mathbf{W} . Then \mathcal{Z} is \mathcal{P} -ODS if and only if $\mathcal{Q}_P \subseteq \overline{\mathcal{Z}}^P$ for all $P \in \mathcal{P}$. A sufficient condition for \mathcal{Z} to be \mathcal{P} -ODS is that $\mathcal{Q} \subseteq \overline{\mathcal{Z}}^P$ for all $P \in \mathcal{P}$, in which case

$$\mathbb{V}_\Delta(\mathcal{Z}; P) - \mathbb{V}_\Delta(\mathcal{Q}; P) = \sum_{t \in \mathbb{T}} c_t^2 D_t(\mathcal{Q}, \mathcal{Z}; P) \geq 0, \quad P \in \mathcal{P}. \quad (3.3.9)$$

In particular, \mathcal{Q} is a \mathcal{P} -ODS description of \mathbf{W} , and (3.3.9) holds with $\mathcal{Z} = \sigma(\mathbf{W})$.

Remark 3.3.5. It is *a priori* not obvious if \mathcal{Q} is given by a representation, i.e., if $\mathcal{Q} = \sigma(\varphi(\mathbf{W}))$ for some measurable mapping φ . In Example 3.2.4 it is, since the arguments can be reformulated to conclude that $\mathcal{Q} = \sigma(\|\mathbf{W}\|)$. \diamond

Instead of working with the entire conditional distribution, it suffices to work with the conditional mean when assuming, e.g., independent additive noise on the outcome.

Proposition 3.3.6. *For fixed $P \in \mathcal{P}$, then $\mathcal{Q}_P = \mathcal{R}_P$ if any of the following statements are true:*

- $F(y|t, \mathbf{W}; P)$ is $\sigma(b_t(\mathbf{W}; P))$ -measurable for each $t \in \mathbb{T}, y \in \mathbb{R}$.
- Y is a binary outcome.
- Y has independent additive noise, i.e., $Y = b_T(\mathbf{W}) + \varepsilon_Y$ with $\varepsilon_Y \perp\!\!\!\perp_P T, \mathbf{W}$.

If $\mathcal{Q}_P = \mathcal{R}_P$ holds, then (3.3.8) also holds for any P -OMS description.

Remark 3.3.7. When $\mathcal{Q}_P = \mathcal{R}_P$ for all $P \in \mathcal{P}$ we have $\mathcal{Q} = \mathcal{R}$. There is thus the same information in the σ -algebra \mathcal{R} generated by the conditional means of the outcome as there is in the σ -algebra \mathcal{Q} generated by the entire conditional distribution of the outcome. The three conditions in Proposition 3.3.6 are sufficient but not necessary to ensure this. \diamond

We also have a result analogous to Lemma 4 of Rotnitzky and Smucler [2020]:

Lemma 3.3.8 (Supplementation with precision). *Fix $P \in \mathcal{P}$ and let $\mathcal{Z}_1 \subseteq \mathcal{Z}_2$ be descriptions of \mathbf{W} such that $T \perp\!\!\!\perp_P \mathcal{Z}_2 | \mathcal{Z}_1$. Then \mathcal{Z}_1 is P -valid if and only if \mathcal{Z}_2 is P -valid. Irrespectively, it always holds that*

$$\mathbb{V}_\Delta(\mathcal{Z}_1; P) - \mathbb{V}_\Delta(\mathcal{Z}_2; P) = \mathbf{c}^\top \text{Var}_P[\mathbf{R}(\mathcal{Z}_1, \mathcal{Z}_2; P)]\mathbf{c} \geq 0$$

where $\mathbf{R}(\mathcal{Z}_1, \mathcal{Z}_2; P) := (R_t(\mathcal{Z}_1, \mathcal{Z}_2; P))_{t \in \mathbb{T}}$ with

$$R_t(\mathcal{Z}_1, \mathcal{Z}_2; P) := \left(\frac{\mathbb{1}(T = t)}{\pi_t(\mathcal{Z}_2; P)} - 1 \right) (b_t(\mathcal{Z}_2; P) - b_t(\mathcal{Z}_1; P)).$$

Writing $R_t = R_t(\mathcal{Z}_1, \mathcal{Z}_2; P)$, the components of the covariance matrix of \mathbf{R} are given by

$$\begin{aligned} \text{Var}_P(R_t) &= \mathbb{E}_P \left[\left(\frac{1}{\pi_t(\mathcal{Z}_1; P)} - 1 \right) \text{Var}_P[b_t(\mathcal{Z}_2; P) | \mathcal{Z}_1] \right], \\ \text{Cov}_P(R_s, R_t) &= -\mathbb{E}_P[\text{Cov}_P(b_s(\mathcal{Z}_2; P)b_t(\mathcal{Z}_2; P) | \mathcal{Z}_1)]. \end{aligned}$$

As a consequence, we obtain the well-known fact that the propensity score is a valid adjustment if \mathbf{W} is, cf. Theorems 1–3 in Rosenbaum and Rubin [1983].

Corollary 3.3.9. *Let $\Pi_P = \sigma(\pi_t(\mathbf{W}; P): t \in \mathbb{T})$. If \mathcal{Z} is a description of \mathbf{W} containing Π_P , then \mathcal{Z} is P -valid and*

$$\mathbb{V}_\Delta(\Pi_P; P) - \mathbb{V}_\Delta(\mathcal{Z}; P) = \mathbf{c}^\top \text{Var}_P[\mathbf{R}(\mathcal{R}_P, \mathcal{Z}; P)]\mathbf{c} \geq 0.$$

3 Efficient adjustment for complex covariates: Gaining efficiency with DOPE

The corollary asserts that while the information contained in the propensity score is valid, it is asymptotically inefficient to adjust for in contrast to all the information of \mathbf{W} . This is in similar spirit to Theorem 2 of Hahn [1998], which states that the efficiency bound $\mathbb{V}_\Delta(\mathbf{W}; P)$ remains unaltered if the propensity is considered as known. However, the corollary also quantifies the difference of the asymptotic efficiencies.

Corollary 3.3.4 asserts that \mathcal{Q} is maximally efficient over all \mathcal{P} -ODS descriptions \mathcal{Z} satisfying $\mathcal{Q} \subseteq \overline{\mathcal{Z}}^P$, and Proposition 3.3.6 asserts that in special cases, \mathcal{Q} reduces to \mathcal{R} . Since \mathcal{R} is \mathcal{P} -OMS, hence \mathcal{P} -valid, it is natural to ask if \mathcal{R} is generally more efficient than \mathcal{Q} . The following example shows that their efficiency bounds may be incomparable uniformly over \mathcal{P} .

Example 3.3.10. Let $0 < \delta < \frac{1}{2}$ be fixed and let \mathcal{P} be the collection of data generating distributions that satisfy:

- $\mathbf{W} \in [\delta, 1 - \delta]$ with a symmetric distribution, i.e., $\mathbf{W} \stackrel{\mathcal{D}}{=} 1 - \mathbf{W}$.
- $T \in \{0, 1\}$ with $\mathbb{E}(T | \mathbf{W}) = \mathbf{W}$.
- $Y = T + g(|\mathbf{W} - \frac{1}{2}|) + v(\mathbf{W})\varepsilon_Y$, where $\varepsilon_Y \perp\!\!\!\perp (T, \mathbf{W})$, $\mathbb{E}[\varepsilon_Y^2] < \infty$, $\mathbb{E}[\varepsilon_Y] = 0$, and where $g: [0, \frac{1}{2} - \delta] \rightarrow \mathbb{R}$ and $v: [\delta, 1 - \delta] \rightarrow [0, \infty)$ are continuous functions.

Letting $\mathbf{Z} = |\mathbf{W} - \frac{1}{2}|$, it is easy to verify directly from Definition 3.3.2 that

$$\mathcal{Q} = \sigma(\mathbf{W}) \neq \sigma(\mathbf{Z}) = \mathcal{R}.$$

It follows that \mathbf{Z} is \mathcal{P} -OMS but not \mathcal{P} -ODS. However, \mathbf{Z} is \mathcal{P}_1 -ODS in the homoscedastic submodel $\mathcal{P}_1 = \{P \in \mathcal{P} | v \equiv 1\}$. In fact, it generates the σ -algebra \mathcal{Q} within this submodel, i.e., $\sigma(\mathbf{Z}) = \vee_{P \in \mathcal{P}_1} \mathcal{Q}_P$. Thus $\mathbb{V}_\Delta(\mathbf{Z}; P) \leq \mathbb{V}_\Delta(\mathbf{W}; P)$ for all $P \in \mathcal{P}_1$.

We refer to the supplementary Section 3.A.6 for complete calculations of the subsequent formulas.

From symmetry it follows that $\pi_1(\mathbf{Z}) = 0.5$ and hence we conclude that $T \perp\!\!\!\perp \mathbf{Z}$. By Lemma 3.3.8, it follows that 0 (the trivial adjustment) is \mathcal{P} -valid, but with $\mathbb{V}_\Delta(\mathbf{Z}; P) \leq \mathbb{V}_\Delta(0; P)$ for all $P \in \mathcal{P}$.

Alternatively, direct computation yields that

$$\begin{aligned} \mathbb{V}_t(0) &= 2 \operatorname{Var}(g(\mathbf{Z})) + 2 \mathbb{E}[v(\mathbf{W})^2] \mathbb{E}[\varepsilon_Y^2] \\ \mathbb{V}_t(\mathbf{Z}) &= \operatorname{Var}(g(\mathbf{Z})) + 2 \mathbb{E}[v(\mathbf{W})^2] \mathbb{E}[\varepsilon_Y^2] \\ \mathbb{V}_t(\mathbf{W}) &= \operatorname{Var}(g(\mathbf{Z})) + \mathbb{E}[v(\mathbf{W})^2 / \mathbf{W}] \mathbb{E}[\varepsilon_Y^2]. \end{aligned}$$

With $\Delta = \chi_t$, the first two equalities confirm that $\mathbb{V}_\Delta(\mathbf{Z}; P) \leq \mathbb{V}_\Delta(0; P)$, $P \in \mathcal{P}$, and the last two yield that indeed $\mathbb{V}_\Delta(\mathbf{Z}; P) \leq \mathbb{V}_\Delta(\mathbf{W}; P)$ for $P \in \mathcal{P}_1$ by applying Jensen's inequality. In fact, these are strict inequalities whenever $g(\mathbf{Z})$ and ε_Y are non-degenerate.

Finally, we show that it is possible for $\mathbb{V}_\Delta(\mathbf{Z}; P) > \mathbb{V}_\Delta(\mathbf{W}; P)$ for $P \notin \mathcal{P}_1$. Let $\tilde{P} \in \mathcal{P}$ be a data generating distribution with $\mathbb{E}_{\tilde{P}}[\varepsilon_Y^2] > 0$, $v(\mathbf{W}) = \mathbf{W}^2$, and with \mathbf{W} uniformly

distributed on $[\delta, 1 - \delta]$. Then

$$\begin{aligned} 2\mathbb{E}_{\tilde{P}}[v(\mathbf{W})^2] &= 2\mathbb{E}_{\tilde{P}}[\mathbf{W}^4] = 2\frac{(1-\delta)^5 - \delta^5}{5(1-2\delta)} \xrightarrow{\delta \rightarrow 0} \frac{2}{5} \\ \mathbb{E}_{\tilde{P}}\left[\frac{v(\mathbf{W})^2}{\mathbf{W}}\right] &= \mathbb{E}_{\tilde{P}}[\mathbf{W}^3] = \frac{(1-\delta)^4 - \delta^4}{4(1-2\delta)} \xrightarrow{\delta \rightarrow 0} \frac{1}{4}. \end{aligned}$$

So for sufficiently small δ , it holds that $\mathbb{V}_{\Delta}(\mathbf{Z}; \tilde{P}) > \mathbb{V}_{\Delta}(\mathbf{W}; \tilde{P})$. The example can also be modified to work for other $\delta > 0$ by taking v to be a sufficiently large power of \mathbf{W} . ♠

Remark 3.3.11. Following Theorem 1 of Benkeser et al. [2020], it is stated that the asymptotic variance $\mathbb{V}_{\Delta}(\mathcal{R}_P; P)$ is generally smaller than $\mathbb{V}_{\Delta}(\mathbf{W}; P)$. The example demonstrates that this requires some assumptions on the outcome distribution, such as the conditions in Proposition 3.3.6. ◇

Remark 3.3.12. Suppose that $\mathcal{P} = \{f_{\theta} \cdot \mu : \theta \in \Theta\}$ is a parametrized family of measures with densities $\{f_{\theta}\}_{\theta \in \Theta}$ with respect to a σ -finite measure μ . Then informally, a *sufficient sub- σ -algebra* is any subset of the observed information for which the remaining information is independent of $\theta \in \Theta$, see Billingsley [2017] for a formal definition. Superficially, this concept seems similar to that of a \mathcal{P} -ODS description. In contrast however, the latter is a subset of the covariate information $\sigma(\mathbf{W})$ rather than the observed information $\sigma(T, \mathbf{W}, Y)$, and it concerns sufficiency for the outcome distribution rather than the entire data distribution. Moreover, the Rao-Blackwell theorem asserts that conditioning an estimator of θ on a sufficient sub- σ -algebra leads to an estimator that is never worse. Example 3.3.10 demonstrates that the situation is more delicate when considering statistical efficiency for adjustment. ◇

3.4 Estimation based on outcome-adapted representations

In this section we develop a general estimation method based on the insights of Section 3.3, and we provide an asymptotic analysis of this methodology under general conditions. Our method modifies the AIPW estimator, which we proceed to discuss in more detail.

We have worked with the propensity score, π_t , and the outcome regression, b_t , as random variables. We now consider, for each $P \in \mathcal{P}$, their function counterparts obtained as regular conditional expectations:

$$\begin{aligned} m_P: \mathbb{T} \times \mathbb{W} &\rightarrow \mathbb{R}, & m_P(t | \mathbf{w}) &:= \mathbb{P}_P(T = t | \mathbf{W} = \mathbf{w}), \\ g_P: \mathbb{T} \times \mathbb{W} &\rightarrow \mathbb{R}, & g_P(t, \mathbf{w}) &:= \mathbb{E}_P[Y | T = t, \mathbf{W} = \mathbf{w}]. \end{aligned}$$

To target the adjusted mean, χ_t , the AIPW estimator utilizes the influence function – given by (3.3.2) – as a score equation. To be more precise, given estimates (\hat{m}, \hat{g}) of the nuisance functions (m_P, g_P) , the AIPW estimator of χ_t is given by

$$\hat{\chi}_t^{\text{aipw}}(\hat{m}, \hat{g}) := \mathbb{P}_n \left[\hat{g}(t, \mathbf{W}) + \frac{\mathbb{1}(T = t)(Y - \hat{g}(t, \mathbf{W}))}{\hat{m}(t | \mathbf{W})} \right], \quad (3.4.10)$$

where $\mathbb{P}_n[\cdot]$ is defined as the empirical mean over n i.i.d. observations from P . The natural AIPW estimator of Δ given by (3.3.5) is then $\hat{\Delta}^{\text{aipw}}(\hat{m}, \hat{g}) = \sum_{t \in \mathbb{T}} c_t \hat{\chi}_t^{\text{aipw}}(\hat{m}, \hat{g})$.

Roughly speaking, the AIPW estimator $\hat{\chi}_t^{\text{aipw}}$ converges to χ_t with a \sqrt{n} -rate if

$$n \cdot \mathbb{P}_n[(\hat{m}(t | \mathbf{W}) - m_P(t | \mathbf{W}))^2] \cdot \mathbb{P}_n[(\hat{g}(t, \mathbf{W}) - g_P(t, \mathbf{W}))^2] \longrightarrow 0. \quad (3.4.11)$$

If (\hat{m}, \hat{g}) are estimated using the same data as \mathbb{P}_n , then the above asymptotic result relies on Donsker class conditions on the spaces containing \hat{m} and \hat{g} . Among others, Chernozhukov et al. [2018] propose circumventing Donsker class conditions by sample splitting techniques such as K -fold cross-fitting.

3.4.1 Representation adjustment and the DOPE

Suppose that the outcome regression function factors through an intermediate representation:

$$g_P(t, \mathbf{w}) = h_P(t, \varphi(\theta_P, \mathbf{w})), \quad t \in \mathbb{T}, \mathbf{w} \in \mathbb{W}, \quad (3.4.12)$$

where $\varphi: \Theta \times \mathbb{W} \rightarrow \mathbb{R}^d$ is a known measurable mapping with unknown parameter $\theta \in \Theta$, and where $h_P: \mathbb{T} \times \mathbb{R}^d \rightarrow \mathbb{R}$ is an unknown function. We use $\mathbf{Z}_\theta = \varphi(\theta, \mathbf{W})$ to denote the corresponding representation parametrized by $\theta \in \Theta$. If a particular covariate value $\mathbf{w} \in \mathbb{W}$ is clear from the context, we also use the implicit notation $\mathbf{z}_\theta = \varphi(\theta, \mathbf{w})$.

Example 3.4.1 (Single-index model). The (*partial*) *single-index model* applies to $\mathbf{W} \in \mathbb{R}^k$ and assumes that (3.4.12) holds with $\varphi(\theta, \mathbf{w}) = \mathbf{w}^\top \theta$, where $\theta \in \Theta \subseteq \mathbb{R}^k$. In other words, it assumes that the outcome regression factors through the *linear predictor* $\mathbf{Z}_\theta = \mathbf{W}^\top \theta$ such that

$$\forall P \in \mathcal{P}: \quad Y = h_P(T, \mathbf{Z}_{\theta_P}) + \varepsilon_Y, \quad \mathbb{E}_P[\varepsilon_Y | T, \mathbf{W}] = 0. \quad (3.4.13)$$

For each treatment $T = t$, the model extends the generalized linear model (GLM) by assuming that the (inverse) link function $h_P(t, \cdot)$ is unknown.

The *semiparametric least squares* (SLS) estimator proposed by Ichimura [1993] estimates θ_P and h_P by performing a two-step regression procedure. Alternative estimation procedures are given in Powell et al. [1989], Delecroix et al. [2003]. ♠

Given an estimator $\hat{\theta}$ of θ_P , independent of (T, \mathbf{W}, Y) , we use the following notation for various functions and estimators thereof:

$$\begin{aligned} g_{\hat{\theta}}(t, \mathbf{w}) &:= \mathbb{E}_P[Y | T = t, \mathbf{Z}_{\hat{\theta}} = \mathbf{z}_{\hat{\theta}}, \hat{\theta}], \\ m_{\hat{\theta}}(t | \mathbf{w}) &:= \mathbb{P}_P(T = t | \mathbf{Z}_{\hat{\theta}} = \mathbf{z}_{\hat{\theta}}, \hat{\theta}), \\ \hat{g}(t, \mathbf{w}) &: \text{an estimator of } g_P(t, \mathbf{w}), \\ \hat{g}_{\hat{\theta}}(t, \mathbf{w}) &: \text{an estimator of } g_{\hat{\theta}}(t, \mathbf{w}) \text{ of the form } \hat{h}(t, \mathbf{z}_{\hat{\theta}}), \end{aligned} \quad (3.4.14a)$$

$$\hat{m}_{\hat{\theta}}(t | \mathbf{w}) : \text{an estimator of } m_{\hat{\theta}}(t | \mathbf{w}) \text{ of the form } \hat{f}(t, \mathbf{z}_{\hat{\theta}}). \quad (3.4.14b)$$

3.4 Estimation based on outcome-adapted representations

In other words, $g_{\hat{\theta}}$ and $m_{\hat{\theta}}$ are the theoretical propensity score and outcome regression, respectively, when using the *estimated* representation $\mathbf{Z}_{\hat{\theta}} = \varphi(\hat{\theta}, \mathbf{W})$. Note that we have suppressed P from the notation on the left-hand side, as we will be working under a fixed P in this section.

Sufficient conditions are well known that ensure

$$\sqrt{n}(\hat{\chi}_t^{\text{aipw}}(\hat{m}, \hat{g}) - \chi_t) \xrightarrow{d} \text{N}(0, \mathbb{V}_t(\mathbf{W}; P)).$$

Such conditions, e.g., Assumption 5.1 in Chernozhukov et al. [2018], primarily entail the condition in (3.4.11). We will leverage (3.4.12) to derive more efficient estimators under similar conditions.

Suppose for a moment that θ_P is known. Since \mathbf{Z}_{θ_P} is P -OMS under (3.4.13), it holds that $\chi_t(\mathbf{Z}_{\theta_P}; P) = \chi_t$. Then under analogous conditions for the estimators $(\hat{m}_{\theta_P}, \hat{g}_{\theta_P})$, we therefore have

$$\sqrt{n}(\hat{\chi}_t^{\text{aipw}}(\hat{m}_{\theta_P}, \hat{g}_{\theta_P}) - \chi_t) \xrightarrow{d} \text{N}(0, \mathbb{V}_t(\mathbf{Z}_{\theta_P}; P)).$$

Under the conditions of Proposition 3.3.6, it holds that $\mathbb{V}_t(\mathbf{Z}_{\theta_P}; P) \leq \mathbb{V}_t(\mathbf{W}; P)$. In other words, $\hat{\chi}_t^{\text{aipw}}(\hat{m}_{\theta_P}, \hat{g}_{\theta_P})$ is asymptotically at least as efficient as $\hat{\chi}_t^{\text{aipw}}(\hat{m}, \hat{g})$.

In general, the parameter θ_P is, of course, unknown, and we therefore consider adjusting for an estimated representation $\mathbf{Z}_{\hat{\theta}}$. For simplicity, we present the special case $\Delta = \chi_t$, but the results are easily extended to general contrasts Δ . Our generic estimation procedure is described in Algorithm 3, where

$$\mathbf{Z}_{\hat{\theta}, i} := \varphi(\hat{\theta}, \mathbf{W}_i)$$

denotes the estimated representation of the i -th observed covariate. We refer to the resulting estimator as the *Debiased Outcome-adapted Propensity Estimator* (DOPE), and it is denoted by $\hat{\chi}_t^{\text{dope}}$.

Algorithm 3 is formulated such that \mathcal{I}_1 , \mathcal{I}_2 , and \mathcal{I}_3 can be arbitrary subsets of $[n]$, and for the asymptotic theory we assume that they are disjoint. However, in practical applications it might be reasonable to use the full sample for every estimation step, i.e., employing the algorithm with $\mathcal{I}_1 = \mathcal{I}_2 = \mathcal{I}_3 = [n]$. In this case, we also imagine that line 4 and line 5 are run simultaneously given that $\hat{\theta}$ may be derived from an outcome regression, cf. the SLS estimator in the single-index model (Example 3.4.1). In the supplementary Section 3.C, we also describe a more advanced cross-fitting scheme for Algorithm 3

Remark 3.4.2. Benkeser et al. [2020] use a similar idea as the DOPE, but their propensity factors through the final outcome regression function instead of a general intermediate representation. That our general formulation of Algorithm 3 contains their collaborative one-step estimator as a special case is seen as follows. Suppose $T \in \{0, 1\}$ is binary and let

$$\Theta = \{\theta = (g_P(0, \cdot), g_P(1, \cdot)): P \in \mathcal{P}\}$$

be the collection of outcome regression functions. Define the intermediate representation as the canonical pairing $\varphi(\theta, \mathbf{W}) = (\theta_1(\mathbf{W}), \theta_2(\mathbf{W}))$. Then, for any outcome regression

Algorithm 3: Debiased Outcome-adapted Propensity Estimator

- 1 **input:** observations $(T_i, \mathbf{W}_i, Y_i)_{i \in [n]}$, index sets $\mathcal{I}_1, \mathcal{I}_2, \mathcal{I}_3 \subseteq [n]$;
 - 2 **options:** method for computing $\hat{\theta}$, regression methods for propensity score and outcome regression of the form (3.4.14);
 - 3 **begin**
 - 4 compute estimate $\hat{\theta}$ based on data $(T_i, \mathbf{W}_i, Y_i)_{i \in \mathcal{I}_1}$;
 - 5 regress outcomes $(Y_i)_{i \in \mathcal{I}_2}$ onto $(T_i, \mathbf{Z}_{\hat{\theta}, i})_{i \in \mathcal{I}_2}$ to obtain $\hat{g}_{\hat{\theta}}(\cdot, \cdot)$;
 - 6 regress treatments $(T_i)_{i \in \mathcal{I}_2}$ onto $(\mathbf{Z}_{\hat{\theta}, i})_{i \in \mathcal{I}_2}$ to obtain $\hat{m}_{\hat{\theta}}(\cdot | \cdot)$;
 - 7 compute AIPW based on data $(T_i, \mathbf{W}_i, Y_i)_{i \in \mathcal{I}_3}$ and nuisance estimates $(\hat{m}_{\hat{\theta}}, \hat{g}_{\hat{\theta}})$:

$$\hat{\chi}_t^{\text{dope}} = \frac{1}{|\mathcal{I}_3|} \sum_{i \in \mathcal{I}_3} \left(\hat{g}_{\hat{\theta}}(t, \mathbf{W}_i) + \frac{\mathbb{1}(T_i = t)(Y_i - \hat{g}_{\hat{\theta}}(t, \mathbf{W}_i))}{\hat{m}_{\hat{\theta}}(t | \mathbf{W}_i)} \right)$$
 - 8 **return** DOPE: $\hat{\chi}_t^{\text{dope}}$.
-

estimate $\hat{\theta}$, we may set $\hat{g}_{\hat{\theta}}(t, \mathbf{w}) = (1 - t)\hat{\theta}_1(\mathbf{w}) + t\hat{\theta}_2(\mathbf{w})$, and the propensity score $\hat{m}_{\hat{\theta}}(t | \mathbf{w})$ factors through the outcome regressions. In this case and with $\mathcal{I}_1 = \mathcal{I}_2 = \mathcal{I}_3$, Algorithm 3 yields the collaborative one-step estimator of Benkeser et al. [2020, Appendix D]. Accordingly, we refer to this special case as DOPE-BCL (Benkeser, Cai and van der Laan). \diamond

3.4.2 Asymptotics of the DOPE

We proceed to discuss the asymptotics of the DOPE. For our theoretical analysis, we assume that the index sets in Algorithm 3 are disjoint. That is, the theoretical analysis relies on sample splitting.

Assumption 3.4.1. *The observations $(T_i, \mathbf{W}_i, Y_i)_{i \in [n]}$ used to compute $\hat{\chi}_t^{\text{dope}}$ are i.i.d. with the same distribution as (T, \mathbf{W}, Y) , and $[n] = \mathcal{I}_1 \cup \mathcal{I}_2 \cup \mathcal{I}_3$ is a partition such that $|\mathcal{I}_3| \rightarrow \infty$ as $n \rightarrow \infty$.*

In our simulations, employing sample splitting did not seem to enhance performance, and hence we regard Assumption 3.4.1 as a theoretical convenience rather than a practical necessity in all cases. Our results can likely also be established under alternative assumptions that avoid sample splitting, in particular Donsker class conditions.

We also henceforth use the convention that each of the quantities π_t, b_t, χ_t , and \mathbb{V}_t are defined conditionally on $\hat{\theta}$, e.g.,

$$\chi_t(\mathbf{Z}_{\hat{\theta}}) = \mathbb{E}[b_t(\mathbf{Z}_{\hat{\theta}}) | \hat{\theta}] = \mathbb{E}[\mathbb{E}[Y | T = t, \mathbf{Z}_{\hat{\theta}}, \hat{\theta}] | \hat{\theta}].$$

The error of the DOPE estimator can then be decomposed as

$$\hat{\chi}_t^{\text{dope}} - \chi_t = (\hat{\chi}_t^{\text{dope}} - \chi_t(\mathbf{Z}_{\hat{\theta}})) + (\chi_t(\mathbf{Z}_{\hat{\theta}}) - \chi_t). \quad (3.4.15)$$

The first term is the error had our target been the adjusted mean when adjusting for the estimated representation $\mathbf{Z}_{\hat{\theta}}$, whereas the second term is the adjustment bias that arises from adjusting for $\mathbf{Z}_{\hat{\theta}}$ rather than \mathbf{W} (or \mathbf{Z}_{θ}).

3.4.2.1 Estimation error conditionally on representation

To describe the asymptotics of the first term of (3.4.15), we consider the decomposition

$$\sqrt{|\mathcal{I}_3|} \widehat{\chi}_t^{\text{dope}} = U_{\hat{\theta}}^{(n)} + R_1 + R_2 + R_3, \quad (3.4.16)$$

where

$$\begin{aligned} U_{\hat{\theta}}^{(n)} &:= \frac{1}{\sqrt{|\mathcal{I}_3|}} \sum_{i \in \mathcal{I}_3} u_i(\hat{\theta}), & u_i(\hat{\theta}) &:= g_{\hat{\theta}}(t, \mathbf{W}_i) + \frac{\mathbb{1}(T_i = t)(Y_i - g_{\hat{\theta}}(t, \mathbf{W}_i))}{m_{\hat{\theta}}(t | \mathbf{W}_i)}, \\ R_1 &:= \frac{1}{\sqrt{|\mathcal{I}_3|}} \sum_{i \in \mathcal{I}_3} r_i^1, & r_i^1 &:= (\widehat{g}_{\hat{\theta}}(t, \mathbf{W}_i) - g_{\hat{\theta}}(t, \mathbf{W}_i)) \left(1 - \frac{\mathbb{1}(T_i = t)}{m_{\hat{\theta}}(t | \mathbf{W}_i)} \right), \\ R_2 &:= \frac{1}{\sqrt{|\mathcal{I}_3|}} \sum_{i \in \mathcal{I}_3} r_i^2, & r_i^2 &:= (Y_i - g_{\hat{\theta}}(t, \mathbf{W}_i)) \left(\frac{\mathbb{1}(T_i = t)}{\widehat{m}_{\hat{\theta}}(t | \mathbf{W}_i)} - \frac{\mathbb{1}(T_i = t)}{m_{\hat{\theta}}(t | \mathbf{W}_i)} \right), \\ R_3 &:= \frac{1}{\sqrt{|\mathcal{I}_3|}} \sum_{i \in \mathcal{I}_3} r_i^3, & r_i^3 &:= (\widehat{g}_{\hat{\theta}}(t, \mathbf{W}_i) - g_{\hat{\theta}}(t, \mathbf{W}_i)) \left(\frac{\mathbb{1}(T_i = t)}{\widehat{m}_{\hat{\theta}}(t | \mathbf{W}_i)} - \frac{\mathbb{1}(T_i = t)}{m_{\hat{\theta}}(t | \mathbf{W}_i)} \right). \end{aligned}$$

We show that the oracle term, $U_{\hat{\theta}}^{(n)}$, drives the asymptotic limit, and that the terms R_1, R_2, R_3 are remainder terms, subject to the conditions stated below.

Assumption 3.4.2. For $(T, \mathbf{W}, Y) \sim P$ satisfying the representation model (3.4.12), it holds that:

- (i) There exists $c > 0$ such that $\max\{|\widehat{m}_{\hat{\theta}} - \frac{1}{2}|, |m_{\hat{\theta}} - \frac{1}{2}|\} \leq \frac{1}{2} - c$.
- (ii) There exists $C > 0$ such that $\mathbb{E}[Y^2 | \mathbf{W}, T] \leq C$.
- (iii) There exists $\delta > 0$ such that $\mathbb{E}[|Y|^{2+\delta}] < \infty$.
- (iv) It holds that $\mathcal{E}_{1,t}^{(n)} := \frac{1}{|\mathcal{I}_3|} \sum_{i \in \mathcal{I}_3} (\widehat{m}_{\hat{\theta}}(t | \mathbf{W}_i) - m_{\hat{\theta}}(t | \mathbf{W}_i))^2 \xrightarrow{P} 0$.
- (v) It holds that $\mathcal{E}_{2,t}^{(n)} := \frac{1}{|\mathcal{I}_3|} \sum_{i \in \mathcal{I}_3} (\widehat{g}_{\hat{\theta}}(t, \mathbf{W}_i) - g_{\hat{\theta}}(t, \mathbf{W}_i))^2 \xrightarrow{P} 0$.
- (vi) It holds that $|\mathcal{I}_3| \cdot \mathcal{E}_{1,t}^{(n)} \mathcal{E}_{2,t}^{(n)} \xrightarrow{P} 0$.

Classical convergence results of the AIPW are proven under similar conditions, but with (vi) replaced by the stronger convergence in (3.4.11). We establish conditional asymptotic results under conditions on the conditional errors $\mathcal{E}_{1,t}^{(n)}$ and $\mathcal{E}_{2,t}^{(n)}$. To the best of our knowledge, the most similar results that we are aware of are those of Benkeser et al. [2020], and our proof techniques are most similar to those of Chernozhukov et al. [2018], Lundborg and Pfister [2023].

We can now state our first asymptotic result for the DOPE.

Theorem 3.4.3. *Under Assumptions 3.4.1 and 3.4.2, it holds that*

$$\mathbb{V}_t(\mathbf{Z}_{\hat{\theta}})^{-1/2} \cdot U_{\hat{\theta}}^{(n)} \xrightarrow{d} \mathbf{N}(0, 1), \quad \text{and} \quad R_i \xrightarrow{P} 0, \quad i = 1, 2, 3,$$

as $n \rightarrow \infty$. As a consequence,

$$\sqrt{|\mathcal{I}_3|} \cdot \mathbb{V}_t(\mathbf{Z}_{\hat{\theta}})^{-1/2} \left(\hat{\chi}_t^{\text{dope}} - \chi_t(\mathbf{Z}_{\hat{\theta}}) \right) \xrightarrow{d} \mathbf{N}(0, 1).$$

In other words, we can expect the DOPE to have an asymptotic distribution, conditionally on $\hat{\theta}$, approximated by

$$\hat{\chi}_t^{\text{dope}}(\hat{g}_{\hat{\theta}}, \hat{m}_{\hat{\theta}}) | \hat{\theta} \quad \overset{as.}{\sim} \quad \mathbf{N}\left(\chi_t(\mathbf{Z}_{\hat{\theta}}), \frac{1}{|\mathcal{I}_3|} \mathbb{V}_t(\mathbf{Z}_{\hat{\theta}})\right).$$

Note that if $|\mathcal{I}_3| = \lfloor n/3 \rfloor$, say, then the asymptotic variance is $\frac{3}{n} \mathbb{V}_t(\mathbf{Z}_{\hat{\theta}})$. Our simulation study indicates that the asymptotic approximation may be valid in some cases without the use of sample splitting, in which case $|\mathcal{I}_3| = n$ and the asymptotic variance is $\frac{1}{n} \mathbb{V}_t(\mathbf{Z}_{\hat{\theta}})$. A direct implementation of the sample splitting procedure thus comes with an efficiency cost. In the supplementary Section 3.C we discuss how the cross-fitting procedure makes use of sample splitting without an efficiency cost.

Given nuisance estimates (\hat{g}, \hat{m}) we consider the empirical variance estimator given by

$$\hat{\mathcal{V}}_t(\hat{g}, \hat{m}) = \frac{1}{|\mathcal{I}_3|} \sum_{i \in \mathcal{I}_3} \hat{u}_i^2 - \left(\frac{1}{|\mathcal{I}_3|} \sum_{i \in \mathcal{I}_3} \hat{u}_i \right)^2 = \frac{1}{|\mathcal{I}_3|} \sum_{i \in \mathcal{I}_3} \left(\hat{u}_i - \frac{1}{|\mathcal{I}_3|} \sum_{j \in \mathcal{I}_3} \hat{u}_j \right)^2, \quad (3.4.17)$$

where

$$\hat{u}_i = \hat{u}_i(\hat{g}, \hat{m}) = \hat{g}(t, \mathbf{W}_i) + \frac{\mathbb{1}(T = t)(Y_i - \hat{g}(t, \mathbf{W}_i))}{\hat{m}(t | \mathbf{W}_i)}.$$

The following theorem states that the variance estimator with nuisance functions $(\hat{g}_{\hat{\theta}}, \hat{m}_{\hat{\theta}})$ is consistent for the asymptotic variance in Theorem 3.4.3.

Theorem 3.4.4. *Under Assumptions 3.4.1 and 3.4.2, it holds that*

$$\hat{\mathcal{V}}_t(\hat{g}_{\hat{\theta}}, \hat{m}_{\hat{\theta}}) - \mathbb{V}_t(\mathbf{Z}_{\hat{\theta}}) \xrightarrow{P} 0$$

as $n \rightarrow \infty$.

3.4.2.2 Asymptotics of representation induced error

We now turn to the discussion of the second term in (3.4.15), i.e., the difference between the adjusted mean for the estimated representation $\mathbf{Z}_{\hat{\theta}}$ and the adjusted mean for the full covariate \mathbf{W} . Under sufficient regularity, the delta method [van der Vaart, 2000, Thm. 3.8] describes the distribution of this error:

3.4 Estimation based on outcome-adapted representations

Proposition 3.4.5. *Assume $(T, \mathbf{W}, Y) \sim P$ satisfies the model (3.4.12) with $\theta_P \in \Theta \subseteq \mathbb{R}^p$. Let $u: \Theta \rightarrow \mathbb{R}$ be the function given by $u(\theta) = \chi_t(\varphi(\theta, \mathbf{W}); P)$ and assume that u is differentiable in θ_P . Suppose that $\hat{\theta}$ is an estimator of θ_P with rate r_n such that*

$$r_n \cdot (\hat{\theta} - \theta_P) \xrightarrow{d} \mathbf{N}(0, \Sigma).$$

Then

$$r_n \cdot (\chi_t(\mathbf{Z}_{\hat{\theta}}) - \chi_t) \xrightarrow{d} \mathbf{N}(0, \nabla u(\theta_P)^\top \Sigma \nabla u(\theta_P))$$

as $n \rightarrow \infty$.

The delta method requires that the adjusted mean, $\chi_t = \chi_t(\varphi(\theta, \mathbf{W}); P)$, is differentiable with respect to $\theta \in \Theta$. The theorem below showcases that this is the case for the single-index model in Example 3.4.1.

Theorem 3.4.6. *Let $(T, \mathbf{W}, Y) \sim P$ be given by the single-index model in Example 3.4.1 with $h_t(\cdot) := h(t, \cdot) \in C^1(\mathbb{R})$. Assume that \mathbf{W} has a distribution with density $p_{\mathbf{W}}$ with respect to Lebesgue measure on \mathbb{R}^d and that $p_{\mathbf{W}}$ is continuous almost everywhere with bounded support. Assume also that the propensity $m(t | \mathbf{w}) = \mathbb{P}(T = t | \mathbf{W} = \mathbf{w})$ is continuous in \mathbf{w} .*

Then $u: \mathbb{R}^d \rightarrow \mathbb{R}$, defined by $u(\theta) = \chi_t(\mathbf{W}^\top \theta; P)$, is differentiable at $\theta = \theta_P$ with

$$\nabla u(\theta_P) = \mathbb{E}_P \left[h'_t(\mathbf{W}^\top \theta_P) \left(1 - \frac{\mathbb{P}(T = t | \mathbf{W})}{\mathbb{P}(T = t | \mathbf{W}^\top \theta_P)} \right) \mathbf{W} \right].$$

The theorem is stated with some restrictive assumptions that simplify the proof, but these are likely not necessary. It should also be possible to extend this result to more general models of the form (3.4.12), but we leave such generalizations for future work. In fact, the proof technique of Theorem 3.4.6 has already found application in Gnecco et al. [2023] in a different context.

The convergences implied in Theorem 3.4.3 and Proposition 3.4.5, with $r_n = \sqrt{|\mathcal{I}_3|}$, suggest that the MSE of the DOPE is of order

$$\begin{aligned} \mathbb{E}[(\hat{\chi}_t^{\text{dope}} - \chi_t)^2] &= \mathbb{E}[(\hat{\chi}_t^{\text{dope}} - \chi_t(\mathbf{Z}_{\hat{\theta}}))^2] + \mathbb{E}[(\chi_t(\mathbf{Z}_{\hat{\theta}}) - \chi_t)^2] \\ &\quad + 2 \mathbb{E}[(\hat{\chi}_t^{\text{dope}} - \chi_t(\mathbf{Z}_{\hat{\theta}}))(\chi_t(\mathbf{Z}_{\hat{\theta}}) - \chi_t)] \\ &\approx \frac{1}{|\mathcal{I}_3|} \left(\mathbb{E}[\mathbb{V}_t(\mathbf{Z}_{\hat{\theta}})] + \nabla u(\theta_P)^\top \Sigma \nabla u(\theta_P) \right) + o(|\mathcal{I}_3|^{-1}) \end{aligned} \quad (3.4.18)$$

The informal approximation ‘ \approx ’ can be turned into an equality by establishing (or simply assuming) uniform integrability, which enables the distributional convergences to be lifted to convergences of moments.

Remark 3.4.7. The expression in (3.4.18) suggests an approximate confidence interval of the form

$$\hat{\chi}_t^{\text{dope}} \pm \frac{z_{1-\alpha}}{\sqrt{|\mathcal{I}_3|}} (\hat{\mathcal{V}}_t + \hat{\mathbf{V}}^\top \hat{\Sigma} \hat{\mathbf{V}}),$$

where $\widehat{\Sigma}$ is a consistent estimator of the asymptotic variance of $\sqrt{|\mathcal{I}_3|}(\hat{\theta} - \theta_P)$, and where $\widehat{\nabla}$ is a consistent estimator of $\nabla_{\theta}\chi_t(\varphi(\theta, \mathbf{W}))|_{\theta=\theta_P}$. However, the requirement of constructing both $\widehat{\Sigma}$ and $\widehat{\nabla}$ adds further to the complexity of the methodology, so it might be preferable to rely on bootstrapping techniques in practice. In the simulation study we return to the question of inference and examine the coverage of a naive interval that does not include the term $\widehat{\nabla}^{\top}\widehat{\Sigma}\widehat{\nabla}$. \diamond

3.5 Experiments

In this section, we showcase the performance of the DOPE on simulated and real data. All code is publicly available on [GitHub](#)³.

3.5.1 Simulation study

We present results of a simulation study based on the single-index model from Example 3.4.1. We demonstrate the performance of the DOPE from Algorithms 3 and 4, and compare with various alternative estimators.

3.5.1.1 Sampling scheme

We simulated datasets consisting of n i.i.d. copies of (T, \mathbf{W}, Y) sampled according to the following scheme:

$$\begin{aligned} \mathbf{W} &= (W^1, \dots, W^d) \sim \text{Unif}(0, 1)^{\otimes d} \\ T | \mathbf{W} &\sim \text{Bern}(0.01 + 0.98 \cdot \mathbb{1}(W^1 > 0.5)) \\ Y | T, \mathbf{W}, \beta_Y &\sim \text{N}(h(T, \mathbf{W}^{\top}\beta_Y), 1) \end{aligned} \quad (3.5.19)$$

where β_Y is sampled once for each dataset with

$$\beta_Y = (1, \tilde{\beta}_Y), \quad \tilde{\beta}_Y \sim \text{N}(0, \mathbf{I}_{d-1}),$$

and where n, d , and h are experimental parameters. The settings considered were $n \in \{300, 900, 2700\}$, $d \in \{4, 12, 36\}$, and with h being one of

$$\begin{aligned} h_{\text{lin}}(t, z) &= t + 3z, & h_{\text{square}}(t, z) &= z^{1+t}, \\ h_{\text{cbrrt}}(t, z) &= (2 + t)z^{1/3}, & h_{\text{sin}}(t, z) &= (3 + t)\sin(\pi z). \end{aligned} \quad (3.5.20)$$

For each setting, $N = 300$ datasets were simulated.

Note that while $\mathbb{E}[T] = 0.01 + 0.98 \cdot \mathbb{P}(W^1 > 0.5) = 0.5$, the propensity score $m(t | \mathbf{w}) = 0.01 + 0.98 \cdot \mathbb{1}(w^1 > 0.5)$ takes the rather extreme values $\{0.01, 0.99\}$. Even though it technically satisfies (strict) positivity, these extreme values of the propensity makes the adjustment for \mathbf{W} a challenging task. For each dataset, the adjusted mean χ_1 (conditional on β_Y) was considered as the target parameter, and the ground truth was numerically computed as the sample mean of 10^7 observations of $h(1, \mathbf{W}^{\top}\beta_Y)$.

³<https://github.com/AlexanderChristgau/OutcomeAdaptedAdjustment>

3.5.1.2 Simulation estimators

This section contains an overview of the estimators used in the simulation. For a complete description see Section 3.B in the supplementary material.

Two settings were considered for outcome regression (OR):

- **Linear:** Ordinary Least Squares (OLS).
- **Neural network:** A feedforward neural network with two hidden layers: a linear layer with one neuron, followed by a fully connected ReLU-layer with 100 neurons. The first layer is a linear bottleneck that enforces the single-index model, and we denote the weights by $\theta \in \mathbb{R}^d$. An illustration of the architecture can be found in the supplementary Section 3.B. For a further discussion of leaning single- and multiple-index models with neural networks, see Parkinson et al. [2023] and references therein.

For propensity score estimation, logistic regression was used across all settings. A ReLU-network with one hidden layer with 100 neurons was also considered for estimation of the propensity score, but it did not enhance the performance of the resulting ATE estimators in this setting. Random forests and other methods were also initially used for both outcome regression and propensity score estimation. They were subsequently excluded because they did not seem to yield any noteworthy insights beyond what was observed for the methods above.

For each outcome regression, two implementations were explored: a stratified regression where Y is regressed onto \mathbf{W} separately for each stratum $T = 1$ and $T = 2$, and a joint regression where Y is regressed onto (T, \mathbf{W}) simultaneously. In the case of joint regression, the neural network architecture represents the regression function as a single index model, given by $Y = h(\alpha T + \theta^\top \mathbf{W}) + \varepsilon_Y$. This representation differs from the description of the regression function specified in the sample scheme (3.5.19), which allows for a more complex interaction between treatment and the index of the covariates. In this sense, the joint regression is misspecified.

Based on these methods for nuisance estimation, we considered the following estimators of the adjusted mean:

- The regression estimator $\hat{\chi}_1^{\text{reg}} := \mathbb{P}_n[\hat{g}(1, \mathbf{W})]$.
- The AIPW estimator given in Equation (3.4.10).
- The DOPE from Algorithm 3, where $\mathbf{Z}_{\hat{\theta}} = \mathbf{W}^\top \hat{\theta}$ and where $\hat{\theta}$ contains the weights of the first layer of the neural network designed for single-index regression. We refer to this estimator as the DOPE-IDX.
- The DOPE-BCL described in Remark 3.4.2, where the propensity score is based on the final outcome regression. See also Benkeser et al. [2020, App. D].

3 Efficient adjustment for complex covariates: Gaining efficiency with DOPE

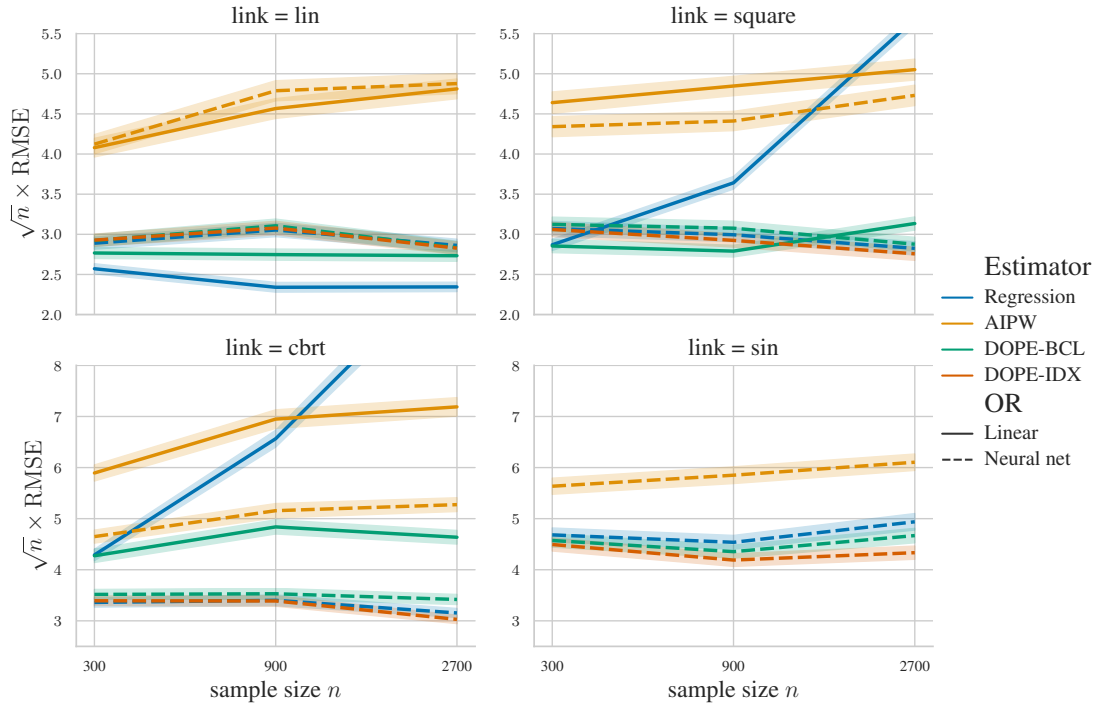


Figure 3.5.3: Root mean square errors for various estimators of χ_1 plotted against sample size. Each data point is an average over 900 datasets, 300 for each choice of $d \in \{4, 12, 36\}$. The bands around each line correspond to ± 1 standard error. For this plot, the outcome regression (OR) was fitted separately for each stratum $T = 1$ and $T = 2$.

We considered two versions of each DOPE estimator: one without sample splitting and another using 4-fold cross-fitting⁴. For the latter, the final empirical mean is calculated based on the hold-out fold, while the nuisance parameters are fitted using the remaining three folds. For each fold $k = 1, \dots, 4$, this means employing Algorithm 3 with $\mathcal{I}_1 = \mathcal{I}_2 = [n] \setminus J_k$ and $\mathcal{I}_3 = J_k$. The two versions showed comparable performance for larger samples. In this section we only present the results for the DOPE without sample splitting, which performed better overall in our simulations. However, a comparison with the cross-fitted version can be found in Section 3.B in the supplement.

Numerical results for the IPW estimator $\hat{\chi}_1^{\text{ipw}} := \mathbb{P}_n[\mathbb{1}(T = t)Y/\hat{m}(1 | \mathbf{W})]$ were also gathered. It performed poorly in our settings, and we have omitted the results for the sake of clarity in the presentation.

3.5.1.3 Empirical performance of estimators

The results for stratified outcome regression are shown in Figure 3.5.3. Each panel, corresponding to a link function in (3.5.20), displays the RMSE for various estimators against sample size. Across different values of $d \in \{4, 12, 36\}$, the results remain consistent, and thus the RMSE is averaged over the 3×300 datasets with varying d . In the upper left panel, where the link is linear, the regression estimator $\hat{\chi}_1^{\text{reg}}$ with OLS performs best as expected. The remaining estimators exhibit similar performance, except for AIPW, which consistently performs poorly across all settings. This can be attributed to extreme values of the propensity score, resulting in a large efficiency bound for AIPW. All estimators seem to maintain an approximate \sqrt{n} -consistent RMSE for the linear link as anticipated.

For the nonlinear links, we observe that the OLS-based regression estimators perform poorly and do not exhibit approximate \sqrt{n} -consistency. For the sine link, the RMSEs for all OLS-based estimators are large, and they are not shown for ease of visualization. For neural network outcome regression, the AIPW still performs the worst across all nonlinear links. The regression estimator and the DOPE estimators share similar performance when used with the neural network, with DOPE-IDX being more accurate overall. Since the neural network architecture is tailored to the single-index model, it is not surprising that the outcome regression works well, and as a result there is less need for debiasing. On the other hand, the debiasing introduced in the DOPE does not hurt the accuracy, and in fact, improves it in this setting.

The results for joint regression of Y on (T, \mathbf{W}) are shown in Figure 3.5.4. The results for the OLS-based estimators provide similar insights as previously discussed, so we focus on the results for the neural network based estimators. The jointly trained neural network is, in a sense, misspecified for the single-index model (except for the linear link), as discussed in Section 3.5.1.2. Thus it is not surprising that the regression estimator fails to adjust effectively for larger sample sizes. What is somewhat unexpected, however, is that the precision of DOPE, especially the DOPE-IDX, does not appear to be compromised by the misspecified outcome regression. In fact, the DOPE-IDX even seems to perform better with joint outcome regression. We suspect that this could be attributed to the joint regression producing more robust predictions for the rare treated subjects with $W^1 \leq 0.5$, for which $m(1 | \mathbf{W}) = 0.01$. The predictions are more robust since the joint regression can leverage some of the information from the many untreated subjects with $W^1 \leq 0.5$ at the cost of introducing systematic bias, which the DOPE-IDX deals with in the debiasing step. While this phenomenon is interesting, a thorough exploration of its exact details, both numerically and theoretically, is a task we believe is better suited for future research.

In summary, the DOPE serves as a middle ground between the regression estimator and the AIPW. It provides an additional safeguard against biased outcome regression, all while avoiding the potential numerical instability entailed by using standard inverse propensity weights.

⁴consistent with Rem. 3.1. in Chernozhukov et al. [2018], which recommends using 4 or 5 folds for cross-fitting.

3 Efficient adjustment for complex covariates: Gaining efficiency with DOPE




Figure 3.5.4: Root mean square errors for various estimators of χ_1 plotted against sample size. Each data point is an average over 900 datasets, 300 for each choice of $d \in \{4, 12, 36\}$. The bands around each line correspond to ± 1 standard error. For this plot, the outcome regression (OR) was fitted jointly onto (T, \mathbf{W}) .

3.5.1.4 Inference

We now consider approximate confidence intervals obtained from the empirical variance estimator \hat{V} defined in (3.4.17). Specifically, we consider intervals of the form $\hat{\chi}_1 \pm \frac{z_{0.975}}{\sqrt{n}} \hat{V}^{-1/2}$, where $z_{0.975}$ is the 0.975 quantile of the standard normal distribution.

Figure 3.5.5 shows the empirical coverage of χ_1 for the AIPW, DOPE-BCL and DOPE-IDX, all based on neural network outcome regression fitted jointly. Based on the number of repetitions, we can expect the estimated coverage probabilities to be accurate up to ± 0.01 after rounding.

The left column shows the results for no sample splitting, whereas the right column shows the results for 4-fold cross-fitting. For all estimators, cross-fitting seems to improve coverage slightly for large sample sizes. The AIPW is a little anti-conservative, especially for the links h_{sin} and h_{cbrt} , but overall maintains approximate coverage. The slight miscoverage is to be expected given the challenging nature with the extreme propensity weights. For the DOPE-BCL and DOPE-IDX, the coverage probabilities are worse, but with the DOPE-IDX having better coverage overall. According to our asymptotic analysis, specifically Theorem 3.4.3, the DOPE-BCL and DOPE-IDX intervals will only achieve asymptotic 95% coverage of the *random* parameters $\chi_1(\hat{h}(t, \mathbf{Z}_{\hat{\theta}}))$ and $\chi_1(\mathbf{Z}_{\hat{\theta}})$, respectively. Thus we do not anticipate these intervals to asymptotically achieve the nominal 95% coverage of χ_1 . However, we show in Figure 3.B.3 in the supplementary Section 3.B that the average widths of the confidence intervals are significantly smaller for the DOPE intervals than the AIPW interval. Hence it is possible that the DOPE intervals can be corrected – as we have also discussed in Remark 3.4.7 – to gain advantage over the AIPW interval, but we leave such explorations for future work.

3 Efficient adjustment for complex covariates: Gaining efficiency with DOPE




Figure 3.5.5: Coverage of asymptotic confidence intervals of the adjusted mean χ_1 , aggregated over $d \in \{4, 12, 36\}$ and with the true β given in (3.5.19).

In summary, the lack of full coverage for the naively constructed intervals are consistent with the asymptotic analysis conditionally on a representation $\mathbf{Z}_{\hat{\theta}}$. In view of this, our asymptotic results might be more realistic and pragmatic than the previously studied regime [Benkeser et al., 2020, Ju et al., 2020], which assumes fast unconditional rates of the nuisance parameters $(\hat{m}_{\hat{\theta}}, \hat{g}_{\hat{\theta}})$.

3.5.2 Application to NHANES data

We consider the mortality dataset collected by the National Health and Nutrition Examination Survey I Epidemiologic Followup Study [Cox et al., 1997], henceforth referred to as the NHANES dataset. The dataset was initially collected as in Lundberg et al. [2020]⁵. The dataset contains several baseline covariates, and our outcome variable is the indicator of death at the end of study. To deal with missing values, we considered both a complete mean imputation as in Lundberg et al. [2020], or a trimmed dataset where covariates with more than 50% of their values missing are dropped and the rest are mean imputed. The latter approach reduces the number of covariates from 79 to 65. The final results were similar for the two imputation methods, so we only report the results for the mean imputed dataset here. In the supplementary Section 3.B we show the results for the trimmed dataset.

The primary aim of this section is to evaluate the different estimation methodologies on a realistic and real data example. For this purpose we consider a treatment variable based on *pulse pressure* and study its effect on mortality. Pulse pressure is defined as the difference between systolic and diastolic blood pressure (BP). High pulse pressure is not only used as an indicator of disease but is also reported to increase the risk of cardiovascular diseases [Franklin et al., 1999]. We investigate the added effect of high pulse pressure when adjusting for other baseline covariates, in particular systolic BP and levels of white blood cells, hemoglobin, hematocrit, and platelets. We do not adjust for diastolic BP, as it determines the pulse pressure when combined with systolic BP, which would therefore lead to a violation of positivity.

A pulse pressure of 40 mmHg is considered normal, and as the pressure increases past 50 mmHg, the risk of cardiovascular diseases is reported to increase. Some studies have used 60 mmHg as the threshold for high pulse pressure [Homan et al., 2024], and thus we consider the following treatment variable corresponding to high pulse pressure:

$$T = \mathbb{1}(\text{pulse pressure} > 60 \text{ mmHg}).$$

For the binary outcome regression we consider logistic regression and a variant of the neural network from Section 3.5.1.2 with an additional ‘sigmoid-activation’ on the output layer. Based on 5-fold cross-validation, logistic regression yields a log-loss of 0.50, whereas the neural network yields a log-loss of 0.48 (smaller is better).

Figure 3.5.6 shows the distribution of the estimated propensity scores used in the AIPW, DOPE-IDX, and DOPE-BCL, based on neural network regression. As expected,

⁵See <https://github.com/suinleelab/treeexplainer-study> for their GitHub repository.

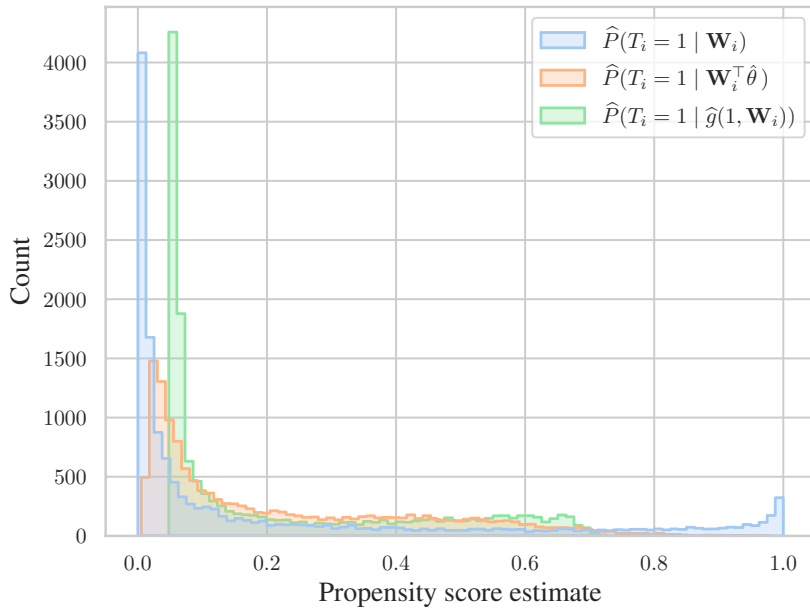


Figure 3.5.6: Distribution of estimated propensity scores based on the full covariate set \mathbf{W} , the single-index $\mathbf{W}^\top \hat{\theta}$, and the outcome predictions $\hat{g}(1, \mathbf{W})$, respectively.

we see that the full covariate set yields propensity scores that essentially violate positivity, with this effect being less pronounced for DOPE-IDX and even less so for DOPE-BCL, where the scores are comparatively less extreme.

Table 3.5.1 shows the estimated treatment effects $\hat{\Delta} = \hat{\chi}_1 - \hat{\chi}_0$ for various estimators together with 95% bootstrap confidence intervals. The estimators are sorted according to bootstrap variance based on $B = 1000$ bootstraps. Because we do not know the true effect, we cannot directly decide which estimator is best. The naive contrast provides a substantially larger estimate than all of the adjustment estimators, which indicates that the added effect of high pulse pressure on mortality when adjusting is different from the unadjusted effect. The IPW estimator, $\hat{\chi}_1^{\text{ipw}}$, is the only estimator to yield a negative estimate of the adjusted effect, and its bootstrap variance is also significantly larger than the other estimators. Thus it is plausible, that the IPW estimator fails to adjust appropriately. This is not surprising given extreme distribution of the propensity weights, as shown in blue in Figure 3.5.6.

The remaining estimators yield comparable estimates of the adjusted effect, with the logistic regression based estimators having a marginally larger estimates than their neural network counterparts. The (bootstrap) standard errors are comparable for the DOPE and regression estimators, but the AIPW estimators have slightly larger standard errors. As a result, the added effect of high pulse pressure on mortality cannot be considered statistically significant for the AIPW estimators, whereas it can for the other estimators.

In summary, the results indicate that the choice treatment effect estimator impacts the

Estimator	Estimate	BS se	BS CI
Regr. (NN)	0.022	0.009	(0.004, 0.037)
Naive contrast	0.388	0.010	(0.369, 0.407)
Regr. (Logistic)	0.026	0.010	(0.012, 0.049)
DOPE-BCL (Logistic)	0.023	0.010	(0.003, 0.040)
DOPE-BCL (NN)	0.021	0.010	(0.001, 0.039)
DOPE-IDX (NN)	0.024	0.012	(0.002, 0.048)
AIPW (NN)	0.023	0.015	(-0.009, 0.051)
AIPW (Logistic)	0.026	0.015	(-0.004, 0.056)
IPW (Logistic)	-0.040	0.025	(-0.103, -0.006)

Table 3.5.1: Adjusted effect estimates of unhealthy pulse pressure on mortality based on NHANES dataset. Estimators are sorted in increasing order after bootstrap (BS) standard error.

final estimate and confidence interval. While it is uncertain which estimator is superior for this particular application, the DOPE estimators seem to offer a reasonable and stable estimate.

3.6 Discussion

In this paper, we address challenges posed by complex data with unknown underlying structure, an increasingly common scenario in observational studies. Specifically, we have formulated a refined and general formalism for studying efficiency of covariate adjustment. This formalism extends and builds upon the efficiency principles derived from causal graphical models, in particular results provided by Rotnitzky and Smucler [2020]. Our theoretical framework has led to the identification of the optimal covariate information for adjustment. From this theoretical groundwork, we introduced the DOPE for general estimation of the ATE with increased efficiency.

Several key areas within this research merit further investigation:

Extension to other targets and causal effects: Our focus has predominantly been on the adjusted mean $\chi_t = \mathbb{E}[\mathbb{E}[Y | T = t, \mathbf{W}]]$. However, the extension of our methodologies to continuous treatments, instrumental variables, or other causal effects, such as the average treatment effect among the treated, is an area that warrants in-depth exploration. We suspect that many of the fundamental ideas can be modified to prove analogous results for such target parameters.

Beyond neural networks: While our examples and simulation study has applied the DOPE with neural networks, its compatibility with other regression methods such as kernel regression and gradient boosting opens new avenues for alternative adjustment estimators. Investigating such estimators can provide practical insights and broaden the applicability of our approach. The Outcome Highly Adapted Lasso [Ju et al., 2020] is

a similar method in this direction that leverages the Highly Adaptive Lasso [Benkeser and van der Laan, 2016] for robust treatment effect estimation.

Causal representation learning integration: Another possible direction for future research is the integration of our efficiency analysis into the broader concept of causal representation learning [Schölkopf et al., 2021]. This line of research, which is not concerned with specific downstream tasks, could potentially benefit from some of the insights of our efficiency analysis and the DOPE framework.

Implications of sample splitting: Our current asymptotic analysis relies on the implementation of three distinct sample splits. Yet, our simulation results suggest that sample splitting might not be necessary. Using the notation of Algorithm 3, we hypothesize that the independence assumption $\mathcal{I}_3 \cap (\mathcal{I}_1 \cup \mathcal{I}_2) = \emptyset$ could be replaced with Donsker class conditions. However, the theoretical justifications for setting \mathcal{I}_1 identical to \mathcal{I}_2 are not as evident, thus meriting additional exploration into this potential simplification.

Valid and efficient confidence intervals: The naive confidence intervals explored in the simulation study, cf. Section 3.5.1.4, do not provide an adequate coverage of the unconditional adjusted mean. As indicated in Remark 3.4.7, it would be interesting to explore methods for correcting the intervals, for example by: (i) adding a bias correction, (ii) constructing a consistent estimator of the unconditional variance, or (iii) using bootstrapping. Appealing to bootstrapping techniques, however, might be computationally expensive, in particular when combined with cross-fitting. To implement (i), one approach could be to generalize Theorem 3.4.6 to establish general differentiability of the adjusted mean with respect to general representations, and then construct an estimator of the gradient. However, one must assess whether this bias correction could negate the efficiency benefits of the DOPE.

Acknowledgments

We thank Leonard Henckel and Anton Rask Lundborg for helpful discussions. AMC and NRH were supported by a research grant (NNF20OC0062897) from Novo Nordisk Fonden.

Supplement to ‘Efficient adjustment for complex covariates: Gaining efficiency with DOPE’

The supplementary material is organized as follows: Section 3.A contains proofs of the results in the main text and a few auxiliary results; Section 3.B contains details on the experiments in Section 3.5; and Section 3.C contains a description of the cross-fitting algorithm.

3.A. Auxiliary results and proofs

The proposition below relaxes the criteria for being a valid adjustment set in a causal graphical model. It follows from the results of Perković et al. [2018], which was pointed out to the authors by Leonard Henckel, and is stated here for completeness.

Proposition 3.A.1. *Let T, \mathbf{Z} , and Y be pairwise disjoint node sets in a DAG $\mathcal{D} = (\mathbf{V}, \mathbf{E})$, and let $\mathcal{M}(\mathcal{D})$ denote the collection of continuous distributions that are Markovian with respect to \mathcal{D} and with $\mathbb{E}|Y| < \infty$. Then \mathbf{Z} is an adjustment set relative to (T, Y) if and only if for all $P \in \mathcal{M}(\mathcal{D})$ and all t*

$$\mathbb{E}_P[Y \mid \text{do}(T = t)] = \mathbb{E}_P \left[\frac{\mathbb{1}(T = t)Y}{\mathbb{P}_P(T = t \mid \text{pa}_{\mathcal{D}}(T))} \right] = \mathbb{E}_P[\mathbb{E}_P[Y \mid T = t, \mathbf{Z}]]$$

Proof. The ‘only if’ direction is straightforward if we use an alternative characterization of \mathbf{Z} being an adjustment set: for any density p of a distribution $P \in \mathcal{P}$ it holds that

$$p(y \mid \text{do}(T = t)) = \int p(\mathbf{y} \mid \mathbf{z}, t) p(\mathbf{z}) d\mathbf{z},$$

where the ‘do-operator’ (and notation) is defined as in, e.g., Peters et al. [2017]. Fubini’s theorem then yields

$$\begin{aligned} \mathbb{E}_P[Y \mid \text{do}(T = t)] &= \int_{\mathbb{R}^l} \mathbf{y} p(\mathbf{y} \mid \text{do}(\mathbf{T} = \mathbf{t})) d\mathbf{y} \\ &= \int_{\mathbb{R}^l} \mathbf{y} \int_{\mathbb{R}^k} p(\mathbf{y} \mid \mathbf{z}, t) p(\mathbf{z}) d\mathbf{z} d\mathbf{y} \\ &= \int_{\mathbb{R}^k} \int_{\mathbb{R}^l} \mathbf{y} p(\mathbf{y} \mid \mathbf{z}, t) d\mathbf{y} p(\mathbf{z}) d\mathbf{z} \\ &= \mathbb{E}_P[\mathbb{E}_P[Y \mid T = t, \mathbf{Z}]], \end{aligned}$$

which is equivalent to the ‘only if’ part. On the contrary, assume that \mathbf{Z} is *not* an adjustment set for (T, Y) . Then Theorem 56 and the proof of Theorem 57 in Perković et al. [2018] imply the existence of a Gaussian distribution $\tilde{P} \in \mathcal{P}$ for which

$$\mathbb{E}_{\tilde{P}}[Y \mid \text{do}(T = 1)] \neq \mathbb{E}_{\tilde{P}}[\mathbb{E}_{\tilde{P}}[Y \mid T = 1, \mathbf{Z}]]$$

This implies the other direction. □

3 Efficient adjustment for complex covariates: Gaining efficiency with DOPE

The following lemma is a generalization of Lemma 27 in Rotnitzky and Smucler [2020].

Lemma 3.A.2. *For any σ -algebras $\mathcal{Z}_1 \subseteq \mathcal{Z}_2 \subseteq \sigma(\mathbf{W})$ it holds, under Assumption 3.1.1, that*

$$\mathbb{E}_P [\pi_t(\mathcal{Z}_2; P)^{-1} | T = t, \mathcal{Z}_1] = \pi_t(\mathcal{Z}_1; P)^{-1}$$

for all $P \in \mathcal{P}$.

Proof. Direct computation yields

$$\begin{aligned} & \mathbb{E}_P [\pi_t(\mathcal{Z}_2; P)^{-1} | T = t, \mathcal{Z}_1] \mathbb{P}_P(T = t | \mathcal{Z}_1) \\ &= \mathbb{E}_P \left[\frac{\mathbb{1}(T = t)}{\pi_t(\mathcal{Z}_2; P)} | \mathcal{Z}_1 \right] = \mathbb{E}_P \left[\frac{\mathbb{E}_P[\mathbb{1}(T = t) | \mathcal{Z}_2]}{\pi_t(\mathcal{Z}_2; P)} | \mathcal{Z}_1 \right] = 1. \end{aligned}$$

□

3.A.1 Proof of Lemma 3.3.1

Since $P \in \mathcal{P}$ is fixed, we suppress it from the notation in the following computations.

From $Y \perp\!\!\!\perp \mathcal{Z}_2 | \mathcal{Z}_1, T$ and $\mathcal{Z}_1 \subseteq \mathcal{Z}_2$, it follows that $\mathbb{E}[Y | \mathcal{Z}_1, T] = \mathbb{E}[Y | \mathcal{Z}_2, T]$. Hence, for each $t \in \mathbb{T}$,

$$b_t(\mathcal{Z}_1) = \mathbb{E}[Y | \mathcal{Z}_1, T = t] = \mathbb{E}[Y | \mathcal{Z}_2, T = t] = b_t(\mathcal{Z}_2). \quad (3.A.1)$$

Therefore $\chi_t(\mathcal{Z}_1) = \mathbb{E}[b_t(\mathcal{Z}_1)] = \mathbb{E}[b_t(\mathcal{Z}_2)] = \chi_t(\mathcal{Z}_2)$. If \mathcal{Z}_2 is description of \mathbf{W} , this identity shows that \mathcal{Z}_2 is \mathcal{P} -valid if and only if \mathcal{Z}_1 is \mathcal{P} -valid.

To compare the asymptotic efficiency bounds for \mathcal{Z}_2 with \mathcal{Z}_1 we use the law of total variance. Note first that

$$\begin{aligned} \mathbb{E}[\psi_t(\mathcal{Z}_2) | T, Y, \mathcal{Z}_1] &= \mathbb{E} \left[\mathbb{1}(T = t)(Y - b_t(\mathcal{Z}_2)) \frac{1}{\pi_t(\mathcal{Z}_2)} + b_t(\mathcal{Z}_2) - \chi_t | T, Y, \mathcal{Z}_1 \right] \\ &= \mathbb{1}(T = t)(Y - b_t(\mathcal{Z}_1)) \mathbb{E} \left[\frac{1}{\pi_t(\mathcal{Z}_2)} | T = t, Y, \mathcal{Z}_1 \right] + b_t(\mathcal{Z}_1) - \chi_t \\ &= \mathbb{1}(T = t)(Y - b_t(\mathcal{Z}_1)) \mathbb{E} \left[\frac{1}{\pi_t(\mathcal{Z}_2)} | T = t, \mathcal{Z}_1 \right] + b_t(\mathcal{Z}_1) - \chi_t \\ &= \mathbb{1}(T = t)(Y - b_t(\mathcal{Z}_1)) \frac{1}{\pi_t(\mathcal{Z}_1)} + b_t(\mathcal{Z}_1) - \chi_t = \psi_t(\mathcal{Z}_1). \end{aligned}$$

Second equality is due to \mathcal{Z}_1 -measurability and (3.A.1), whereas the third equality follows from $Y \perp\!\!\!\perp \mathcal{Z}_2 | \mathcal{Z}_1, T$ and the last inequality is due to Lemma 3.A.2. On the other hand,

$$\begin{aligned} & \mathbb{E}[\text{Var}(\psi_t(\mathcal{Z}_2) | T, Y, \mathcal{Z}_1)] \\ &= \mathbb{E} \left[\text{Var} \left(\mathbb{1}(T = t)(Y - b_t(\mathcal{Z}_1)) \frac{1}{\pi_t(\mathcal{Z}_2)} + b_t(\mathcal{Z}_1) - \chi_t(P) | T, Y, \mathcal{Z}_1 \right) \right] \\ &= \mathbb{E} \left[\mathbb{1}(T = t)(Y - b_t(\mathcal{Z}_1))^2 \text{Var} \left(\frac{1}{\pi_t(\mathcal{Z}_2)} | T = t, Y, \mathcal{Z}_1 \right) \right] \\ &= \mathbb{E} \left[\pi_t(\mathcal{Z}_1) \text{Var}(Y | T = t, \mathcal{Z}_1) \text{Var} \left(\frac{1}{\pi_t(\mathcal{Z}_2)} | T = t, \mathcal{Z}_1 \right) \right]. \end{aligned}$$

Combined we have that

$$\begin{aligned}\mathbb{V}_t(\mathcal{Z}_2) &= \text{Var}(\psi_t(\mathcal{Z}_2)) = \text{Var}(\mathbb{E}[\psi_t(\mathcal{Z}_2) | T, Y, \mathcal{Z}_1]) + \mathbb{E}[\text{Var}[\psi_t(\mathcal{Z}_2) | T, Y, \mathcal{Z}_1]] \\ &= \mathbb{V}_t(\mathcal{Z}_1) + \mathbb{E}\left[\pi_t(\mathcal{Z}_1) \text{Var}(Y | T = t, \mathcal{Z}_1) \text{Var}\left(\frac{1}{\pi_t(\mathcal{Z}_2)} \middle| T = t, Y, \mathcal{Z}_1\right)\right].\end{aligned}$$

To prove the last part of the lemma, we first note that $\psi_t(\mathcal{Z}_2)$ and $\psi_{t'}(\mathcal{Z}_2)$ are conditionally uncorrelated for $t \neq t'$. This follows from (3.A.1), from which the conditional covariance reduces to

$$\text{Cov}(\psi_t(\mathcal{Z}_2), \psi_{t'}(\mathcal{Z}_2) | T, Y, \mathcal{Z}_1) = \mathbb{1}(T = t)\mathbb{1}(T = t') \text{Cov}(*, * | T, Y, \mathcal{Z}_1) = 0.$$

Thus, letting $\psi(\mathcal{Z}_2) = (\psi_t(\mathcal{Z}_2))_{t \in \mathbb{T}}$ and using the computation for \mathbb{V}_t , we obtain that:

$$\begin{aligned}\mathbb{V}_\Delta(\mathcal{Z}_2) &= \text{Var}(\mathbf{c}^\top \psi(\mathcal{Z}_2)) \\ &= \text{Var}(\mathbb{E}[\mathbf{c}^\top \psi(\mathcal{Z}_2) | T, Y, \mathcal{Z}_1]) + \mathbb{E}[\text{Var}[\mathbf{c}^\top \psi(\mathcal{Z}_2) | T, Y, \mathcal{Z}_1]] \\ &= \mathbb{V}_\Delta(\mathcal{Z}_1) + \sum_{t \in \mathbb{T}} c_t^2 \mathbb{E}\left[\pi_t(\mathcal{Z}_1) \text{Var}(Y | T = t, \mathcal{Z}_1) \text{Var}\left(\frac{1}{\pi_t(\mathcal{Z}_2)} \middle| T = t, Y, \mathcal{Z}_1\right)\right].\end{aligned}$$

This concludes the proof. \square

3.A.2 Proof of Theorem 3.3.3

(i) Assume that $\mathcal{Z} \subseteq \sigma(\mathbf{W})$ is a description of \mathbf{W} such that $\mathcal{Q}_P \subseteq \bar{\mathcal{Z}}$. Observe that almost surely,

$$\mathbb{P}_P(Y \leq y | T, \mathbf{W}) = \sum_{t \in \mathbb{T}} \mathbb{1}(T = t) F(y | t, \mathbf{W}; P)$$

is $\sigma(T, \mathcal{Q}_P)$ -measurable and hence also $\sigma(T, \bar{\mathcal{Z}})$ -measurable. It follows that $\mathbb{P}_P(Y \leq y | T, \mathbf{W})$ is also a version of $\mathbb{P}_P(Y \leq y | T, \bar{\mathcal{Z}})$. As a consequence, $\mathbb{P}_P(Y \leq y | T, \mathbf{W}) = \mathbb{P}_P(Y \leq y | T, \mathcal{Z})$ almost surely. From Doob's characterization of conditional independence [Kallenberg, 2021, Theorem 8.9], we conclude that $Y \perp\!\!\!\perp_P \mathbf{W} | T, \mathcal{Z}$, or equivalently that \mathcal{Z} is P -ODS. The 'in particular' follows from setting $\mathcal{Z} = \mathcal{Q}_P$.

(ii) Assume that $\mathcal{Z} \subseteq \sigma(\mathbf{W})$ is P -ODS. Using Doob's characterization of conditional independence again, $\mathbb{P}_P(Y \leq y | T, \mathbf{W}) = \mathbb{P}_P(Y \leq y | T, \mathcal{Z})$ almost surely. Under Assumption 3.1.1, this entails that $F(y | T = t, \mathbf{W}; P) = \mathbb{P}_P(Y \leq y | T = t, \mathcal{Z})$ almost surely, and hence $F(y | T = t, \mathbf{W}; P)$ must be $\bar{\mathcal{Z}}$ -measurable. As the generators of \mathcal{Q}_P are $\bar{\mathcal{Z}}$ -measurable, we conclude that $\mathcal{Q}_P \subseteq \bar{\mathcal{Z}}$.

(iii) Let \mathcal{Z} be a P -ODS description. From (i) and (ii) it follows that $\mathcal{Q}_P \subseteq \bar{\mathcal{Z}}$, and since $\mathcal{Z} \subseteq \sigma(\mathbf{W})$ it also holds that $Y \perp\!\!\!\perp_P \bar{\mathcal{Z}} | T, \mathcal{Q}_P$. Since $\mathbb{V}_\Delta(\mathcal{Z}; P) = \mathbb{V}_\Delta(\bar{\mathcal{Z}}; P)$, Lemma 3.3.1 gives the desired conclusion. \square

3.A.3 Proof of Corollary 3.3.4

Theorem 3.3.3 (i,ii) implies that \mathcal{Z} is P -ODS if and only if $\bar{\mathcal{Z}}^P$ contains \mathcal{Q}_P for all $P \in \mathcal{P}$.

3 Efficient adjustment for complex covariates: Gaining efficiency with DOPE

If $\mathcal{Q} \subseteq \bar{\mathcal{Z}}^P$ for all $P \in \mathcal{P}$ then $\mathcal{Q}_P \subseteq \bar{\mathcal{Z}}^P$ for all $P \in \mathcal{P}$ by definition, and \mathcal{Z} is \mathcal{P} -ODS. Equation (3.3.9) follows from the same argument as Theorem 3.3.3 (iii). That is, $\mathcal{Q} \subseteq \bar{\mathcal{Z}}^P$, and since \mathcal{Q} is P -ODS and $\mathcal{Z} \subseteq \sigma(\mathbf{W})$ we have $Y \perp\!\!\!\perp_P \bar{\mathcal{Z}}^P \mid T, \mathcal{Q}$. Lemma 3.3.1 can be applied for each $P \in \mathcal{P}$ to obtain the desired conclusion. \square

3.A.4 Proof of Proposition 3.3.6

Since $\mathcal{R}_P \subseteq \mathcal{Q}_P$ always holds, so it suffices to prove that $\mathcal{Q}_P \subseteq \mathcal{R}_P$. If $F(y \mid t, \mathbf{W}; P)$ is $\sigma(b_t(\mathbf{W}; P))$ -measurable, then $\mathcal{Q}_P \subseteq \mathcal{R}_P$ follows directly from definition. If Y is a binary outcome, then the conditional distribution is determined by the conditional mean, i.e., $F(y \mid t, \mathbf{W}; P)$ is $\sigma(b_t(\mathbf{W}; P))$ -measurable. If $Y = b_T(\mathbf{W}) + \varepsilon_Y$ with $\varepsilon_Y \perp\!\!\!\perp T, \mathbf{W}$, then $F(y \mid t, \mathbf{w}; P) = F_P(y - b_t(\mathbf{w}))$ where F_P is the distribution function for the distribution of ε_Y . Again, $F(y \mid t, \mathbf{W}; P)$ becomes $\sigma(b_t(\mathbf{W}; P))$ -measurable.

To prove the last part, let \mathcal{Z} be a P -OMS description. Since $b_t(\mathbf{W}; P) = b_t(\mathcal{Z}; P)$ is \mathcal{Z} -measurable for every $t \in \mathbb{T}$, it follows that $\mathcal{R}_P \subseteq \bar{\mathcal{Z}}$. Hence the argument of Theorem 3.3.3 (iii) establishes Equation (3.3.8) when $\mathcal{Q}_P = \mathcal{R}_P$. \square

3.A.5 Proof of Lemma 3.3.8

Since $P \in \mathcal{P}$ is fixed throughout this proof, we suppress it from notation. From $\mathcal{Z}_1 \subseteq \mathcal{Z}_2 \subseteq \sigma(\mathbf{W})$ and $T \perp\!\!\!\perp_P \mathcal{Z}_2 \mid \mathcal{Z}_1$, it follows that $\pi_t(\mathcal{Z}_1) = \pi_t(\mathcal{Z}_2)$, and hence

$$\chi_t(\mathcal{Z}_1) = \mathbb{E}[\pi_t(\mathcal{Z}_1)^{-1} \mathbb{1}(T = t) Y] = \mathbb{E}[\pi_t(\mathcal{Z}_2)^{-1} \mathbb{1}(T = t) Y] = \chi_t(\mathcal{Z}_2).$$

This establishes the first part. For the second part, we use that $\pi_t(\mathcal{Z}_1) = \pi_t(\mathcal{Z}_2)$ and $\chi_t(\mathcal{Z}_1) = \chi_t(\mathcal{Z}_2)$ to see that

$$\begin{aligned} \psi_t(\mathcal{Z}_1) - \psi_t(\mathcal{Z}_2) &= b_t(\mathcal{Z}_1) - b_t(\mathcal{Z}_2) + \frac{\mathbb{1}(T = t)}{\pi_t(\mathcal{Z}_2)} (b_t(\mathcal{Z}_2) - b_t(\mathcal{Z}_1)) \\ &= \left(\frac{\mathbb{1}(T = t)}{\pi_t(\mathcal{Z}_2)} - 1 \right) (b_t(\mathcal{Z}_2) - b_t(\mathcal{Z}_1)) =: R_t(\mathcal{Z}_1, \mathcal{Z}_2). \end{aligned}$$

Since $\mathbb{E}[R_t(\mathcal{Z}_1, \mathcal{Z}_2) \mid \mathcal{Z}_2] = 0$ we have that $\mathbb{E}[\psi_t(\mathcal{Z}_2) R_t(\mathcal{Z}_1, \mathcal{Z}_2)] = 0$, which means that $\psi_t(\mathcal{Z}_2)$ and $R_t(\mathcal{Z}_1, \mathcal{Z}_2)$ are uncorrelated. Note that from $\pi_t(\mathcal{Z}_1) = \pi_t(\mathcal{Z}_2)$ it also follows that

$$b_t(\mathcal{Z}_1) = \mathbb{E}\left[\frac{Y \mathbb{1}(T = t)}{\pi_t(\mathcal{Z}_1)} \mid \mathcal{Z}_1\right] = \mathbb{E}\left[\mathbb{E}\left[\frac{Y \mathbb{1}(T = t)}{\pi_t(\mathcal{Z}_2)} \mid \mathcal{Z}_2\right] \mid \mathcal{Z}_1\right] = \mathbb{E}[b_t(\mathcal{Z}_2) \mid \mathcal{Z}_1]$$

This means that $\mathbb{E}[(b_t(\mathcal{Z}_2) - b_t(\mathcal{Z}_1))^2 | \mathcal{Z}_1] = \text{Var}(b_t(\mathcal{Z}_2) | \mathcal{Z}_1)$. These observations let us compute that

$$\begin{aligned}
\mathbb{V}_t(\mathcal{Z}_1) &= \text{Var}(\psi_t(\mathcal{Z}_1)) \\
&= \text{Var}(\psi_t(\mathcal{Z}_2)) + \text{Var}(R_t(\mathcal{Z}_1, \mathcal{Z}_2)) \\
&= \text{Var}(\psi_t(\mathcal{Z}_2)) + \mathbb{E}[\mathbb{E}[R_t(\mathcal{Z}_1, \mathcal{Z}_2)^2 | \mathcal{Z}_2]] \\
&= \mathbb{V}_t(\mathcal{Z}_2) + \mathbb{E}\left[(b_t(\mathcal{Z}_2) - b_t(\mathcal{Z}_1))^2 \mathbb{E}\left[\frac{\mathbb{1}(T=t)^2}{\pi_t(\mathcal{Z}_1)^2} - \frac{2\mathbb{1}(T=t)}{\pi_t(\mathcal{Z}_1)} + 1 \mid \mathcal{Z}_2\right]\right] \\
&= \mathbb{V}_t(\mathcal{Z}_2) + \mathbb{E}\left[\text{Var}(b_t(\mathcal{Z}_2) | \mathcal{Z}_1) \left(\frac{1}{\pi_t(\mathcal{Z}_1)} - 1\right)\right]. \tag{3.A.2}
\end{aligned}$$

This establishes the formula in the case $\Delta = \chi_t$.

For general Δ , note first that \mathbf{R} and $(\psi_t(\mathcal{Z}_2))_{t \in \mathbb{T}}$ are uncorrelated since $\mathbb{E}[\mathbf{R} | \mathcal{Z}_2] = 0$. Therefore we have

$$\begin{aligned}
\mathbb{V}_\Delta(\mathcal{Z}_1) &= \text{Var}\left(\sum_{t \in \mathbb{T}} c_t \psi_t(\mathcal{Z}_1)\right) \\
&= \text{Var}\left(\sum_{t \in \mathbb{T}} c_t \psi_t(\mathcal{Z}_2)\right) + \text{Var}\left(\sum_{t \in \mathbb{T}} c_t R_t(\mathcal{Z}_1, \mathcal{Z}_2)\right) \\
&= \mathbb{V}_\Delta(\mathcal{Z}_2) + \mathbf{c}^\top \text{Var}(\mathbf{R}) \mathbf{c}.
\end{aligned}$$

It now remains to establish the covariance expressions for the last term, since the expression of $\text{Var}(R_t(\mathcal{Z}_1, \mathcal{Z}_2))$ was found in (3.A.2). To this end, note first that for any $s, t \in \mathbb{T}$ with $s \neq t$,

$$\mathbb{E}\left[\left(\frac{\mathbb{1}(T=s)}{\pi_s(\mathcal{Z}_2)} - 1\right) \left(\frac{\mathbb{1}(T=t)}{\pi_t(\mathcal{Z}_2)} - 1\right) \mid \mathcal{Z}_2\right] = -1.$$

Thus we finally conclude that

$$\begin{aligned}
\text{Cov}(R_s, R_t) &= -\mathbb{E}[(b_s(\mathcal{Z}_2) - b_s(\mathcal{Z}_1))(b_t(\mathcal{Z}_2) - b_t(\mathcal{Z}_1))] \\
&= -\mathbb{E}[\text{Cov}(b_s(\mathcal{Z}_2), b_t(\mathcal{Z}_2) | \mathcal{Z}_1)],
\end{aligned}$$

as desired. □

3.A.6 Computations for Example 3.3.10

Using symmetry, we see that

$$\begin{aligned}
\pi_1(\mathbf{Z}) &= \mathbb{E}[T | \mathbf{Z}] = \mathbb{E}[\mathbf{W} | \mathbf{Z}] \\
&= (0.5 - \mathbf{Z})\mathbb{P}(\mathbf{W} \leq 0.5 | \mathbf{Z}) + (0.5 + \mathbf{Z})\mathbb{P}(\mathbf{W} > 0.5 | \mathbf{Z}) = \frac{1}{2}.
\end{aligned}$$

From $T \perp\!\!\!\perp \mathbf{Z}$ we observe that

$$\begin{aligned}
b_t(\mathbf{W}) &= t + \mathbb{E}[g(\mathbf{Z}) | T = t, \mathbf{W}] = t + g(\mathbf{Z}), \\
b_t(\mathbf{Z}) &= t + \mathbb{E}[g(\mathbf{Z}) | T = t, \mathbf{Z}] = t + g(\mathbf{Z}), \\
b_t(0) &= t + \mathbb{E}[g(\mathbf{Z}) | T = t] = t + \mathbb{E}[g(\mathbf{Z})].
\end{aligned}$$

Plugging these expressions into the asymptotic variance yields:

$$\begin{aligned}
 \mathbb{V}_t(0) &= \text{Var} \left(b_t(0) + \frac{\mathbb{1}(T=0)}{\pi_t(0)}(Y - b_t(0)) \right) \\
 &= \pi_t(0)^{-2} \text{Var}(\mathbb{1}(T=0)(Y - b_t(0))) \\
 &= 4\mathbb{E}[\mathbb{1}(T=0)(g(\mathbf{Z}) - \mathbb{E}[g(\mathbf{Z})] + v(\mathbf{W})\varepsilon_Y)^2] \\
 &= 4\mathbb{E}[\mathbf{W}(g(\mathbf{Z}) - \mathbb{E}[g(\mathbf{Z})])^2] + 2\mathbb{E}[v(\mathbf{W})\varepsilon_Y^2] \\
 &= 2 \text{Var}(g(\mathbf{Z})) + 2\mathbb{E}[v(\mathbf{W})^2]\mathbb{E}[\varepsilon_Y^2]
 \end{aligned}$$

$$\begin{aligned}
 \mathbb{V}_t(\mathbf{Z}) &= \text{Var} \left(b_t(\mathbf{Z}) + \frac{\mathbb{1}(T=t)}{\pi_t(\mathbf{Z})}(Y - b_t(\mathbf{Z})) \right) \\
 &= \text{Var}(g(\mathbf{Z}) + 2\mathbb{1}(T=t)v(\mathbf{W})\varepsilon_Y) \\
 &= \text{Var}(g(\mathbf{Z})) + 2\mathbb{E}[v(\mathbf{W})^2]\mathbb{E}[\varepsilon_Y^2]
 \end{aligned}$$

$$\begin{aligned}
 \mathbb{V}_t(\mathbf{W}) &= \text{Var} \left(b_t(\mathbf{W}) + \frac{\mathbb{1}(T=t)}{\pi_t(\mathbf{W})}(Y - b_t(\mathbf{W})) \right) \\
 &= \text{Var} \left(g(\mathbf{Z}) + \frac{\mathbb{1}(T=t)v(\mathbf{W})\varepsilon_Y}{\mathbf{W}} \right) \\
 &= \text{Var}(g(\mathbf{Z})) + \mathbb{E} \left[\frac{v(\mathbf{W})^2}{\mathbf{W}} \right] \mathbb{E}[\varepsilon_Y^2].
 \end{aligned}$$

□

Proof of Theorem 3.4.3

We first prove that $R_i \xrightarrow{P} 0$ for $i = 1, 2, 3$. The third term is taken care of by combining Assumption 3.4.2 (i, iv) with Cauchy-Schwarz:

$$\begin{aligned}
 |R_3| &\leq \sqrt{|\mathcal{I}_3|} \sqrt{\frac{1}{|\mathcal{I}_3|} \sum_{i \in \mathcal{I}_3} (\hat{m}_{\hat{\theta}}(t | \mathbf{W}_i)^{-1} - m_{\hat{\theta}}(t | \mathbf{W}_i)^{-1})^2} \sqrt{\mathcal{E}_{2,t}^{(n)}} \\
 &\leq \ell_c \sqrt{|\mathcal{I}_3| \mathcal{E}_{1,t}^{(n)} \mathcal{E}_{2,t}^{(n)}} \xrightarrow{P} 0,
 \end{aligned}$$

where $\ell_c > 0$ is the Lipschitz constant of $x \mapsto x^{-1}$ on $[c, 1 - c]$.

For the first term, R_1 , note that conditionally on $\hat{\theta}$ and the observed estimated representations $\mathbf{Z}_{\bullet} := (\mathbf{Z}_{\hat{\theta},i})_{i \in \mathcal{I}_3} = (\varphi(\hat{\theta}, \mathbf{W}_i))_{i \in \mathcal{I}_3}$, the summands are conditionally i.i.d. due to Assumption 3.4.1. They are also conditionally mean zero since

$$\begin{aligned}
 \mathbb{E}[r_i^1 | \mathbf{Z}_{\bullet}, \hat{\theta}] &= (\hat{g}_{\hat{\theta}}(t, \mathbf{W}_i) - g_{\hat{\theta}}(t, \mathbf{W}_i)) \left(1 - \frac{\mathbb{E}[\mathbb{1}(T_i = t) | \mathbf{Z}_{\bullet}, \hat{\theta}]}{m_{\hat{\theta}}(t | \mathbf{W}_i)} \right) \\
 &= (\hat{g}_{\hat{\theta}}(t, \mathbf{W}_i) - g_{\hat{\theta}}(t, \mathbf{W}_i)) \left(1 - \frac{\mathbb{E}[\mathbb{1}(T_i = t) | \mathbf{Z}_{\hat{\theta},i}, \hat{\theta}]}{m_{\hat{\theta}}(t | \mathbf{W}_i)} \right) = 0,
 \end{aligned}$$

for each $i \in \mathcal{I}_3$. Using Assumption 3.4.2 (i), we can bound the conditional variance by

$$\begin{aligned} \text{Var}\left(r_i^1 \mid \mathbf{Z}_\bullet, \hat{\theta}\right) &= \frac{(\hat{g}_\theta(t, \mathbf{W}_i) - g_\theta(t, \mathbf{W}_i))^2}{m_\theta(t \mid \mathbf{W})^2} \text{Var}\left(\mathbb{1}(T = t) \mid \mathbf{Z}_\bullet, \hat{\theta}\right) \\ &= \frac{(\hat{g}_\theta(t, \mathbf{W}_i) - g_\theta(t, \mathbf{W}_i))^2}{m_\theta(t \mid \mathbf{W})} (1 - m_\theta(t \mid \mathbf{W})) \\ &\leq \frac{(1-c)}{c} (\hat{g}_\theta(t, \mathbf{W}_i) - g_\theta(t, \mathbf{W}_i))^2 \end{aligned}$$

Now it follows from Assumption 3.4.2 (v) and the conditional Chebyshev's inequality that

$$\mathbb{P}\left(|R_1| > \epsilon \mid \mathbf{Z}_\bullet, \hat{\theta}\right) \leq \epsilon^{-2} \text{Var}(R_1 \mid \mathbf{Z}_\bullet, \hat{\theta}) \leq \frac{1-c}{\epsilon^2 c} \mathcal{E}_{2,t}^{(n)} \xrightarrow{P} 0,$$

from which we conclude that $R_1 \xrightarrow{P} 0$.

The analysis of R_2 is similar. Observe that the summands are conditionally i.i.d. given $\mathbf{Z}_\bullet, \hat{\theta}$, and $T_\bullet := (T_i)_{i \in \mathcal{I}_3}$. They are also conditionally mean zero since

$$\begin{aligned} \mathbb{E}\left[r_i^2 \mid \mathbf{Z}_\bullet, T_\bullet, \hat{\theta}\right] &= \left(\mathbb{E}[Y_i \mid \mathbf{Z}_\bullet, T_\bullet, \hat{\theta}] - g_\theta(t, \mathbf{W}_i)\right) \left(\frac{\mathbb{1}(T_i = t)}{\hat{m}_\theta(\mathbf{W}_i)} - \frac{\mathbb{1}(T_i = t)}{m_\theta(\mathbf{W}_i)}\right) \\ &= \left(\mathbb{E}[Y_i \mid \varphi(\hat{\theta}, \mathbf{W}_i), T_i, \hat{\theta}] - g_\theta(t, \mathbf{W}_i)\right) \left(\frac{\mathbb{1}(T_i = t)}{\hat{m}_\theta(\mathbf{W}_i)} - \frac{\mathbb{1}(T_i = t)}{m_\theta(\mathbf{W}_i)}\right) = 0, \end{aligned}$$

where we use that on the event $(T_i = t)$,

$$\mathbb{E}[Y_i \mid \mathbf{Z}_{\hat{\theta},i}, T_i, \hat{\theta}] = \mathbb{E}[Y_i \mid \mathbf{Z}_{\hat{\theta},i}, T_i = t, \hat{\theta}] = g_\theta(t, \mathbf{W}_i).$$

To bound the conditional variance, note first that Assumption 3.4.1 and Assumption 3.4.2 imply

$$\text{Var}(Y_i \mid \mathbf{Z}_{\hat{\theta},i}, T_i, \hat{\theta}) \leq \mathbb{E}[Y_i^2 \mid \mathbf{Z}_{\hat{\theta},i}, T_i, \hat{\theta}] = \mathbb{E}[\mathbb{E}[Y_i^2 \mid \mathbf{W}_i, T_i, \hat{\theta}] \mid \mathbf{Z}_{\hat{\theta},i}, T_i, \hat{\theta}] \leq C,$$

which in return implies that

$$\begin{aligned} \text{Var}\left(r_i^2 \mid \mathbf{Z}_\bullet, T_\bullet, \hat{\theta}\right) &= \mathbb{1}(T_i = t) \left(\frac{1}{\hat{m}_\theta(\mathbf{W}_i)} - \frac{1}{m_\theta(\mathbf{W}_i)}\right)^2 \text{Var}\left(Y_i \mid \mathbf{Z}_{\hat{\theta},i}, T_i, \hat{\theta}\right) \\ &\leq C \left(\frac{1}{\hat{m}_\theta(\mathbf{W}_i)} - \frac{1}{m_\theta(\mathbf{W}_i)}\right)^2. \end{aligned}$$

It now follows from the conditional Chebyshev's inequality that

$$\begin{aligned} \mathbb{P}(|R_2| > \epsilon \mid \mathbf{Z}_\bullet, T_\bullet, \hat{\theta}) &\leq \epsilon^{-2} \text{Var}(R_2 \mid \mathbf{Z}_\bullet, T_\bullet, \hat{\theta}) \\ &\leq \frac{C}{|\mathcal{I}_3|} \sum_{i \in \mathcal{I}_3} \left(\frac{1}{\hat{m}_\theta(\mathbf{W}_i)} - \frac{1}{m_\theta(\mathbf{W}_i)}\right)^2 \leq C \ell_c^2 \mathcal{E}_{1,t}^{(n)}. \end{aligned}$$

3 Efficient adjustment for complex covariates: Gaining efficiency with DOPE

Assumptions 3.4.2 (iv) implies convergence to zero in probability, and from this we conclude that also $R_2 \xrightarrow{P} 0$.

We now turn to the discussion of the oracle term $U_{\hat{\theta}}^{(n)}$. From Assumption 3.4.1 it follows that $(T_i, \mathbf{W}_i, Y_i)_{i \in \mathcal{I}_3} \perp\!\!\!\perp \hat{\theta}$, and hence the conditional distribution $U_{\hat{\theta}}^{(n)} | \hat{\theta} = \theta$ is the same as $U_{\theta}^{(n)}$. We show that this distribution is asymptotically Gaussian uniformly over $\theta \in \Theta$. To see this, note that for each $\theta \in \Theta$, the terms of $U_{\theta}^{(n)}$ are i.i.d. with mean $\chi_t(\mathbf{Z}_{\theta})$ and variance $\mathbb{V}_t(\mathbf{Z}_{\theta})$. The conditional Jensen's inequality and Assumption 3.4.2 (i) imply that

$$\begin{aligned} \mathbb{E}[|u_i(\theta)|^{2+\delta}] &\leq C_{2+\delta}(\mathbb{E}[|g_{\theta}(t, \mathbf{W})|^{2+\delta}] + c^{-(1+\delta)}\mathbb{E}[|Y - g_{\theta}(t, \mathbf{W})|^{2+\delta}]) \\ &\leq C'\mathbb{E}[|Y|^{2+\delta}], \end{aligned}$$

where $C' > 0$ is a constant depending on c and δ . Thus we can apply the Lindeberg-Feller CLT, as stated in Shah and Peters [2020, Lem. 18], to conclude that

$$S_n(\theta) := \sqrt{|\mathcal{I}_3|}\mathbb{V}_t(\mathbf{Z}_{\theta})^{-1/2}(U_{\theta}^{(n)} - \chi_t(\mathbf{Z}_{\theta})) \xrightarrow{d/\Theta} \mathbf{N}(0, 1),$$

where the convergence holds uniformly over $\theta \in \Theta$. With $\Phi(\cdot)$ denoting the CDF of a standard normal, this implies that

$$\begin{aligned} \sup_{t \in \mathbb{R}} |\mathbb{P}(S_n(\hat{\theta}) \leq t) - \Phi(t)| &= \sup_{t \in \mathbb{R}} \left| \int (\mathbb{P}(S_n(\hat{\theta}) \leq t | \hat{\theta} = \theta) - \Phi(t)) P_{\hat{\theta}}(d\theta) \right| \\ &\leq \sup_{\theta \in \Theta} \sup_{t \in \mathbb{R}} |\mathbb{P}(S_n(\theta) \leq t) - \Phi(t)| \longrightarrow 0. \end{aligned} \quad (3.A.3)$$

This shows that $S_n(\hat{\theta}) \xrightarrow{d} \mathbf{N}(0, 1)$ as desired. The last part of the theorem is a simple consequence of Slutsky's theorem. \square

Proof of Theorem 3.4.4

For each $i \in \mathcal{I}_3$, let $\hat{u}_i = u_i(\hat{\theta}) + r_i^1 + r_i^2 + r_i^3$ be the decomposition from (3.4.16). For the squared sum, it is immediate from Theorem 3.4.3 that

$$\left(\frac{1}{|\mathcal{I}_3|} \sum_{i \in \mathcal{I}_3} \hat{u}_i \right)^2 = \left(\frac{1}{|\mathcal{I}_3|} \sum_{i \in \mathcal{I}_3} u_i(\hat{\theta}) \right)^2 + o_P(1).$$

For the sum of squares, we first expand the squares as

$$\hat{u}_i^2 = u_i(\hat{\theta})^2 + 2(r_i^1 + r_i^2 + r_i^3)u_i(\hat{\theta}) + (r_i^1 + r_i^2 + r_i^3)^2. \quad (3.A.4)$$

We show that the last two terms are convergent to zero in probability. For the cross-term, we note by Cauchy-Schwarz that

$$\left| \frac{1}{|\mathcal{I}_3|} \sum_{i \in \mathcal{I}_3} (r_i^1 + r_i^2 + r_i^3)u_i(\hat{\theta}) \right| \leq \left(\frac{1}{|\mathcal{I}_3|} \sum_{i \in \mathcal{I}_3} (r_i^1 + r_i^2 + r_i^3)^2 \right)^{1/2} \left(\frac{1}{|\mathcal{I}_3|} \sum_{i \in \mathcal{I}_3} u_i(\hat{\theta})^2 \right)^{1/2}.$$

We show later that the sum $\frac{1}{|\mathcal{I}_3|} \sum_{i \in \mathcal{I}_3} u_i^2$ is convergent in probability. Thus, to show that the last two terms of (3.A.4) converge to zero, it suffices to show that $\frac{1}{|\mathcal{I}_3|} \sum_{i \in \mathcal{I}_3} (r_i^1 + r_i^2 + r_i^3)^2 \xrightarrow{P} 0$. To this end, we observe that

$$\frac{1}{|\mathcal{I}_3|} \sum_{i \in \mathcal{I}_3} (r_i^1 + r_i^2 + r_i^3)^2 \leq \frac{3}{|\mathcal{I}_3|} \sum_{i \in \mathcal{I}_3} (r_i^1)^2 + (r_i^2)^2 + (r_i^3)^2$$

The last term is handled similarly to R_3 in the proof of Theorem 3.4.3,

$$\begin{aligned} \frac{1}{|\mathcal{I}_3|} \sum_{i \in \mathcal{I}_3} (r_i^3)^2 &= \frac{1}{|\mathcal{I}_3|} \sum_{i \in \mathcal{I}_3} (\hat{g}_\theta(t, \mathbf{W}_i) - g_\theta(t, \mathbf{W}_i))^2 \left(\frac{\mathbb{1}(T_i = t)}{\hat{m}_\theta(t | \mathbf{W}_i)} - \frac{\mathbb{1}(T_i = t)}{m_\theta(t | \mathbf{W}_i)} \right)^2 \\ &\leq \frac{1}{|\mathcal{I}_3|} \left(\sum_{i \in \mathcal{I}_3} (\hat{g}_\theta(t, \mathbf{W}_i) - g_\theta(t, \mathbf{W}_i))^2 \right) \left(\sum_{i \in \mathcal{I}_3} \left(\frac{\mathbb{1}(T_i = t)}{\hat{m}_\theta(t | \mathbf{W}_i)} - \frac{\mathbb{1}(T_i = t)}{m_\theta(t | \mathbf{W}_i)} \right)^2 \right) \\ &\leq \ell_c^2 n \mathcal{E}_1^{(n)} \mathcal{E}_{2,t}^{(n)} \xrightarrow{P} 0, \end{aligned}$$

where we have used the naive inequality $\sum a_i^2 b_i^2 \leq (\sum a_i^2)(\sum b_i^2)$ rather than Cauchy-Schwarz. From the proof of Theorem 3.4.3, we have that

$$\mathbb{E}[(r_i^1)^2 | \mathbf{Z}_\bullet, \hat{\theta}] = \text{Var}[r_i^1 | \mathbf{Z}_\bullet, \hat{\theta}] \leq \frac{1-c}{c} (\hat{g}_\theta(t, \mathbf{W}_i) - g_\theta(t, \mathbf{W}_i))^2,$$

and hence the conditional Markov's inequality yields

$$\begin{aligned} \mathbb{P} \left(\frac{1}{|\mathcal{I}_3|} \sum_{i \in \mathcal{I}_3} (r_i^1)^2 > \epsilon | \mathbf{Z}_\bullet, \hat{\theta} \right) &\leq \epsilon^{-1} \mathbb{E} \left[\frac{1}{|\mathcal{I}_3|} \sum_{i \in \mathcal{I}_3} (r_i^1)^2 | \mathbf{Z}_\bullet, \hat{\theta} \right] \\ &= \epsilon^{-1} \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left[(r_i^1)^2 | \mathbf{Z}_\bullet, \hat{\theta} \right] \\ &\leq \frac{1-c}{\epsilon c} \mathcal{E}_{2,t}^{(n)} \xrightarrow{P} 0. \end{aligned}$$

Thus we also conclude that $\frac{1}{|\mathcal{I}_3|} \sum_{i \in \mathcal{I}_3} (r_i^1)^2 \xrightarrow{P} 0$. Analogously, the final remainder $\frac{1}{|\mathcal{I}_3|} \sum_{i \in \mathcal{I}_3} (r_i^2)^2$ can be shown to converge to zero in probability by leveraging the argument for $R_2 \xrightarrow{P} 0$ in the proof of Theorem 3.4.3.

Combining the arguments above we conclude that

$$\hat{\mathcal{V}}_t(\hat{g}, \hat{m}) = \frac{1}{|\mathcal{I}_3|} \sum_{i \in \mathcal{I}_3} \hat{u}_i^2 - \left(\frac{1}{|\mathcal{I}_3|} \sum_{i \in \mathcal{I}_3} \hat{u}_i \right)^2 = \frac{1}{|\mathcal{I}_3|} \sum_{i \in \mathcal{I}_3} u_i(\hat{\theta})^2 - \left(\frac{1}{|\mathcal{I}_3|} \sum_{i \in \mathcal{I}_3} u_i(\hat{\theta}) \right)^2 + o_P(1)$$

As noted in the proof of Theorem 3.4.3, for each $\theta \in \Theta$ the terms $\{u_i(\theta)\}_{i \in \mathcal{I}_3}$ are i.i.d with $\mathbb{E}[|u_i(\theta)|^{2+\delta}] \leq C' \mathbb{E}[|Y|^{2+\delta}] < \infty$. Hence, the uniform law of large numbers, as stated in Shah and Peters [2020, Lem. 19], implies that

$$\frac{1}{|\mathcal{I}_3|} \sum_{i \in \mathcal{I}_3} u_i(\theta) - \mathbb{E}[u_1(\theta)] \xrightarrow{P/\Theta} 0, \quad \frac{1}{|\mathcal{I}_3|} \sum_{i \in \mathcal{I}_3} u_i^2(\theta) - \mathbb{E}[u_1(\theta)^2] \xrightarrow{P/\Theta} 0,$$

3 Efficient adjustment for complex covariates: Gaining efficiency with DOPE

where the convergence in probability holds uniformly over $\theta \in \Theta$. Since $\text{Var}(u_1(\theta)) = \text{Var}(\psi_t(\mathbf{Z}_\theta)) = \mathbb{V}_t(\mathbf{Z}_\theta)$, this lets us conclude that

$$S_n(\theta) := \frac{1}{|\mathcal{I}_3|} \sum_{i \in \mathcal{I}_3} u_i(\theta)^2 - \left(\frac{1}{|\mathcal{I}_3|} \sum_{i \in \mathcal{I}_3} u_i(\theta) \right)^2 - \mathbb{V}_t(\mathbf{Z}_\theta) \xrightarrow{P/\Theta} 0.$$

We now use that convergence in distribution is equivalent to convergence in probability for deterministic limit variables, see Christgau et al. [2023b, Cor. B.4.] for a general uniform version of the statement. Since Assumption 3.4.1 implies that the conditional distribution $S_n(\hat{\theta}) | \hat{\theta} = \theta$ is the same as $S_n(\theta)$, the computation in (3.A.3) with $\Phi(\cdot) = \mathbb{1}(\cdot \geq 0)$ yields that $S_n(\hat{\theta}) \xrightarrow{d} 0$. This lets us conclude that

$$\hat{\mathcal{V}}_t(\hat{g}, \hat{m}) = S_n(\hat{\theta}) + o_P(1) \xrightarrow{P} 0,$$

which finishes the proof. \square

3.A.7 Proof of Theorem 3.4.6

We fix $P \in \mathcal{P}$ and suppress it from notation throughout the proof. Given $\theta \in \mathbb{R}^d$, the single-index model assumption (3.4.13) implies that

$$b_t(\mathbf{W}^\top \theta) = \mathbb{E}[Y | T = t, \mathbf{W}^\top \theta] = \mathbb{E}[h_t(\mathbf{W}^\top \theta_P) | T = t, \mathbf{W}^\top \theta].$$

Now since $h_t \in C^1(\mathbb{R})$, we may write

$$\begin{aligned} h_t(\mathbf{W}^\top \theta) &= h_t(\mathbf{W}^\top \theta_P) + R(\theta, \theta_P), \\ R(\theta, \theta_P) &= h'_t(\mathbf{W}^\top \theta_P) \mathbf{W}^\top (\theta - \theta_P) + r(\mathbf{W}^\top (\theta - \theta_P)) \mathbf{W}^\top (\theta - \theta_P) \end{aligned}$$

for a continuous function $r: \mathbb{R} \rightarrow \mathbb{R}$ satisfying that $\lim_{\epsilon \rightarrow 0} r(\epsilon) = 0$. It follows that

$$\mathbb{E}[h_t(\mathbf{W}^\top \theta_P) | T = t, \mathbf{W}^\top \theta] = h_t(\mathbf{W}^\top \theta) - \mathbb{E}[R(\theta, \theta_P) | T = t, \mathbf{W}^\top \theta],$$

The theoretical bias induced by adjusting for $\mathbf{W}^\top \theta$ instead of \mathbf{W} , or equivalently $\mathbf{W}^\top \theta_P$, is therefore

$$\begin{aligned} \chi_t(\mathbf{W}^\top \theta; P) - \chi_t(\mathbf{W}) &= \chi_t(\mathbf{W}^\top \theta; P) - \chi_t(\mathbf{W}^\top \theta_P) \\ &= \mathbb{E} [h_t(\mathbf{W}^\top \theta) - \mathbb{E}[R(\theta, \theta_P) | T = t, \mathbf{W}^\top \theta] - h_t(\mathbf{W}^\top \theta_P)] \\ &= \mathbb{E}[R(\theta, \theta_P) - \mathbb{E}[R(\theta, \theta_P) | T = t, \mathbf{W}^\top \theta]] \\ &= \mathbb{E} \left[R(\theta, \theta_P) \cdot \left(1 - \frac{\mathbb{1}(T = t)}{\mathbb{P}(T = t | \mathbf{W}^\top \theta)} \right) \right] \\ &= C(\theta, \theta_P)^\top (\theta - \theta_P), \end{aligned}$$

where

$$C(\theta; \theta_P) = \mathbb{E} \left[(h'_t(\mathbf{W}^\top \theta) + r(\mathbf{W}^\top (\theta - \theta_P))) \left(1 - \frac{\mathbb{1}(T = t)}{\mathbb{P}(T = t | \mathbf{W}^\top \theta)} \right) \mathbf{W} \right].$$

To show that u is differentiable at $\theta = \theta_P$, it therefore suffices to show that the mapping $\theta \mapsto C(\theta; \theta_P)$ is continuous at θ_P .

We first show that $\mathbb{P}(T = t | \mathbf{W}^\top \theta)$ is continuous at θ_P almost surely, which follows after applying a coordinate change such that θ becomes a basis vector. To be more precise, we may choose a neighborhood $\mathcal{U} \subset \mathbb{R}^d$ of θ_P and $d - 1$ continuous functions

$$q_2, \dots, q_d: \mathcal{U} \longrightarrow \mathcal{S}^{d-1}$$

such that $Q(\theta) := (\|\theta\|^{-1}\theta, q_2(\theta), \dots, q_d(\theta))$ is an orthogonal matrix for every $\theta \in \mathcal{U}$. For example, the first d vectors of the Gram-Schmidt process applied to $(\theta, \mathbf{e}_1, \dots, \mathbf{e}_d)$ is continuous and yields an orthonormal basis for $\theta \in \mathbb{R}^d \setminus \text{span}(\mathbf{e}_d)$, so this works in case $\theta_P \neq \pm \mathbf{e}_d$. Let $Z(\theta) = Q(\theta)^\top \mathbf{W}$ and note that $Z(\theta)_1 = \|\theta\|^{-1} \theta^\top \mathbf{W} = \|\theta\|^{-1} \mathbf{W}^\top \theta$. Then by iterated expectations,

$$\begin{aligned} \mathbb{P}(T = t | \mathbf{W}^\top \theta) &= \mathbb{E}[m(t | \mathbf{W}) | \mathbf{W}^\top \theta] \\ &= \mathbb{E}[m(t | Q(\theta)Z(\theta)) | Z(\theta)_1] \\ &= \frac{\int_{\mathbb{R}^{d-1}} m(t | \widetilde{\mathbf{W}}(\theta, z)) p_{\mathbf{W}}(\widetilde{\mathbf{W}}(\theta, z)) dz}{\int_{\mathbb{R}^{d-1}} p_{\mathbf{W}}(\widetilde{\mathbf{W}}(\theta, z)) dz}, \quad \widetilde{\mathbf{W}}(\theta, z) := Q(\theta) \begin{pmatrix} Z(\theta)_1 \\ z \end{pmatrix} \end{aligned}$$

Each integrand is bounded and continuous over $\theta \in \mathcal{U}$. From dominated convergence it follows that $\mathbb{P}(T = t | \mathbf{W}^\top \theta)$ is continuous at θ_P almost surely. It follows that almost surely,

$$\begin{aligned} &(h'_t(\mathbf{W}^\top \theta) + r(\mathbf{W}^\top(\theta - \theta_P))) \left(1 - \frac{\mathbb{1}(T = t)}{\mathbb{P}(T = t | \mathbf{W}^\top \theta)}\right) \mathbf{W} \\ &\longrightarrow h'_t(\mathbf{W}^\top \theta_P) \left(1 - \frac{\mathbb{1}(T = t)}{\mathbb{P}(T = t | \mathbf{W}^\top \theta_P)}\right) \mathbf{W}, \quad \text{for } \theta \rightarrow \theta_P. \end{aligned}$$

By dominated convergence again we conclude that

$$C(\theta; \theta_P) \longrightarrow \mathbb{E} \left[h'_t(\mathbf{W}^\top \theta_P) \left(1 - \frac{\mathbb{1}(T = t)}{\mathbb{P}(T = t | \mathbf{W}^\top \theta_P)}\right) \mathbf{W} \right], \quad \theta \rightarrow \theta_P.$$

This shows that $\theta \mapsto C(\theta; \theta_P)$ is continuous at θ_P , and hence we conclude that u is differentiable at $\theta = \theta_P$ with gradient

$$\begin{aligned} \nabla u(\theta_P) &= \mathbb{E} \left[h'_t(\mathbf{W}^\top \theta_P) \left(1 - \frac{\mathbb{1}(T = t)}{\mathbb{P}(T = t | \mathbf{W}^\top \theta_P)}\right) \mathbf{W} \right] \\ &= \mathbb{E} \left[h'_t(\mathbf{W}^\top \theta_P) \left(1 - \frac{\mathbb{P}(T = t | \mathbf{W})}{\mathbb{P}(T = t | \mathbf{W}^\top \theta_P)}\right) \mathbf{W} \right]. \end{aligned}$$

□




Figure 3.B.1: Neural network architecture for single-index model.

3.B. Details of simulation study

Our experiments were conducted in Python [Van Rossum et al., 1995]. The linear and logistic regression was imported from the `scikit-learn` package [Pedregosa et al., 2011], and the neural network for the DOPE-IDX was implemented using `pytorch` [Paszke et al., 2019].

The neural network architecture is illustrated in Figure 3.B.1. The network was optimized using MSE loss and the ADAM optimizer with `lr=1e-3` and `n_iter=1200` in the simulation experiment. For the NHANES application the settings were similar, but with BCELoss and `n_iter=3000`. We discovered that the optimization procedure could get stuck in a local minimum with significantly lower training score. To avoid getting stuck in a such a ‘bad local minimum’, it was possible to refit the network with random initialization several times and pick the model with highest training score. The frequency of bad local minimums is roughly 15%, with variation between each setting. Thus, refitting the network, say 5 times, is enough to avoid a bad minimum $> 99, 99\%$ of the time. To save the computational cost of refitting many times for each simulation, we initialized the first layer with an initial value $\theta_{init} = \beta + \text{Unif}(-0.1, 0.1)^{\otimes d}$. The logistic regression for the propensity score was fitted without ℓ_2 -penalty and optimized using the `lbfgs` optimizer. The propensity score was clipped to the interval $(0.01, 0.99)$ for all estimators of the adjusted mean.

Figure 3.B.2 shows the cross-fitted DOPE estimators versus their full sample counterparts in the simulation setup of Section 3.5.1. We observe that cross-fitting generally seems to decrease performance for small sample sizes, but the discrepancy diminishes for larger sample sizes.

Table 3.B.1 corresponds to Table 3.5.1 in the main manuscript, but where covariates with more than 50% missing data have been removed rather than imputed. Except for the naive contrast (which has dropped two rows) the ordering according to bootstrap variance is the same.

Figure 3.B.3 shows the same coverage values as Figure 3.5.5, but also includes the

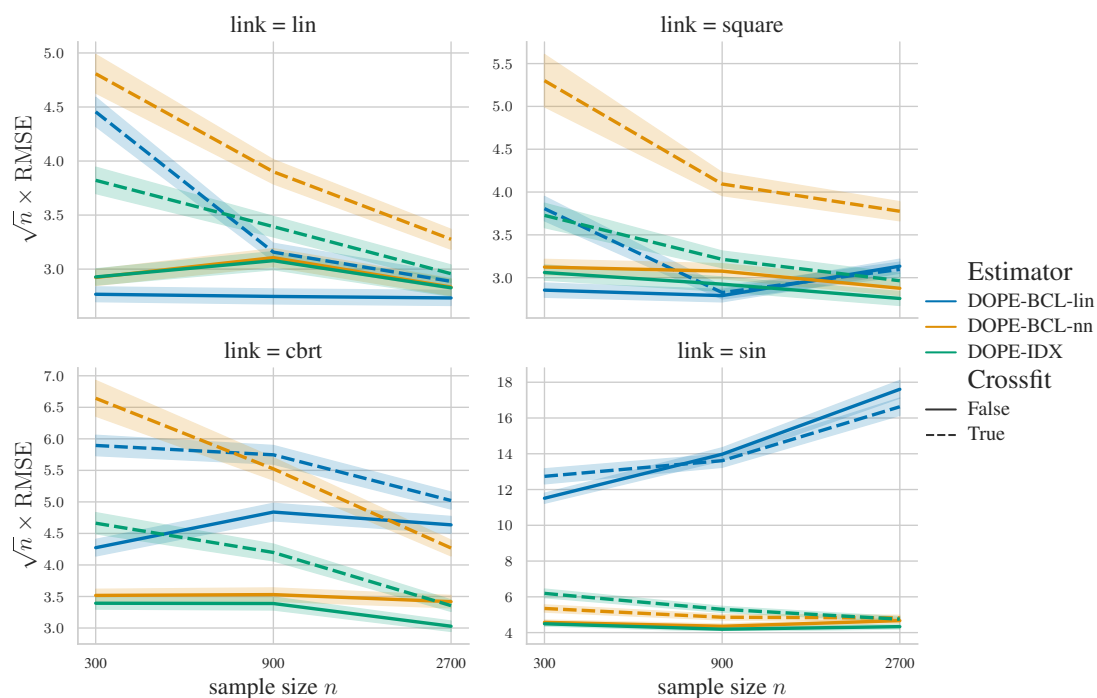


Figure 3.B.2: Comparison of DOPE estimators with and without sample splitting.

Estimator	Estimate	BS se	BS CI
Regr. (NN)	0.020	0.008	(0.002, 0.035)
Regr. (Logistic)	0.027	0.009	(0.012, 0.048)
DOPE-BCL (Logistic)	0.024	0.010	(0.004, 0.040)
Naive contrast	0.388	0.010	(0.369, 0.407)
DOPE-BCL (NN)	0.018	0.010	(-0.003, 0.036)
DOPE-IDX (NN)	0.023	0.012	(-0.001, 0.047)
AIPW (NN)	0.017	0.016	(-0.017, 0.046)
AIPW (Logistic)	0.022	0.016	(-0.012, 0.051)
IPW (Logistic)	-0.047	0.027	(-0.119, -0.01)

Table 3.B.1: Treatment effect estimates for NHANES dataset, where covariates with more than 50% missing values have been removed.

average lengths of the confidence intervals in the parentheses. We note that the lengths of the confidence intervals for the DOPE estimators are much smaller than those of the AIPW estimator.

3 Efficient adjustment for complex covariates: Gaining efficiency with DOPE

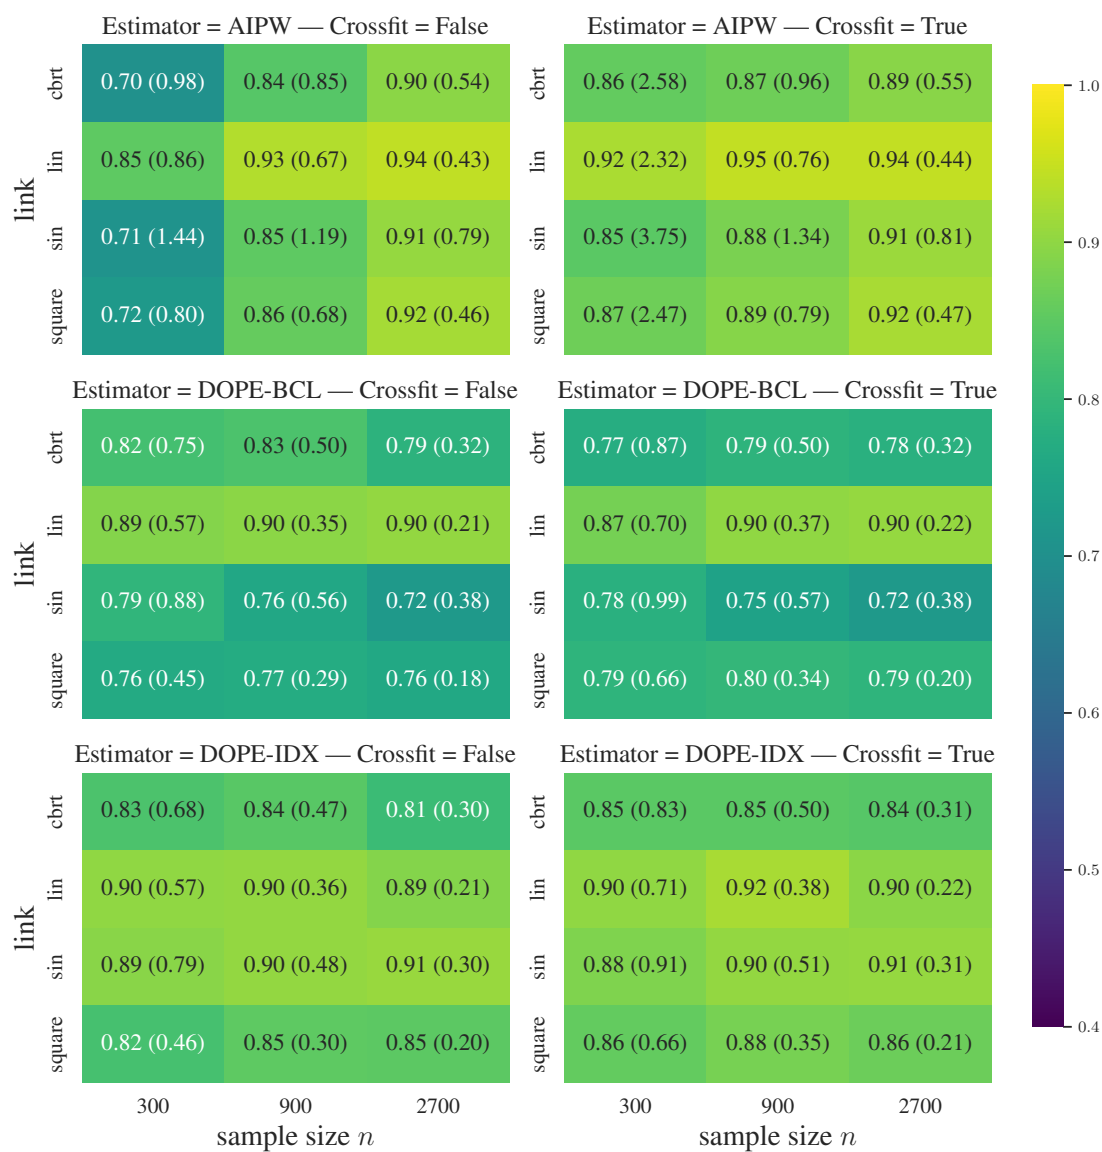


Figure 3.B.3: Coverage of asymptotic confidence intervals of the adjusted mean χ_1 , aggregated over $d \in \{4, 12, 36\}$ and with the true β given in (3.5.19). The numbers in the parentheses denote the corresponding average width of the confidence interval.

3.C. Extension to cross-fitting

Algorithm 4: Cross-fitted DOPE

1 input: observations $(T_i, \mathbf{W}_i, Y_i)_{i \in [n]}$, partition $[n] = J_1 \cup \dots \cup J_K$;
2 options: integer $1 \leq m \leq K - 2$ and options for Algorithm 3;
3 for $k = 1, \dots, K$ **do**
4 Set $\mathcal{I}_3 = J_k$, $\mathcal{I}_1 = \bigcup_{l=k+1}^{k+m} J_l$ and $\mathcal{I}_2 = [n] \setminus (\mathcal{I}_1 \cup \mathcal{I}_3)$;
5 Compute $\hat{\chi}_{t,k}^{\text{dope}}$ as the output of Algorithm 3;
6 Compute variance estimate $\hat{\mathcal{V}}_{t,k}$ given as in (3.4.17);
7 return $\hat{\chi}_t^x := \frac{1}{K} \sum_{k=1}^K \hat{\chi}_{t,k}^{\text{dope}}$ and $\hat{\mathcal{V}}_t^x := \frac{1}{K} \sum_{k=1}^K \hat{\mathcal{V}}_{t,k}$

A cross-fitting procedure for the DOPE is described in Algorithm 4, which computes both a cross-fitted version of the DOPE and its variance estimator. Here the indices of the folds are understood to cycle modulo K such that $J_{K+1} = J_1$ and so forth. This version of cross-fitting with three index sets has also been referred to as ‘double cross-fitting’ by Zivich and Breskin [2021].

We note that the ‘standard arguments’ for establishing convergence of the cross-fitted estimator cannot be applied directly to our case. This is because, for each fold $k \in [K]$, the corresponding oracle terms $U_{\hat{\theta}_k}^{(n)}$ do not only involve the data indexed by \mathcal{I}_3^k , but also depend on $\hat{\theta}_k$ which is estimated from data indexed by \mathcal{I}_1^k . Hence the oracle terms are not independent. However, we believe that this dependency should be negligible, and perhaps the convergence can be established under a more refined theoretical analysis.

4 Assumption-lean Aalen regression

ALEXANDER MANGULAD CHRISTGAU AND NIELS RICHARD HANSEN

Abstract

We propose a novel and model-free estimand to measure the conditional association between a time-to-event and an exposure given covariates. The estimand may be interpreted as a weighted hazard difference, and can be viewed as an assumption-lean generalization of the cumulative exposure coefficient in the Aalen additive hazards model. We develop an estimation method based on the principles of double machine learning, which is algorithm-agnostic and incorporates cross-fitting. We prove that our method is doubly robust in the sense that, when the nuisance functions are learned with modest rate conditions, the resulting estimator converges to the true estimand at a \sqrt{n} -rate. A simulation study is conducted, showcasing that our estimation procedure yields an efficient and robust estimator of our estimand.

4.1 Introduction

Statistical inference for the conditional association between a time-to-event and an exposure is the primary objective in many bio-medical applications. A critical task for such applications is to first define an appropriate measure of association. Formulating a general and interpretable measure is intricate and remains an active research topic.

The hazard ratio is one of the most popular and reported measures of association in clinical research, with its application being a subject of ongoing discussion since the influential work of Cox [1972]. Despite its popularity, the hazard ratio, and the often accompanying Cox model, are subject to several significant drawbacks, as highlighted by Hernán [2010], Martinussen and Vansteelandt [2013], Stensrud and Hernán [2020], among others. These drawbacks include non-collapsibility of the hazard ratio and that it is unsuitable for quantifying causal effects. The pitfalls of model misspecification is another notable drawback. When used as a conditional association measure within the framework of the Cox model, its interpretation is problematic without strong model assumptions such as proportionality.

To address the issue of model misspecification, Vansteelandt et al. [2022] propose a more general measure of association that reduces to the log (cumulative) hazard ratio under the Cox model. Generalizing a model-based measure of association with a model-free measure has now become known as assumption-lean inference [Berk et al., 2019].

This approach subsumes model-based and estimand-based methods, and can be advantageous when each method has its own merits, see Vansteelandt and Dukes [2022] and the surrounding discussion for more details.

In this work we develop an analogous approach, but based on the hazard difference as an effect measure. The hazard difference is collapsible for the *additive hazards model*, introduced by Aalen [Aalen, 1980, 1989]. As mentioned, this desirable property does not hold for the hazard ratio and the Cox model [Martinussen and Vansteelandt, 2013, Daniel et al., 2021]. There are other appealing properties of the hazard difference, but the discussion is subtle, see, for example, Didelez and Stensrud [2022], Martinussen et al. [2020]. Regardless, there are several applications where additive hazards are more plausible than proportional ones; we refer to Breslow and Day [1987], Lin and Ying [1994], Kravdal [1997], Lee S. McDaniel and Chappell [2019], Bischofberger et al. [2023].

Assumptions on the non-baseline term	Ratio-based target	Difference-based target
Parametric	Cox model [Cox, 1972]	Additive risk model in Lin and Ying [1994]
Semi-parametric	Partial Cox model [Sasieni, 1992, Huang, 1999, Zhong et al., 2022]	Partially additive hazards [McKeague and Sasieni, 1994, Dukes et al., 2019, Hou et al., 2023]
Assumption-lean	[Vansteelandt et al., 2022]	This work

Table 4.1.1: Overview of selected works on inference for the conditional association between a time-to-event and an exposure, defined using varying levels of model assumptions. All models include a nonparametric baseline term that does not depend on covariates.

Despite some of their well-behaved theoretical properties, additive hazards are “somewhat overlooked in practice” according to Martinussen and Scheike [2006, p. 103]. A common criticism of additive hazard models is that they can lead to hazard estimates that are negative and are thus not valid hazards. To circumvent this issue, we define a model-free target parameter based on the *Local Covariance Measure (LCM)* introduced by Christgau et al. [2023b]. We show that, within a partially additive hazards model, it coincides with the cumulative effect of the exposure. However, our target does not rely on such modeling assumptions and we discuss its interpretations in more general settings. Table 4.1.1 gives an overview of our target parameter in relation to some existing works. We show how the target may be estimated using flexible machine learning methods, and prove asymptotic normality under suitable rate conditions on the nuisance parameters.

4.1.1 A hazard difference estimand

To motivate our general framework and results we first present a simplified example.

Example 4.1.1. Let $T \in [0, 1]$ denote a survival time, let $X \in \{0, 1\}$ denote a binary baseline exposure and let Z denote additional baseline covariates. With $\mathbf{H}_t(X, Z)$ denoting the conditional cumulative hazard given (X, Z) , the conditional survival function is $\mathbb{P}(T > t | X, Z) = e^{-\mathbf{H}_t(X, Z)}$, and

$$\Theta_t(Z) := \log \frac{\mathbb{P}(T > t | X = 0, Z)}{\mathbb{P}(T > t | X = 1, Z)} = \mathbf{H}_t(1, Z) - \mathbf{H}_t(0, Z) \quad (4.1.1)$$

quantifies the, possibly heterogeneous, exposure effect as a function of Z . Suppose that

$$\mathbf{H}_t(X, Z) = \int_0^t \mathbf{h}_s(X, Z) ds$$

is given in terms of the conditional hazard $\mathbf{h}_s(X, Z)$, then

$$\Theta_t(Z) = \int_0^t \underbrace{\mathbf{h}_s(1, Z) - \mathbf{h}_s(0, Z)}_{=: \theta_s(Z)} ds = \int_0^t \theta_s(Z) ds$$

is the integrated difference of conditional hazards. We can always express the conditional hazard as

$$\mathbf{h}_s(X, Z) = X\theta_s(Z) + \mathbf{h}_s(0, Z), \quad (4.1.2)$$

and the additive decomposition (4.1.2) does not impose any model restrictions unless we restrict one of the functions $\theta_s(Z)$ or $\mathbf{h}_s(0, Z)$.

We seek an estimand that summarizes the values of $\theta_s(Z)$ over s and Z . To this end, let $w_s(Z)$ denote a *weight process* with $w_s(Z) \geq 0$ and $\int_0^1 \mathbb{E}[w_s(Z)] ds = 1$, and consider the *weighted hazard difference estimand*

$$\gamma_1^w := \int_0^1 \mathbb{E}[\theta_s(Z)w_s(Z)] ds.$$

The real-valued parameter γ_1^w is the endpoint of the functional estimand

$$\gamma_t^w := \int_0^t \mathbb{E}[\theta_s(Z)w_s(Z)] ds, \quad t \in [0, 1].$$

For $w_s(Z) = 1$, the estimand simplifies to the expected cumulative hazard difference $\gamma_t^1 = \mathbb{E}(\Theta_t(Z))$, but other weights are also of interest. For the *partially additive hazards model*, where θ_s is independent of Z , we see that $\gamma_t^w = \int_0^t \theta_s \mathbb{E}[w_s(Z)] ds$ is a cumulative weighted hazard difference. Efficient and (rate) doubly robust estimation was considered for the partially additive model in Duker et al. [2019] and Hou et al. [2023] using semiparametric estimation theory – in particular when $\theta_t = \theta$ is independent of t and $\gamma_t^w = \theta t$. In this paper we introduce an assumption-lean additive hazard estimand and corresponding estimation theory that, as a special case, applies to model (4.1.2) without any model restrictions.

4 Assumption-lean Aalen regression

To give an explicit formula for our assumption-lean estimand for model (4.1.2) within this simplified example, let $Y_t = \mathbb{1}(T \geq t)$ denote the at-risk indicator, and define what we call the *Aalen covariance* weights

$$w_s^{\text{AC}}(Z) := \frac{Y_s \pi_s(Z)(1 - \pi_s(Z))}{\mathbb{E}[Y_s \pi_s(Z)(1 - \pi_s(Z))]}, \quad (4.1.3)$$

where $\pi_s(Z) = \mathbb{P}(X = 1 \mid T \geq s, Z)$. The resulting estimand can be expressed as

$$\gamma_t = \int_0^t \mathbb{E}[\theta_s(Z) w_s^{\text{AC}}(Z)] ds = \int_0^t \frac{\mathbb{E}[\theta_s(Z) Y_s \pi_s(Z)(1 - \pi_s(Z))]}{\mathbb{E}[Y_s \pi_s(Z)(1 - \pi_s(Z))]} ds. \quad (4.1.4)$$

The weights (4.1.3) and the formula (4.1.4) may at first appear unmotivated. In Section 3 in Vansteelandt and Dukes [2022] we find weights similar to (4.1.3) used to represent assumption-lean GLM regression estimands, and also in Section 2 in Vansteelandt et al. [2022] similar weights appear in the representation of assumption-lean Cox regression estimands. ♠

In Example 4.2.6 we show that the Aalen covariance weights and the estimand (4.1.4) appear naturally as an assumption-lean Aalen regression estimand. The estimand γ_t given by (4.1.4) is an example of the *Aalen covariance measure* (ACM) – the general estimand we introduce in Definition 4.2.1. This is, in turn, an example of a *Local Covariance Measure* as introduced by Christgau et al. [2023b]. We show in this paper that ACM can be defined and interpreted as an assumption-lean estimand for any real-valued and time-dependent X -process and time-dependent Z -process, thus generalizing model (4.1.2) considerably.

4.1.2 Organization of paper

We organize the rest of the paper as follows. In Section 4.2, we first introduce our general setup and define our target estimand. We then discuss interpretations of the estimand based on theoretical properties and computations in concrete examples. Conditions on the censoring mechanism are postponed to the end of the section to avoid a technical discussion from the onset. In Section 4.3, we discuss an estimation methods for our target estimand, in particular we propose the x-ACM, which is based on double machine learning estimator. In Section 4.4, we prove that the x-ACM is asymptotically Gaussian with a \sqrt{n} -rate under suitable conditions. The most significant conditions are rate-conditions on nuisance estimators, which are similar to those for other (rate) doubly robust estimators. In Section 4.5, conduct a simulation study that demonstrates the effect of the x-ACM. We conclude with a discussion in Section 4.6 outlining potential future research directions.

The supplementary material consists of Section 4.A, which contains proofs of all our main results and related auxiliary lemmas, and Section 4.B, which includes a discussion of our framework in the context of multiplicative hazards models.

4.2 Target estimand and interpretations

4.2.1 General setup

We consider a survival time $T^* > 0$ censored by a random variable $C \in [0, 1]$, such that the observed time is $T := T^* \wedge C$ together with the indicator $\Delta := \mathbb{1}(T^* \leq C)$. From these we define the observed counting process $N_t = \mathbb{1}(T \leq t, \Delta = 1)$, and the at-risk indicator $Y_t = \mathbb{1}(T \geq t)$. In addition to (T, Δ) , we also observe a covariate process $Z = (Z_t)$ and an exposure process $X = (X_t)$ of particular interest. The covariate process can, in principle, take values in any measurable space, whereas the exposure process is assumed real-valued and left-continuous.

We postpone an elaborate discussion of assumptions on the censoring mechanism to Section 4.2.4. In summary, we introduce Assumption 4.2.1, which is an instance of *independent right-censoring* in the sense of Andersen et al. [1993], and which will remain an assumption throughout the remainder of the manuscript.

For any stochastic process $W = (W_t)$, we let $\bar{W}_t = (W_s)_{s < t}$ denote the history up to time $t > 0$, but not including the timepoint t .¹ Note that by left-continuity, $X_t = X_{t-} := \lim_{s \rightarrow t-} X_s$ is determined by the history \bar{X}_t .

The conditional cumulative hazard \mathbf{H}_t , given the processes X and Z , is defined by

$$\mathbf{H}_t(\bar{X}_t, \bar{Z}_t) := -\log \mathbb{P}(T^* > t \mid \bar{X}_t, \bar{Z}_t).$$

The overall objective is to quantify how $\mathbf{H}_t(\bar{X}_t, \bar{Z}_t)$ depends on the exposure process \bar{X}_t . Such a dependence can, in general, be complicated and depend on the value of the covariate process \bar{Z}_t . Our approach is *assumption-lean* meaning that we do not make strong model assumptions on how \mathbf{H}_t depends on \bar{X}_t and \bar{Z}_t . We will assume, though, that $\mathbf{H}_t(\bar{X}_t, \bar{Z}_t)$ is given by a conditional hazard function, denoted by \mathbf{h}_t , that is, $\mathbf{H}_t(\bar{X}_t, \bar{Z}_t) = \int_0^t \mathbf{h}_s(\bar{X}_s, \bar{Z}_s) ds$.

4.2.2 Aalen Covariance Measure

Before we define our model-free estimand in the most general case, we need to introduce two auxiliary stochastic processes: the predictable projection of X and the Z -compensated counting process.

A stochastic process is said to be *càglàd* if its sample paths are left-continuous and have right-limits at all times. We henceforth assume that X is càglàd and square-integrable. Then, the *predictable projection* of X onto the histories of N and Z can be defined as the unique² càglàd process, denoted by $\Pi = (\Pi_t)$, satisfying that

$$\Pi_t = \mathbb{E}[X_t \mid \bar{N}_t, \bar{Z}_t]. \quad (4.2.5)$$

¹This convention differs from other works such as [Lok, 2017, 2008], which do include the timepoint t in the history. For any left-continuous process, this convention does not affect the information content of the history. In particular, there is no distinction between the conventions when conditioning on \bar{X}_t , see also Remark 4.2.3

²Technically, it is unique up to *evanescence*.

4 Assumption-lean Aalen regression

For more details, see for example Corollary 7.6.8 in Cohen and Elliott [2015]. We will primarily need the process Π_t on the event $(Y_t = 1) = (T \geq t)$, where it takes the value $\pi_t(\bar{Z}_t) := \mathbb{E}[X_t | T \geq t, \bar{Z}_t]$.³

Introducing

$$h_t(\bar{Z}_s) = \mathbb{E}[\mathbf{h}_s(\bar{X}_s, \bar{Z}_s) | T \geq t, \bar{Z}_t],$$

the *innovation theorem*, see, e.g., Theorem II.T14 in Brémaud [1981], gives that h_t is the conditional hazard given Z that satisfies the relation

$$H_t(\bar{Z}_t) := \int_0^t h_s(\bar{Z}_s) ds = -\log \mathbb{P}(T^* > t | \bar{Z}_t).$$

In terms of the conditional hazard given Z we define the Z -compensated counting process

$$M_t := N_t - \int_0^t Y_s h_s(\bar{Z}_s) ds = N_t - H_{t \wedge T}(\bar{Z}_{t \wedge T}).$$

We can now define our target estimand.

Definition 4.2.1 (Aalen Covariance Measure). For each $t \in [0, 1]$ let

$$\rho(t) := \mathbb{E}[Y_t(X_t - \Pi_t)^2] = \mathbb{E}[Y_t(X_t - \pi_t(\bar{Z}_t))^2], \quad (4.2.6)$$

and define, whenever $\rho(t) > 0$, the *residual process* $G = (G_t)$ by

$$G_t := \frac{Y_t(X_t - \Pi_t)}{\rho(t)} = \frac{Y_t(X_t - \pi_t(\bar{Z}_t))}{\rho(t)}. \quad (4.2.7)$$

If $\rho(t) > 0$ for all $t \in [0, 1]$, the *Aalen Covariance Measure (ACM)* is defined as the functional estimand $\gamma = (\gamma_t)$ given by

$$\gamma_t = \mathbb{E} \left[\int_0^t G_s dM_s \right] \quad (4.2.8)$$

for $t \in [0, 1]$. ♣

We discuss the positivity assumption $\rho(t) > 0$ further in Remark 4.4.1. To appreciate Definition 4.2.1 we compute the ACM for the partially additive hazards model.

Example 4.2.2 (Partially additive hazards model). Suppose that

$$\mathbf{h}_t(\bar{X}_t, \bar{Z}_t) = \theta_t X_t + g_t(\bar{Z}_t), \quad (4.2.9)$$

for some $\theta_t \in \mathbb{R}$ and a function $g_t \in L^1(P_{\bar{Z}_t})$, where $P_{\bar{Z}_t}$ denotes the distribution of \bar{Z}_t . The model given by (4.2.9) is the *partially additive hazards model*. It specifies the direct

³We technically define π_t via the Doob-Dynkin lemma such that $Y_t \Pi_t = Y_t \pi_t(\bar{Z}_t)$ holds *surely*. Thus we can ensure that $\pi_t(\bar{Z}_t)$ is a càglàd version of $\mathbb{E}[X_t | T \geq t, \bar{Z}_t]$.

effect of X in a separate term with time-varying coefficient θ_t . For this model the innovation theorem yields the explicit expression

$$h_t(\bar{Z}_t) = \theta_t \pi_t(\bar{Z}_t) + g_t(\bar{Z}_t). \quad (4.2.10)$$

Introducing the (X, Z) -compensated counting process, $\mathbf{M}_t := N_t - \mathbf{H}_{t \wedge T}(\bar{X}_{t \wedge T}, \bar{Z}_{t \wedge T})$ we get by (4.2.10) that

$$M_t = \mathbf{M}_t + \int_0^t \theta_s (X_s - \pi_s(\bar{Z}_s)) ds.$$

Since \mathbf{M}_t is a mean zero martingale and the residual process is predictable,

$$\begin{aligned} \gamma_t &= \mathbb{E} \left[\int_0^t G_s d\mathbf{M}_s + \int_0^t G_s \theta_s (X_s - \pi_s(\bar{Z}_s)) ds \right] \\ &= \mathbb{E} \left[\int_0^t \theta_s \frac{Y_s (X_s - \pi_s(\bar{Z}_s))^2}{\rho(s)} ds \right] = \int_0^t \theta_s ds. \end{aligned} \quad (4.2.11)$$

Hence, the ACM reduces to the cumulative direct effect $\Theta_t = \int_0^t \theta_s ds$ for the partially additive hazards model. \spadesuit

We note that both Dukes et al. [2019], Hou et al. [2023] consider models that fall under the partially additive model given by (4.2.9), and they propose estimation of the cumulative effect Θ_t with doubly robust estimation methods. The example above shows that our ACM estimand reduces to the cumulative effect for the partially additive model, but we will not make this model assumption. We present additional representations and interpretations of the ACM in Section 4.2.3

Remark 4.2.3. The ACM is an example of a *Local Covariance Measure* (LCM), which was introduced by Christgau et al. [2023b] in a more general setting of counting processes, where all definitions and results are formulated in terms of filtrations and intensities. Since this work is focused on survival analysis, the results are formulated using histories and conditional hazards to make them more accessible. To describe the connection between the notation used in this paper and in Christgau et al. [2023b], introduce the two filtrations $\mathcal{F}_t := \sigma(Z_s, N_s; s \leq t)$ and $\mathcal{G}_t := \sigma(X_s, Z_s, N_s; s \leq t)$. Then $\mathcal{G}_{t-} = \sigma(\bar{X}_t, \bar{Z}_t, \bar{N}_t)$ and $\mathcal{F}_{t-} = \sigma(\bar{Z}_t, \bar{N}_t)$, and $\lambda_t := Y_t h_t(\bar{Z}_t)$ and $\boldsymbol{\lambda}_t := Y_t \mathbf{h}_t(\bar{X}_t, \bar{Z}_t)$ are the \mathcal{F}_{t-} and \mathcal{G}_{t-} -predictable intensities, respectively, for the counting process N . The martingale argument in Example 4.2.2 uses that \mathbf{M}_t is a \mathcal{G}_t -martingale, which relies on the implicit Assumption 4.2.1 on the censoring mechanism. The supplementary Section 4.A elaborates on the notation, which is also used in the proofs.

In Christgau et al. [2023b], the LCM was introduced for any \mathcal{G}_t -predictable residual process G_t satisfying $\mathbb{E}[G_t | \mathcal{F}_{t-}] = 0$. The residual process given by (4.2.7) is a specific example of such a process. The LCM was introduced in Christgau et al. [2023b] to quantify deviations from the hypothesis that N is *locally independent* of X given \mathcal{F}_t . The ACM, being an example of a LCM, enjoys the same interpretation: under local independence, γ is equal to the zero function. The ACM was, however, not considered in Christgau et al. [2023b], and we proceed to show that the ACM enjoys the interpretation as an effect measure generalizing the cumulative effect in a partially additive hazards model. \diamond

4.2.3 Interpretations of the ACM as an effect measure

By Proposition 2.6 in Christgau et al. [2023b], the LCM, and therefore also the ACM, has the following equivalent representation for each $t \in [0, 1]$,

$$\gamma_t = \int_0^t \text{Cov}(G_s, Y_s(\mathbf{h}_s(\bar{X}_s, \bar{Z}_s) - h_s(\bar{Z}_s))) ds \quad (4.2.12)$$

$$= \int_0^t \mathbb{E}[G_s(\mathbf{h}_s(\bar{X}_s, \bar{Z}_s) - h_s(\bar{Z}_s))] ds, \quad (4.2.13)$$

where (4.2.13) follows from G_s having mean zero and $Y_s \in \{0, 1\}$. Example 4.2.2 showed that the ACM equals the cumulative direct effect under the partially additive hazards model. The following proposition asserts that in general – when (4.2.9) may not hold – the ACM equals the cumulated coefficient in the L^2 -projection of the conditional hazard \mathbf{h}_t onto a partially additive model.

Proposition 4.2.4. *Consider the minimization problem*

$$\underset{\vartheta, g}{\text{minimize:}} \quad \mathbb{E} \left[\int_0^1 Y_t(\mathbf{h}_t(\bar{X}_t, \bar{Z}_t) - \vartheta_t X_t - g_t(\bar{Z}_t))^2 dt \right] \quad (4.2.14)$$

over all measurable functions $\vartheta: [0, 1] \rightarrow \mathbb{R}$ and collections of functions $(g_t)_{t \in [0, 1]}$ such that $g_t(\bar{Z}_t)$ is progressively measurable. Let $\vartheta^* = (\vartheta_t^*)$ be given by

$$\vartheta_t^* := \text{Cov}(G_t, Y_t(\mathbf{h}_t(\bar{X}_t, \bar{Z}_t) - h_t(\bar{Z}_t)))$$

for $t \in [0, 1]$.

Then there exists g^* such that (ϑ^*, g^*) is a solution to (4.2.14), and ϑ^* is unique up to modifications on a null set of $[0, 1]$. In particular, the ACM can be expressed as

$$\gamma_t = \int_0^t \vartheta_s ds, \quad t \in [0, 1],$$

for any solution (ϑ, g) to (4.2.14).

All proofs are deferred to the supplement.

Example 4.2.5. Continuing Example 4.2.2 we see that for the partially additive hazards model given by (4.2.9), the minimizer given by Proposition 4.2.4 is

$$\begin{aligned} \vartheta_t^* &= \text{Cov}(G_t, Y_t(\mathbf{h}_t(\bar{X}_t, \bar{Z}_t) - h_t(\bar{Z}_t))) \\ &= \text{Cov}(G_t, \theta_t Y_t(X_t - \pi_t(\bar{Z}_t))) \\ &= \frac{\theta_t \mathbb{E}[Y_t(X_t - \pi_t(\bar{Z}_t))^2]}{\rho(s)} = \theta_t. \end{aligned}$$

Thus, as also established directly in Example 4.2.2, for the partially additive hazards model, $\gamma_t = \Theta_t = \int_0^t \theta_s ds$. ♠

Example 4.2.6. As a generalization of the introductory Example 4.1.1, suppose that $X_t \in \{0, 1\}$ and that

$$\mathbf{h}_t(\bar{X}_t, \bar{Z}_t) = X_t \theta_t(\bar{Z}_t) + g_t(\bar{Z}_t).$$

If $\theta_t(\bar{Z}_t)$ does not depend on \bar{Z}_t , this model is also a partially additive hazards model. In general, we see that

$$h_t(\bar{Z}_t) = \pi_t(\bar{Z}_t) \theta_t(\bar{Z}_t) + g_t(\bar{Z}_t)$$

where $\pi_t(\bar{Z}_t) = \mathbb{P}(X_t = 1 \mid T \geq t, \bar{Z}_t)$. Hence,

$$\begin{aligned} \rho(t) &= \mathbb{E} [Y_t (X_t - \pi_t(\bar{Z}_t))^2] \\ &= \mathbb{E} [Y_t \mathbb{E} [(X_t - \pi_t(\bar{Z}_t))^2 \mid T \geq t, \bar{Z}_t]] \\ &= \mathbb{E} [Y_t \pi_t(\bar{Z}_t) (1 - \pi_t(\bar{Z}_t))]. \end{aligned}$$

The minimizer given by Proposition 4.2.4 is

$$\begin{aligned} \vartheta_t^* &= \text{Cov}(G_t, Y_t (\mathbf{h}_t(\bar{X}_t, \bar{Z}_t) - h_t(\bar{Z}_t))) \\ &= \text{Cov}(G_t, \theta_t(\bar{Z}_t) Y_t (X_t - \pi_t(\bar{Z}_t))) \\ &= \mathbb{E} \left[\theta_t(\bar{Z}_t) \frac{Y_t \mathbb{E} [(X_t - \pi_t(\bar{Z}_t))^2 \mid T \geq t, \bar{Z}_t]}{\rho(t)} \right] \\ &= \mathbb{E} \left[\theta_t(\bar{Z}_t) \frac{Y_t \pi_t(\bar{Z}_t) (1 - \pi_t(\bar{Z}_t))}{\mathbb{E} [Y_t \pi_t(\bar{Z}_t) (1 - \pi_t(\bar{Z}_t))]} \right] \\ &= \mathbb{E} [\theta_t(\bar{Z}_t) w_t^{\text{AC}}(\bar{Z}_t)], \end{aligned}$$

where the Aalen covariance weights are

$$w_t^{\text{AC}}(Z) = \frac{Y_t \pi_t(\bar{Z}_t) (1 - \pi_t(\bar{Z}_t))}{\mathbb{E} [Y_t \pi_t(\bar{Z}_t) (1 - \pi_t(\bar{Z}_t))]} \quad (4.2.15)$$

In the special case with baseline exposure and baseline covariates, as treated in Example 4.1.1, these weights reduce to the weights in (4.1.3), and the estimand given by (4.1.4) is, indeed, the ACM.

To give a more specific example of a model that is *not* a partially additive model, suppose that $Z_t \in \mathbb{R}^d$ and that $\theta_t(\bar{Z}_t) = \theta^T Z_t$ for $\theta \in \mathbb{R}^d$. Then

$$\gamma_t^w = \theta^T \int_0^t \mathbb{E} [Z_s w_s^{\text{AC}}(\bar{Z}_s)] ds = \theta^T \omega_t^{\text{AC}} = \sum_{k=1}^d \theta_k \omega_{k,t}^{\text{AC}}$$

is a linear combination of the θ_k -parameters with $\omega_{k,t}^{\text{AC}} = \int_0^t \mathbb{E} [Z_{k,s} w_s^{\text{AC}}(\bar{Z}_s)] ds$. ♠

Proposition 4.2.4 allows us to interpret the ACM as a cumulative direct effect in the best L^2 model-approximation by a partially additive hazards model. This underlines that the ACM is well-defined and retains some level of interpretability for (small) deviations

from the partially additive model. In contrast, the model-based target parameter in (4.2.11) is *a priori* undefined when the partially additive model is misspecified.

Beyond the partially additive model, as treated in Examples 4.2.2 and 4.2.5, we cannot expect the ACM to have a simple closed-form expression. Even for the simple baseline model (4.1.2), the innocent looking representation

$$\gamma_t = \int_0^t \mathbb{E} [\theta_s(Z) w_s^{\text{AC}}(Z)] ds,$$

that follows from Example 4.2.6, involves the weights and in turn the predictable projections $\pi_t(Z)$. We elaborate on the computation of $\pi_t(Z)$ for a multiplicative (baseline) hazards model in supplementary Section 4.B, and we discuss the representation of the ACM further for the Cox proportional hazards model in Example 4.6.2.

4.2.4 Censoring

Let $N_t^* := \mathbb{1}(T^* \leq t)$ denote the uncensored counting process, and let $N_t^c = \mathbb{1}(C \leq t)$ denote the counting process for censoring. Recall the definitions of the filtrations \mathcal{F}_t and \mathcal{G}_t in Remark 4.2.3, and define similarly the filtrations $\mathcal{F}_t^* := \sigma(Z_s, N_s^*; s \leq t)$ and $\mathcal{G}_t^* := \sigma(X_s, Z_s, N_s^*; s \leq t)$.

We assume two conditions on the censoring mechanism, both are forms of *independent right-censoring* in the sense of Andersen et al. [1993]. We refer to Martinussen and Scheike [2006], Andersen et al. [1993] for interpretations and intuition behind this type of censoring. Formulated in terms of local independence as in [Røysland et al., 2022], we specifically assume that N_t^* is locally independent of N_t^c given both \mathcal{F}_t and \mathcal{G}_t . Alternatively, the censoring condition may be formulated as in Assumption 4.2.1 below.

An equivalent definition of the conditional hazards, h_s and \mathbf{h}_s , is that

$$N_t^* - \int_0^t \mathbb{1}(T^* \geq s) h_s(\bar{Z}_s) ds \quad \text{and} \quad N_t^* - \int_0^t \mathbb{1}(T^* \geq s) \mathbf{h}_s(\bar{X}_t, \bar{Z}_s) ds$$

are martingales with respect to the filtrations \mathcal{F}_t^* and \mathcal{G}_t^* , respectively. Independent right-censoring means that the conditional hazards also yield valid compensators of the observed counting process in the following sense.

Assumption 4.2.1 (Independent right-censoring). *It holds that*

$$N_t - \int_0^t \mathbb{1}(T \geq s) h_s(\bar{Z}_s) ds \quad \text{and} \quad N_t - \int_0^t \mathbb{1}(T \geq s) \mathbf{h}_s(\bar{X}_t, \bar{Z}_s) ds$$

are martingales with respect to the filtrations \mathcal{F}_t and \mathcal{G}_t , respectively.

All of our results regarding estimation and asymptotic theory will rely on Assumption 4.2.1, and no other assumptions on censoring.

It is natural to ask if the ACM depends on the censoring distribution. For the partially additive hazards model (4.2.9), the ACM is determined by the conditional hazard \mathbf{h} and is thus unaffected by the censoring distribution in this special case. However, the more

general interpretations of the ACM in (4.2.12) and in Proposition 4.2.4 depend *a priori* on the censoring distribution via the at-risk indicator Y_t . In view of (4.2.12), the ACM is unaffected by censoring if for all $t \in [0, 1]$,

$$\frac{\mathbb{E}[u(\bar{X}_t, \bar{Z}_t) | C \geq t, T^* \geq t]}{\mathbb{E}[v(X_t, \bar{Z}_t) | C \geq t, T^* \geq t]} = \frac{\mathbb{E}[u(\bar{X}_t, \bar{Z}_t) | T^* \geq t]}{\mathbb{E}[v(X_t, \bar{Z}_t) | T^* \geq t]},$$

where

$$\begin{aligned} u(\bar{X}_t, \bar{Z}_t) &= (X_t - \pi_t(\bar{Z}_t))(\mathbf{h}_t(\bar{X}_t, \bar{Z}_t) - h_t(\bar{Z}_t)), \\ v(X_t, \bar{Z}_t) &= (X_t - \pi_t(\bar{Z}_t))^2. \end{aligned}$$

This holds, for example, when $C \perp\!\!\!\perp X, Z, T^*$.

4.3 Estimation

We let $[n] = \{1, \dots, n\}$ for $n \in \mathbb{N}$ and assume that a sample of n observations, $(\mathbf{W}^{(i)})_{i \in [n]} = (T^{(i)}, \Delta^{(i)}, X^{(i)}, Z^{(i)})_{i \in [n]}$, is available as independent copies of the template observation $\mathbf{W} := (T, \Delta, X, Z)$. For each index $i \in [n]$, we use a superscript notation, $^{(i)}$, for the corresponding quantities introduced in Section 4.2. For example, $M_t^{(i)} = N_t^{(i)} - H_{t \wedge T^{(i)}}(\bar{Z}_{t \wedge T^{(i)}}^{(i)})$ is the $Z^{(i)}$ -compensated counting process, where $N_t^{(i)} = \mathbb{1}(T^{(i)} \leq t, \Delta^{(i)} = 1)$.

We will use *cross-fitting* to construct an estimator of the ACM, but for simplicity we first describe our estimator based on single sample split $[n] = \mathcal{I}_1 \cup \mathcal{I}_2$, where \mathcal{I}_1 and \mathcal{I}_2 are disjoint sets. For convenience we let $n_1 := |\mathcal{I}_1|$ and $n_2 := |\mathcal{I}_2|$, and we assume that $\min\{n_1, n_2\} \rightarrow \infty$ as $n \rightarrow \infty$. The data indexed by \mathcal{I}_1 , henceforth called the *training data*, is used to compute estimates $\hat{\pi}_t$ and \hat{h}_t of the functions π_t and h_t .

The remaining data $(\mathbf{W}_i)_{i \in \mathcal{I}_2}$ is used to compute the estimate of ACM by the following three-step procedure:

- (a) Compute an estimate of $\rho(\cdot)$, as defined by (4.2.6). Such an estimate could be

$$\tilde{\rho}(t) = \frac{1}{n_2} \sum_{i \in \mathcal{I}_2} Y_t^{(i)} (X_t^{(i)} - \hat{\pi}_t(\bar{Z}_t^{(i)}))^2. \quad (4.3.16)$$

Since we need to divide by an estimate of $\rho(\cdot)$, it will be convenient to consider a clipped estimate of the form $\hat{\rho}(\cdot) = \max\{\tilde{\rho}(\cdot), c\}$ for some $c > 0$. We return to this discussion in both the general theory and for practical considerations.

- (b) Compute estimates of the residual process and $Z^{(i)}$ -compensated counting process by

$$\begin{aligned} \hat{G}_t^{(i)} &= \hat{\rho}(t)^{-1} Y_t^{(i)} (X_t^{(i)} - \hat{\pi}_t(\bar{Z}_t^{(i)})), \\ \hat{M}_t^{(i)} &= N_t^{(i)} - \int_0^t Y_s \hat{h}_s(\bar{Z}_s^{(i)}) ds, \end{aligned} \quad (4.3.17)$$

for each index $i \in \mathcal{I}_2$.

(c) Compute the ACM estimator, denoted $\hat{\gamma} = (\hat{\gamma}_t)$, by

$$\hat{\gamma}_t = \frac{1}{n_2} \sum_{i \in \mathcal{I}_2} \int_0^t \hat{G}_s^{(i)} dM_s^{(i)}, \quad t \in [0, 1]. \quad (4.3.18)$$

While the formula (4.3.18) is the same for the LCM estimator proposed in Christgau et al. [2023b], there is an important technical distinction: the residual estimate $\hat{G}_t^{(i)}$ cannot be determined from the training data and i -th observation alone, since the estimate $\hat{\rho}$ is based on all of the data indexed by \mathcal{I}_2 . Thus the residual estimates are not conditionally independent given the training data, which makes the theoretical analysis of the estimator a bit more involved. See the discussion in Section 4.6 for an elaboration regarding this difference.

A general procedure for computing the ACM estimator is described in Algorithm 5. There are additional numerical challenges when computing the ACM estimator in practice. The estimator (4.3.18) is, for instance, defined as an integral, and its practical evaluation may require numerical integration. We return to this practical discussion in Section 4.5.1.

Algorithm 5: Single split estimate of ACM

- 1 **input:** sample $(T^{(i)}, \Delta^{(i)}, X^{(i)}, Z^{(i)})_{i \in [n]}$ and partition $\mathcal{I}_1 \cup \mathcal{I}_2 = [n]$;
 - 2 **options:** regression methods for estimation of $h_t(\cdot)$ and $\pi_t(\cdot)$;
 - 3 **begin**
 - 4 fit conditional hazard $\hat{h}_t(\cdot)$ by regressing $(T^{(i)}, \Delta^{(i)})_{i \in \mathcal{I}_1}$ onto $(\bar{Z}_t^{(i)})_{i \in \mathcal{I}_1}$;
 - 5 fit projection $\hat{\pi}_t(\cdot)$ by regressing $(X_t^{(i)})_{i \in \mathcal{I}_1}$ onto $(T^{(i)}, \Delta^{(i)}, \bar{Z}_t^{(i)})_{i \in \mathcal{I}_1}$;
 - 6 compute $\hat{\rho}(\cdot)$ using $\hat{\pi}_t(\cdot)$ and $(T^{(i)}, \Delta^{(i)}, X^{(i)}, Z^{(i)})_{i \in \mathcal{I}_2}$;
 - 7 compute $(\hat{G}^{(i)}, M^{(i)})_{i \in \mathcal{I}_2}$ according to (4.3.17);
 - 8 **output:** the ACM estimate $\hat{\gamma}_t = \frac{1}{n_2} \sum_{i \in \mathcal{I}_2} \int_0^t \hat{G}_s^{(i)} dM_s^{(i)}$;
-

4.3.1 Cross-fitted ACM estimator

The ACM estimator defined in (4.3.18), and described in Algorithm 5, is based on a single data split $[n] = \mathcal{I}_1 \cup \mathcal{I}_2$ and only uses the data indexed by \mathcal{I}_2 for estimation of the target. The more sophisticated technique of 2-fold cross-fitting produces an additional estimate by swapping the roles of \mathcal{I}_1 and \mathcal{I}_2 and aggregates the two estimates by averaging. More generally, K -fold cross-fitting partitions the data into K disjoint folds, and cyclically uses each fold for estimation of the ACM. We assume that the folds have sizes at least $\lfloor n/K \rfloor$ and at most $\lceil n/K \rceil$. The procedure is described in detail in Algorithm 6, and the resulting cross-fitted estimator is referred to as the x-ACM and is denoted by $\check{\gamma}$. We refer to Chernozhukov et al. [2018] and references therein for a more general description of cross-fitting.

Algorithm 6: K -fold cross-fitted X-ACM

1 **input:** sample $(T^{(i)}, \Delta^{(i)}, X^{(i)}, Z^{(i)})_{i \in [n]}$ and partition $J_1 \cup \dots \cup J_K = [n]$;
2 **options:** options for Algorithm 5;
3 **begin**
4 **for** $k = 1, \dots, K$ **do**
5 compute $\hat{\gamma}^k$ using Algorithm 5 with $\mathcal{I}_1 = \bigcup_{\ell \neq k} J_\ell$ and $\mathcal{I}_2 = J_k$;
6 **output:** the X-ACM estimate $\check{\gamma} = \frac{1}{K} \sum_{k=1}^K \hat{\gamma}^k$

4.4 Asymptotic theory

We present in this section asymptotic representations of the estimation errors, $\hat{\gamma} - \gamma$ and $\check{\gamma} - \gamma$, for the sample split and cross-fit estimators. We leverage these representations to derive asymptotic Gaussian process limits, which can be used for statistical inference. The asymptotic properties of the general LCM estimator were analyzed in Christgau et al. [2023b] for the specific purpose of understanding type I and type II errors of the *Local Covariance Test* – a test of local independence based on the LCM estimator. In Christgau et al. [2023b] it was established, under suitable regularity conditions, that:

- The asymptotic limit of the LCM estimator is a Gaussian martingale under the null hypothesis $\gamma = 0$.
- The LCM estimator is \sqrt{n} -consistent using the *additive residual process* $X_t - \Pi_t$.

These results cannot directly be applied to our proposed ACM estimator, as the residual estimates in (4.3.17) are constructed with a mutual dependence via $\hat{\rho}$. However, by tailoring the theoretical analysis to the residual process given by (4.2.7), and its corresponding estimates in (4.3.17), we establish the asymptotic Gaussian process limit of $\sqrt{n_2}(\hat{\gamma} - \gamma)$, also under the alternative $\gamma \neq 0$.

It will be instructive and useful for the general analysis to first understand the asymptotics of the ACM estimator in the simplified scenario where ρ is assumed known. In this case we can largely use the results from Christgau et al. [2023b], and we can then in Section 4.4.2 combine these results with a decomposition of the full estimation error to arrive at the general results when ρ is estimated.

4.4.1 Asymptotics with a ρ -oracle

In this section we analyze the asymptotics of the ρ -oracle estimator

$$\tilde{\gamma}_t := \frac{1}{n_2} \sum_{i \in \mathcal{I}_2} \int_0^t \frac{Y_t^{(i)}(X_t^{(i)} - \hat{\pi}_t(\bar{Z}_t^{(i)}))}{\rho(s)} dM_s^{(i)}, \quad t \in [0, 1]. \quad (4.4.19)$$

That is, we assume that ρ is known and not estimated. Since the terms in the sum are conditionally i.i.d. given $(\hat{\pi}_t, \hat{h}_t)$ – and these nuisance estimates are obtained from the

4 Assumption-lean Aalen regression

training data only – the estimator $\tilde{\gamma}_t$ fits into the framework of Christgau et al. [2023b]. While Christgau et al. [2023b] did not establish an asymptotic limit in the alternative $\gamma \neq 0$, we can follow their initial asymptotic analysis. To this end, it will be convenient to first define the processes

$$E_t := Y_t(X_t - \Pi_t) = \rho(t)G_t, \quad (4.4.20)$$

$$\hat{E}_t := Y_t(X_t - \hat{\pi}_t(\bar{Z}_t)) = \hat{\rho}(t)\hat{G}_t. \quad (4.4.21)$$

The following assumption is essentially a reformulation of Assumption 4.1 in [Christgau et al., 2023b].

Assumption 4.4.1 (Boundedness). *There exist constants $C_{\mathbf{h}}, C_E, c_\rho > 0$ such that for each $t \in [0, 1]$:*

- (i) $Y_t \max\{\mathbf{h}_t(\bar{X}_t, \bar{Z}_t), \hat{\mathbf{h}}_t(\bar{Z}_t)\} \leq C_{\mathbf{h}}$.
- (ii) $\max\{|E_t|, |\hat{E}_t|\} \leq C_E$.
- (iii) $\rho(t) \geq c_\rho$.

A few remarks regarding Assumption 4.4.1 are in order.

Remark 4.4.1 (Remarks on boundedness assumptions).

- (a) From the innovation theorem it follows that

$$Y_t \mathbf{h}_t(\bar{X}_t, \bar{Z}_t) \leq C_{\mathbf{h}} \implies Y_t h_t(\bar{Z}_t) \leq C_{\mathbf{h}}.$$

Since hazards are non-negative, the assumption (i) also entails the bounds $Y_t |\mathbf{h}_t(\bar{X}_t, \bar{Z}_t) - h_t(\bar{Z}_t)| \leq C_{\mathbf{h}}$ and $Y_t |h_t(\bar{Z}_t) - \hat{h}_t(\bar{Z}_t)| \leq C_{\mathbf{h}}$.

- (b) If the exposure process takes values in an interval, that is, $X_t \in [a, b]$, then also $\Pi_t \in [a, b]$ and it is reasonable to impose $\hat{\pi}_t \in [a, b]$. In this case, (ii) is satisfied with $C_E = b - a$.

- (c) Note that

$$\rho(t) = \mathbb{P}(T \geq t) \mathbb{E}[(X_t - \Pi_t)^2 | T \geq t].$$

Thus Assumption 4.4.1 (iii) is equivalent to the two conditions: $\mathbb{P}(T \geq 1) > 0$ and $\inf_{t \in [0, 1]} \mathbb{E}[(X_t - \Pi_t)^2 | T \geq t] > 0$. The first is natural in order to perform statistical inference up to time $t = 1$. The second condition should also not be surprising: even for the partially additive model we would need $\mathbb{E}[(X_t - \Pi_t)^2 | T \geq t] > 0$ for all $t \in [0, 1]$ to identify the coefficient θ_t .

In case that (iii) appears to be violated for the observed data, it is always possible to restrict the observation window and target the ACM over a shorter timespan.

◇

To establish convergence it is, of course, also necessary to control the estimation error of the nuisance parameters.

Assumption 4.4.2 (Analogous to Assumption 4.2 in Christgau et al. [2023b]). *With*

$$a_n := \int_0^1 \mathbb{E} \left[Y_t (\hat{\pi}_t(\bar{Z}_t) - \pi_t(\bar{Z}_t))^2 \right] dt, \quad \text{and} \quad b_n := \int_0^1 \mathbb{E} \left[Y_t (\hat{h}_t(\bar{Z}_t) - h_t(\bar{Z}_t))^2 \right] dt,$$

it holds that

$$\max\{a_n, b_n, n_2 \cdot a_n b_n\} \longrightarrow 0$$

as $n \rightarrow \infty$.

Under these assumptions, it follows that $\tilde{\gamma}_t$ is asymptotically equivalent to the corresponding oracle term, where the nuisance estimates $(\hat{\pi}_t, \hat{h}_t)$ are replaced with (π_t, h_t) .

Proposition 4.4.2. *Under Assumptions 4.4.1 and 4.4.2, it holds that*

$$\sup_{t \in [0,1]} |\sqrt{n_2} \cdot \tilde{\gamma}_t - U_t| \xrightarrow{P} 0$$

as $n \rightarrow \infty$, where (U_t) is the oracle process given by

$$U_t = \frac{1}{\sqrt{n_2}} \sum_{i \in \mathcal{I}_2} \int_0^t G_s^{(i)} dM_s^{(i)}, \quad t \in [0, 1].$$

The proof of Proposition 4.4.2 is largely based on the results of Christgau et al. [2023b], with a few minor modifications. Under Assumptions 4.4.1 and 4.4.2, it is possible to establish that the process $U - \sqrt{n_2} \gamma$ converges in distribution to a continuous Gaussian process by appealing to a central limit theorem in Skorokhod space. However, as we are primarily interested in the actual ACM estimator, $\hat{\gamma}$, we return to the asymptotics of the process (U_t) in the proof of Theorem 4.4.6, see Section 4.A.6 in the supplement, where it is analyzed in combination with another non-vanishing term.

Remark 4.4.3. It should be possible to weaken the rate requirements to the empirical errors. To wit, if we define

$$\begin{aligned} \tilde{a}_n &:= \frac{1}{n_2} \sum_{i \in \mathcal{I}_2} \int_0^1 Y_t^{(i)} (\hat{\pi}_t(\bar{Z}_t^{(i)}) - \pi_t(\bar{Z}_t^{(i)}))^2 dt, \\ \tilde{b}_n &:= \frac{1}{n_2} \sum_{i \in \mathcal{I}_2} \int_0^1 Y_t^{(i)} (\hat{h}_t(\bar{Z}_t^{(i)}) - h_t(\bar{Z}_t^{(i)}))^2 dt, \end{aligned}$$

then Assumption 4.4.2 can likely be replaced by the weaker condition that

$$\max\{\tilde{a}_n, \tilde{b}_n, n_2 \cdot \tilde{a}_n \tilde{b}_n\} \xrightarrow{P} 0, \quad n \rightarrow \infty.$$

In fact, $(\tilde{a}_n, \tilde{b}_n) \xrightarrow{P} 0$ is equivalent with $(a_n, b_n) = (\mathbb{E}[\tilde{a}_n], \mathbb{E}[\tilde{b}_n]) \rightarrow 0$ under Assumption 4.4.1 due to uniform integrability. The scaled product error, $n_2 \cdot \tilde{a}_n \tilde{b}_n$, is not *a priori* uniformly integrable, so its convergence does not imply that $n_2 \cdot a_n b_n \rightarrow 0$. \diamond

4.4.2 Asymptotics when ρ is estimated

To analyze the ACM estimator in the general case, where ρ is estimated, we consider the decomposition

$$\sqrt{n_2} \cdot \hat{\gamma} = \sqrt{n_2} \cdot \tilde{\gamma} + B, \quad (4.4.22)$$

where $\tilde{\gamma}$ is the ρ -oracle estimator defined as in (4.4.19), and where B is the process given by

$$B_t = \frac{1}{\sqrt{n_2}} \sum_{i \in \mathcal{I}_2} \int_0^t \left(\frac{1}{\hat{\rho}(s)} - \frac{1}{\rho(s)} \right) \hat{E}_s^{(i)} dM_s^{(i)}.$$

In Proposition 4.4.2 we found an asymptotically equivalent oracle process for $\sqrt{n_2} \cdot \tilde{\gamma}$. We proceed to find an asymptotically equivalent oracle process for B , and then we identify the limit distribution of the sum of the oracle processes. The decompositions of $\tilde{\gamma}$ and B into oracle terms and remainder terms can be found in the proofs of Proposition 4.4.2 and Theorem 4.4.5.

In view of Assumption 4.4.1 (iii), we will for our general theory, and in the residual estimates $\hat{G}^{(i)}$, consider the clipped estimator $\hat{\rho}$ given by

$$\hat{\rho}(t) = \max\{c_{\hat{\rho}}, \tilde{\rho}(t)\}, \quad (4.4.23)$$

for a $c_{\hat{\rho}} > 0$ and where $\tilde{\rho}$ is the empirical estimate from (4.3.16). We assume that $c_{\hat{\rho}}$ is chosen such that $\lim_{n \rightarrow \infty} c_{\hat{\rho}} = 0$ and $c_{\hat{\rho}}^{-1} = o(\sqrt{n_2})$, or in other words, $c_{\hat{\rho}}$ tends to zero, but slower than $1/\sqrt{n_2}$. In particular, this implies that $c_{\hat{\rho}} < c_{\rho}$ eventually, and thus the event $(\hat{\rho} \neq \tilde{\rho})$ has vanishing probability.

To control the error of $\hat{\rho}$ with a sufficiently fast rate, we require not only that $\hat{\pi}_t$ is estimated consistently, but that it can be estimated with a $n^{1/4}$ -rate.

Assumption 4.4.3. *With $a'_n := \int_0^1 \mathbb{E}[Y_t(\hat{\pi}_t(\bar{Z}_t) - \pi_t(\bar{Z}_t))^4] dt$, it holds that $n_2 a'_n \rightarrow 0$ as $n \rightarrow \infty$.*

This assumption is analogous to standard assumptions in the partially linear model, cf. Assumption 4.1(e)(ii) in Chernozhukov et al. [2018]. However, we have assumed the rate on the slightly more restrictive 4-norm. We believe that this rate requirement could be relaxed to $\sqrt{n_2} a_n \rightarrow 0$ under (potentially) additional mild assumptions. The reason that a'_n simplifies the asymptotic analysis is that it controls the estimation error of $\hat{\rho}$ in terms of the 2-norm.

Proposition 4.4.4. *Under Assumptions 4.4.1 and 4.4.3, it holds that*

$$\limsup_{n \rightarrow \infty} n_2 \int_0^1 \mathbb{E}[(\hat{\rho}(t) - \rho(t))^2] dt \leq \frac{C_E^4}{4}. \quad (4.4.24)$$

The following result describes the asymptotics of B .

Theorem 4.4.5. *Under Assumptions 4.4.1, 4.4.2, and 4.4.3, it holds that*

$$\sup_{t \in [0,1]} |B_t - V_t| \xrightarrow{P} 0,$$

where $V = (V_t)$ is the process given by

$$V_t = \sqrt{n_2} \gamma_t - \frac{1}{\sqrt{n_2}} \sum_{i \in \mathcal{I}_2} \int_0^t \frac{E_s^{(i)2}}{\rho(s)} \mathbb{E}[G_s(\mathbf{h}_s(\bar{X}_s, \bar{Z}_s) - h_s(\bar{Z}_s))] ds.$$

The proof of Theorem 4.4.5 does not readily follow from any of the results in Christgau et al. [2023b], and can be dissected into two novel steps: first control the approximation

$$\left(\frac{1}{\hat{\rho}(s)} - \frac{1}{\rho(s)} \right) \left(X_s^{(i)} - \hat{\pi}_s(\bar{Z}_s^{(i)}) \right) dM_s^{(i)} \approx \frac{\rho(s) - \hat{\rho}(s)}{\rho(s)^2} (X_s^{(i)} - \Pi_s^{(i)}) dM_s^{(i)},$$

asymptotically (Proposition 4.A.2), and then use a functional law of large numbers to reduce all stochasticity to the stochastics from $\rho - \hat{\rho}$ (Proposition 4.A.3).

To summarize, we have considered the decomposition $\sqrt{n_2} \cdot \hat{\gamma} = \sqrt{n_2} \cdot \tilde{\gamma} + B$, and established under Assumptions 4.4.1, 4.4.2, and 4.4.3 that

$$\sqrt{n_2} \cdot \tilde{\gamma} = U + o_P(n^{-1/2}) \quad \text{and} \quad B = V + o_P(n^{-1/2}).$$

In view of (4.2.13) we can regard γ as a signed measure on $[0, 1]$ with Radon-Nikodym derivative

$$\frac{d\gamma_s}{ds} := \mathbb{E}[G_s(\mathbf{h}_s(\bar{X}_s, \bar{Z}_s) - h_s(\bar{Z}_s))],$$

which gives that

$$U_t + V_t - \sqrt{n_2} \gamma_t = \frac{1}{\sqrt{n_2}} \sum_{i \in \mathcal{I}_2} \left(\int_0^t G_s^{(i)} dM_s^{(i)} - \int_0^t E_s^{(i)} G_s^{(i)} d\gamma_s \right). \quad (4.4.25)$$

The integrand ‘ $G_s^{(i)} dM_s^{(i)} - E_s^{(i)} G_s^{(i)} d\gamma_s$ ’ corresponds to the oracle term of a partially linear model on an ‘infinitesimal scale’. Theorem 4.A.6 in the supplement establishes asymptotic normality of (4.4.25) based on a central limit theorem in Skorokhod space.

Combining Proposition 4.4.2, Theorem 4.4.5, and Theorem 4.A.6, we obtain asymptotic Gaussianity of the ACM estimator and the x-ACM, $\check{\gamma}$, described in Section 4.3.1.

Theorem 4.4.6. *Let $\Gamma = (\Gamma_t)_{t \in [0,1]}$ be a continuous Gaussian process with mean zero and the same covariance function as the process $(\int_0^t G_s dM_s - \int_0^t E_s G_s d\gamma_s)_{t \in [0,1]}$.*

Under Assumptions 4.4.1, 4.4.2, and 4.4.3, it holds that

$$\sqrt{n_2}(\hat{\gamma} - \gamma) \xrightarrow{d} \Gamma \quad \text{and} \quad \sqrt{n}(\check{\gamma} - \gamma) \xrightarrow{d} \Gamma$$

with respect to the uniform topology as $n \rightarrow \infty$.

4 Assumption-lean Aalen regression

Note that the cross-fitted estimate is scaled with \sqrt{n} rather than $\sqrt{n_2}$, and is thus more efficient than the single-split estimator $\hat{\gamma}$.

Remark 4.4.7. Having established asymptotic Gaussianity of the ACM estimator, it is natural to ask if the (co)variance function of Γ can be estimated in order to perform statistical inference for γ . Let \mathcal{V} be the variance function of the Gaussian process Γ . Under the null hypothesis that $\gamma = 0$, Γ is a martingale and

$$\mathcal{V}(t) = \text{Var} \left(\int_0^t G_s dM_s \right) = \mathbb{E} \left[\int_0^t G_s^2 dN_s \right],$$

which can be estimated by

$$\frac{1}{n_2} \sum_{i \in \mathcal{I}_2} \left(\int_0^t \hat{G}_s^{(i)} dM_s^{(i)} \right)^2 \quad \text{or} \quad \frac{1}{n_2} \sum_{i \in \mathcal{I}_2} \int_0^t (\hat{G}_s^{(i)})^2 dN_s^{(i)},$$

cf. Proposition 4.7 in Christgau et al. [2023b]. However, if the main purpose is to test the null hypothesis, it suffices to use the LCM with the additive residual process $X_t - \Pi_t$ that omits the scaling by ρ and thus avoids estimation of ρ . Then, a test of this null can be carried out using the X-LCT from Christgau et al. [2023b]. However, if a pointwise confidence band for the ACM estimator is desired, we surmise that

$$\hat{\mathcal{V}}(t) = \frac{1}{n_2} \sum_{i \in \mathcal{I}_2} \left(\int_0^t \hat{G}_s^{(i)} dM_s^{(i)} - \int_0^t \hat{E}_s^{(i)} \hat{G}_s^{(i)} d\hat{\gamma}_s \right)^2$$

is a consistent estimator of $\mathcal{V}(t)$, though a proof is outstanding. \diamond

4.5 Simulation study

Simulations were conducted in the special case with time-independent X and Z , and with administrative censoring $C = 1$. Two different settings were considered for the data generating process:

(\mathcal{S}_{lin}) An Aalen additive model with

$$\begin{aligned} Z &= (Z_1, \dots, Z_d) \sim \text{Unif}(0, 1)^{\otimes d}, \\ X | Z &\sim \text{Unif}(Z_1, Z_1 + 1), \\ \mathbf{h}_t^{\text{A}}(X, Z) &= 2t(1 + X + \beta^\top Z), \end{aligned}$$

where the coefficient $\beta = (\beta_1, \dots, \beta_d)$ is sampled once and independently for each dataset, with $(\beta_1, \dots, \beta_4) \sim \text{Dirichlet}(\mathbf{1}_4)$ and $\beta_i = 0$ for $5 \leq i \leq d$.

(\mathcal{S}_{par}) A partially additive model with

$$\begin{aligned} Z &= (Z_1, \dots, Z_d) \sim \text{N}(0, 1)^{\otimes d}, \\ X | Z &\sim \text{Unif}(\varphi(Z_1), \varphi(Z_1) + 1), \\ \mathbf{h}_t^{\text{pA}}(X, Z) &= 2t(1 + X + \varphi(Z_1)), \end{aligned}$$

where $\varphi(t) = \exp(-2t^2)$ is a scaled Gaussian density.

Note that both settings adhere to the partially additive model (4.2.9), and in this case the ACM coincides with the cumulative direct effect $\Theta_t = \int_0^t 2s ds = t$. The uncensored survival time T^* can in this case be simulated by

$$T^* := \frac{\sqrt{E}}{\sqrt{1 + X + f(Z)}},$$

where $E \perp\!\!\!\perp X, Z$ is standard exponentially distributed, and where $f(z)$ equals $\beta^\top z$ for the setting \mathcal{S}_{lin} , and equals $\varphi(z_1)$ for the setting \mathcal{S}_{par} .

The simulations were performed with covariate dimensions $d \in \{4, 16\}$ and datasets of sample sizes $n \in \{200, 600, 1800\}$. For each setting $N = 500$ datasets were simulated.

4.5.1 Estimators and implementation

The simulation study was implemented in Python. As a benchmark, an estimator of the exposure coefficient in the Aalen additive model was considered, which was implemented using `AalenAdditiveFitter` in the *lifelines* package [Davidson-Pilon, 2019]. This estimator is referred to as ‘Aalen estimator’ and denoted by $\hat{A} = (\hat{A}_t)$.

Estimation of the ACM requires methods for estimating the predictable projection and the conditional hazard. For estimation of π_t we implemented the following procedure.

1. Instantiate a grid $\mathcal{T} = (\frac{k}{n_{\mathcal{T}}-1} : k = 0, 1, \dots, n_{\mathcal{T}} - 1)$. Simulations were conducted with $n_{\mathcal{T}} = 20$.
2. Let $[n] = \mathcal{I}_1 \cup \mathcal{I}_2$ be a given sample-split, and let $x_{it} := (t, Y_t^{(i)}, Z_{1t}^{(i)}, \dots, Z_{dt}^{(i)})$, where $Z_{jt}^{(i)}$ denote the j -th covariate for the i -th individual (independent of t in this setting). Then format:
 - covariate matrices $\mathbb{X}_{\text{train}} := (x_{it})_{(i,t) \in \mathcal{I}_1 \times \mathcal{T}}$ and $\mathbb{X}_{\text{test}} := (x_{it})_{(i,t) \in \mathcal{I}_2 \times \mathcal{T}}$ of dimensions $(|\mathcal{I}_1| \cdot n_{\mathcal{T}}) \times (d + 2)$ and $(|\mathcal{I}_2| \cdot n_{\mathcal{T}}) \times (d + 2)$, respectively.
 - a response vector $\mathbb{Y}_{\text{train}} := (X_t^{(i)})_{(t,i) \in \mathcal{T} \times \mathcal{I}_1}$ of length $n_{\mathcal{T}} \cdot |\mathcal{I}_1|$.
3. Regress $\mathbb{Y}_{\text{train}}$ onto $\mathbb{X}_{\text{train}}$ using any regression method. For this we considered OLS and *gradient boosting* [Friedman, 2001, Hastie et al., 2009], which are denoted by a subscript `lin` and `gb`, respectively. These were implemented using classes `LinearRegression` and `GradientBoostingRegressor` from the Scikit-learn package [Pedregosa et al., 2011].
4. Use the fitted methods to predict response values for \mathbb{X}_{test} to produce estimates of the predictable projection over the grid \mathcal{T} for each test observation.
5. Interpolate to general time points $t \in [0, 1]$ using the estimated value at time $t_{\leq} = \max\{s \in \mathcal{T} \mid s \leq t\}$.

It is also possible to subset matrices in step (2) to pairs (i, t) such that $(T^{(i)} > t)$. This corresponds to directly targeting the quantity $\pi_t(Z) = \mathbb{E}[X \mid Y_t = 1, Z]$, rather than

targeting $\mathbb{E}[X | Y_t, Z_t]$ as done in the steps (3) and (4). Both methods were tried initially, with no noteworthy differences in the analysis result observed. For the interpolation in step (5), it is important that the method preserves predictability, which would not be ensured by, e.g., *linear interpolation*.

For conditional hazard estimation, we considered two off-the-shelf implementations: `GradientBoostingSurvivalAnalysis` from the `scikit-survival` module [Pölsterl, 2020] and the `AalenAdditiveFitter` from the *lifelines* package [Davidson-Pilon, 2019]. The two methods are again indicated by a subscript `gb` and `lin`, respectively.

The estimate $\hat{\rho}$ was computed with the clipping value $c_{\hat{\rho}} = 0.005$. However, the estimate was set to ‘nan’ rather than $c_{\hat{\rho}}$ whenever $\tilde{\rho}(t) < c_{\hat{\rho}}$. This is similar to how the `AalenAdditiveFitter` returns `nan` for times t such that $\sum_i Y_t^{(i)} < 3d$.

The x-ACM was implemented with 4-fold cross-fitting according to Algorithm 6. In addition, a version of the ACM estimator without any sample splitting was implemented, and this estimator is referred to as the N-ACM. Thus, 4 different estimators of the ACM were implemented: N-ACM_{lin}, X-ACM_{lin}, N-ACM_{gb}, and X-ACM_{gb}.

4.5.2 Results

Figure 4.5.1 shows the scaled root mean square error RMSE for various estimators with respect to the cumulative direct effect at time $t = 67/127 \approx 0.622$. This timepoint was the largest, for which all estimators yielded well-defined estimates for all the simulated datasets (the Aalen estimator yielded `nan` at time $t = 68/127$ for a dataset in the additive setting with $n = 200$ and $d = 16$). We observed initially that the ACM estimators based on linear regression methods, namely N-ACM_{lin} and X-ACM_{lin}, behaved similarly to the Aalen estimate \hat{A}_t . In view of this fact, and since our focus is on estimation of the ACM using flexible learning methods, the results for linear methods are omitted for more clear visualization.

In the left column of Figure 4.5.1, corresponding to the setting \mathcal{S}_{lin} , we observe that the Aalen estimator and X-ACM_{gb} have similar performance and appear to be \sqrt{n} -consistent. It is expected that both these estimators are \sqrt{n} -consistent, but it is perhaps surprising that the X-ACM_{gb} is as efficient as the Aalen estimator in the linear setting. The small discrepancy might be due to numerical approximations when fitting the estimators to the same common grid, or difference in bias-variance tradeoff because of (implicit) penalization.

The estimator N-ACM_{gb} without sample splitting performs significantly worse, although the discrepancy seems to diminish for larger sample sizes. In theory, sample-splitting is used to control the empirical process terms⁴, but these terms can also be handled by Donsker conditions or algorithmic stability, see for example the discussions in Chernozhukov et al. [2018], Chen et al. [2022]. This indicates, that such Donsker conditions are violated in our setup, and that the gradient boosting methods are not sufficiently stable. In practice, overfitting can lead to biases that would otherwise not

⁴In our analysis, these are the terms denoted by $R^{(1)}, R^{(2)}, R^{(4)}, r_1, \tilde{R}^{(2)}$, and $\tilde{R}^{(4)}$ found in the supplementary Section 4.A.

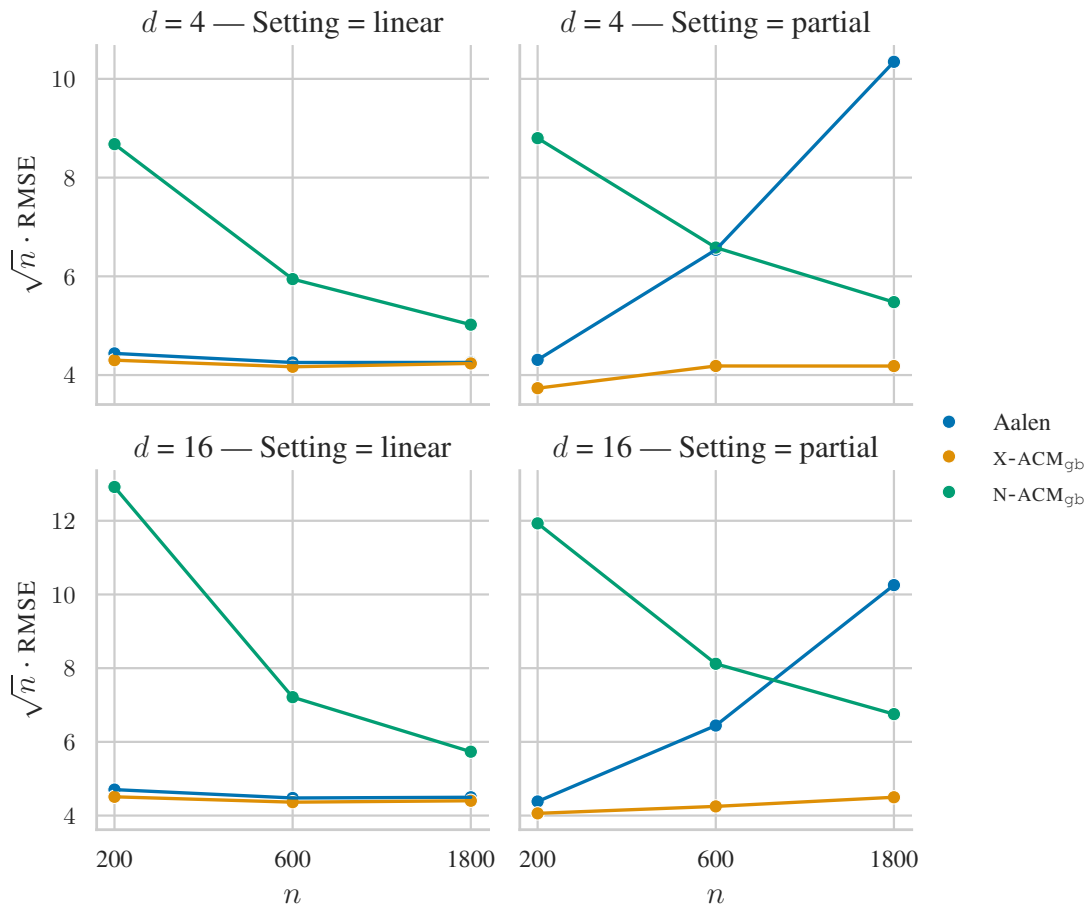


Figure 4.5.1: Scaled RMSE for various estimators with respect to the cumulative direct effect.

appear with sample-splitting [Chernozhukov et al., 2018]. Since the gradient boosted methods are fitted with the same hyper parameters, and these were tuned towards the larger sample sizes, it is plausible that they overfit in the smaller sample sizes, resulting in a systematic bias.

We observe that each of the ACM estimators perform similarly across the linear setting \mathcal{S}_{lin} and in the partial setting \mathcal{S}_{par} . This is different for the Aalen estimator \hat{A}_t , which fails severely in estimating the effect Θ_t . This is to be expected, as the Aalen additive model is misspecified, and as a consequence, \hat{A}_t cannot reasonable account for the nonlinear effect $\varphi(Z_1)$.

In summary, we conclude that X-ACM_{gb} performs well across all settings, indicating that it provides an efficient and robust estimator of the cumulative direct effect.

4.6 Discussion

The ACM was proposed as a model-free estimand to quantify the conditional relationship between a time-to-event and an exposure, controlling for given covariates. While it is a specific instance of the more general LCM introduced by Christgau et al. [2023b], we have highlighted several unique properties and results that apply to the ACM – and not necessarily to the LCM – which were not considered in Christgau et al. [2023b].

The ACM may be understood as a weighted hazard difference estimand, and we showed in Proposition 4.2.4 that it coincides with the cumulative exposure coefficient in the partially additive hazards model (4.2.9) that best approximates the true hazard in L^2 . As a result, the ACM measures the cumulative direct effect within partially additive hazards model, and in particular within the Aalen additive model.

The asymptotic analysis for our proposed estimator, the X-ACM, was to some extent based on results from Christgau et al. [2023b]. However, a novel examination of the term B in (4.4.22) was necessary to account for the estimation error in the function ρ appearing in the denominator. Notably, the remainder term $D^{(2)}$ from Christgau et al. [2023b], analogous to $R^{(4)}$ in (4.A.3), does not vanish for the ACM when ρ is estimated. Moreover, the decomposition in (2.4.17) differs from typical decompositions of estimators related to the *partially linear model*, considered in, e.g., [Robinson, 1988, Lundborg and Pfister, 2023, Hines et al., 2023], because these decompositions are not feasible when integrating over time.

Estimating ρ using the training data. Suppose that the function ρ is estimated using the training data, which corresponds to replacing \mathcal{I}_2 with \mathcal{I}_1 in (4.3.16). Although the residual estimates would then fit within the framework of Christgau et al. [2023b], the analysis does not simplify as the term $D^{(2)}$ in Christgau et al. [2023b] still remains non-vanishing. Nevertheless, the analysis in Section 4.4.2 can be adapted to derive an oracle decomposition similar to (4.4.25), but with the integrals $\int_0^t E_s^{(i)} G_s^{(i)} d\gamma_s$ summed over $i \in \mathcal{I}_1$. This makes it unfeasible to aggregate the cross-fitted estimates, resulting in a smaller effective sample when using this approach. Additionally, controlling the estimation error of ρ in this case requires Donsker class conditions or algorithmic stability, cf. the term r_1 in the proof of Proposition 4.4.4. Our simulations in Section 4.5 suggest

that such conditions may not hold in practice, at least when gradient boosting is used.

The cross-fitting scheme. The x-ACM, described in Section 4.3.1, employs cross-fitting similarly to the DML1 estimator described in Chernozhukov et al. [2018]. While it is conceivable that a cross-fitting scheme corresponding to DML2 – where numerators and denominators are aggregated separately across sample splits – could be implemented, the integration over time complicates this approach. If possible, however, such an approach might lead to increased numerical stability in small samples, cf. Remark 3.1 in Chernozhukov et al. [2018].

Relaxations of assumptions. We already discussed, in Remark 4.4.3, the possibility of relaxing the rate requirements to hold for the empirical errors instead of the expected errors. The assumptions on boundedness, Assumptions 4.4.1 (i) and (ii), can be found in similar works, cf. Theorem 2 in Vansteelandt et al. [2022] or Assumption 1 in Hou et al. [2023]. It is, however, likely that they can be replaced by suitable tail bounds or bounds on conditional variances, but at the cost of more technical proofs. Finally, Assumption 4.4.3 states that Π can be estimated at an order $o(n^{-1/4})$ in terms of the 4-norm. While the order $o(n^{-1/4})$ is to be expected, see [Chernozhukov et al., 2018], it is conceivable that the condition can be relaxed to hold for the 2-norm rather than the slightly more restrictive 4-norm.

Alternative estimands. A limitation of the ACM is the inverse weighting by the function ρ , potentially making it difficult to estimate accurately. An interesting direction for future work is to consider alternative ways of weighting the hazard difference. This may lead to estimands that are easier to estimate, but still interpretable. Such a compromise was considered by Vansteelandt et al. [2022], who propose a hazard ratio estimand that deliberately avoids inverse weighting by the conditional density of the exposure. However, their estimand still requires inverse weighting by the cumulative hazard, which imposes a limitation similar to the inverse weighting with ρ .

The related works Dukes et al. [2019], Hou et al. [2023], cf. Table 4.1.1, consider estimation of Θ_t in (4.2.9), with time-independent exposure and covariates, based on (efficient) orthogonal score methodology. It would be interesting to investigate how the x-ACM performs relatively to their methodology, both for a well-specified and misspecified model. Furthermore, supposing that their methodologies can be shown to target general model-free estimands, it would be interesting to compare these estimands with the ACM.

Acknowledgments

We are grateful to Anton Rask Lundborg for helpful discussions. AMC and NRH were supported by a research grant (NNF20OC0062897) from Novo Nordisk Fonden.

Supplement to ‘Assumption-lean Aalen regression’

The supplementary material consists of Section 4.A, which contains proofs of the main results and related auxiliary lemmas, and Section 4.B, which includes a discussion of the ACM in the context of multiplicative hazards models.

4.A. Auxillary results and proofs

4.A.1 Additional notation for proofs

While the main manuscript is formulated using histories and the conditional hazards h_t and \mathbf{h}_t , it will be convenient to work with filtrations and intensities in the proofs, as also discussed in Remark 4.2.3. Recall that we define the filtrations $\mathcal{F}_t := \sigma(Z_s, N_s; s \leq t)$ and $\mathcal{G}_t := \sigma(X_s, Z_s, N_s; s \leq t)$. The stochastic processes $\lambda = (\lambda_t)$ and $\boldsymbol{\lambda} = (\boldsymbol{\lambda}_t)$, defined by

$$\lambda_t := Y_t h_t(\bar{Z}_t) \quad \text{and} \quad \boldsymbol{\lambda}_t := Y_t \mathbf{h}_t(\bar{X}_t, \bar{Z}_t),$$

are then, under independent right-censoring as described in Section 4.2.4, the \mathcal{F}_t - and \mathcal{G}_t -intensity, respectively, of N_t . We remark that, by definition, this means that $M_t = N_t - \int_0^t \lambda_s ds$ is an \mathcal{F}_t -martingale and that

$$M_t := N_t - \int_0^t \lambda_s ds$$

is a \mathcal{G}_t -martingale. By the innovation theorem, it also holds that $\mathbb{E}[\boldsymbol{\lambda}_t | \mathcal{F}_{t-}] = \lambda_t$. This notation is consistent with Christgau et al. [2023b].

For the proofs in the asymptotic analysis we will additionally use the following notation:

$$\begin{aligned} \hat{\lambda}_t^{(i)} &:= Y_t^{(i)} \hat{h}_t(\bar{Z}^{(i)}) \\ \hat{\Pi}_t^{(i)} &:= \hat{\pi}_t(\bar{Z}^{(i)}) \end{aligned}$$

for $i \in [n]$ and $t \in [0, 1]$. Finally, we let $\|\cdot\|_p$ denote the norm on $L^p([0, 1] \times \Omega, dt \otimes \mathbb{P})$, so that, for example,

$$b_n = \int_0^1 \mathbb{E}[(\hat{\lambda}_t - \lambda_t)^2] dt = \|\hat{\lambda} - \lambda\|_2^2.$$

4.A.2 Proof of Proposition 4.2.4

We first consider the minimization objective for a fixed timepoint $t \in [0, 1]$.

Lemma 4.A.1. *For fixed $t \in [0, 1]$, it holds that*

$$\begin{cases} \vartheta_t^* = \text{Cov}(G_t, \boldsymbol{\lambda}_t - \lambda_t) \\ g_t^* = h_t - \text{Cov}(G_t, \boldsymbol{\lambda}_t - \lambda_t) \pi_t \end{cases} \quad (4.A.1)$$

are solutions to

$$\underset{\vartheta_t \in \mathbb{R}, g_t \in L^0(P_{\bar{Z}_t})}{\text{minimize}} \quad \mathbb{E} \left[Y_t (\mathbf{h}_t(\bar{X}_t, \bar{Z}_t) - (\vartheta_t X_t + g_t(\bar{Z}_t)))^2 \right]$$

where $P_{\bar{Z}_t}$ denotes the distribution of \bar{Z}_t . Moreover, the solution ϑ_t^* is unique.

Proof. Recall that on the event $(Y_t = 1) = (T \geq t)$, the predictable projection Π_t takes the value $\pi_t(\bar{Z}_t) = \mathbb{E}[X_t | T \geq t, \bar{Z}_t]$. Since $\pi_t, h_t \in L^0(P_{\bar{Z}_t})$, the minimization problem is equivalent to minimizing

$$\begin{aligned} u(\vartheta_t, \ell_t) &:= \mathbb{E} \left[Y_t (\mathbf{h}_t(\bar{X}_t, \bar{Z}_t) - h_t(\bar{Z}_t) - \vartheta_t (X_t - \pi_t(\bar{Z}_t)) - \ell_t(\bar{Z}_t))^2 \right] \\ &= \mathbb{E} \left[(\boldsymbol{\lambda}_t - \lambda_t - \vartheta_t Y_t (X_t - \Pi_t) - Y_t \ell_t(\bar{Z}_t))^2 \right], \end{aligned}$$

which corresponds to the substitution $g_t = \ell_t + h_t - \vartheta_t \pi_t$. Note that

$$\mathbb{E}[\boldsymbol{\lambda}_t - \lambda_t | \mathcal{F}_{t-}] = 0 = \mathbb{E}[X_t - \Pi_t | \mathcal{F}_{t-}].$$

Now, expanding the square in $u(\vartheta_t, \ell_t)$ and using that $\mathbb{E}[\ell_t(\bar{Z}_t)^2] \geq 0$, we obtain that

$$u(\vartheta_t, \ell_t) \geq \mathbb{E}[(\boldsymbol{\lambda}_t - \lambda_t)^2] + \vartheta_t^2 \mathbb{E}[Y_t (X_t - \Pi_t)^2] - 2\vartheta_t \mathbb{E}[(\boldsymbol{\lambda}_t - \lambda_t)(X_t - \Pi_t)].$$

Equality is attained when $\ell_t = 0$, which is equivalent to $g_t = h_t - \vartheta_t \pi_t \in L^0(P_{\bar{Z}_t})$. The right-hand side is a quadratic in ϑ_t and its unique minimizer is given by

$$\frac{\mathbb{E}[(\boldsymbol{\lambda}_t - \lambda_t)(X_t - \Pi_t)]}{\mathbb{E}[Y_t (X_t - \Pi_t)^2]} = \text{Cov}(G_t, \boldsymbol{\lambda}_t - \lambda_t).$$

This shows that (ϑ_t^*, g_t^*) in (4.A.1) are indeed minimizers and that ϑ^* is unique. \square

Returning to the proof of Proposition 4.2.4, we let (ϑ_t^*, g_t^*) be the pointwise minimizers from (4.A.1) for each $t \in [0, 1]$. For any (ϑ, g) as in Proposition 4.2.4, it follows directly from Tonelli's theorem and Lemma 4.A.1 that

$$\begin{aligned} \text{minimization objective} &= \int_0^1 \mathbb{E} \left[Y_t (\mathbf{h}_t(\bar{X}_t, \bar{Z}_t) - \vartheta_t X_t - g_t(\bar{Z}_t))^2 \right] dt \\ &\geq \int_0^1 \mathbb{E} \left[Y_t (\mathbf{h}_t(\bar{X}_t, \bar{Z}_t) - \vartheta_t^* X_t - g_t^*(\bar{Z}_t))^2 \right] dt. \end{aligned} \quad (4.A.2)$$

We conclude that (ϑ^*, g^*) is indeed a minimizer to (4.2.14), since ϑ_t is measurable and $g_t^* \in L^0(P_{\bar{Z}_t})$ for each $t \in [0, 1]$. Note that equality occurs in (4.A.2) if and only if

$$\mathbb{E} \left[Y_t (\mathbf{h}_t(\bar{X}_t, \bar{Z}_t) - \vartheta_t X_t - g_t(\bar{Z}_t))^2 \right] = \mathbb{E} \left[Y_t (\mathbf{h}_t(\bar{X}_t, \bar{Z}_t) - \vartheta_t^* X_t - g_t^*(\bar{Z}_t))^2 \right]$$

for almost all $t \in [0, 1]$. By the uniqueness of ϑ^* in Lemma 4.A.1, this is equivalent to $\vartheta_t = \vartheta_t^*$ for almost all $t \in [0, 1]$. The last part of the proposition now follows since $\gamma_t = \int_0^t \vartheta_s^* ds$, according to (4.2.12). \square

4.A.3 Proof of Proposition 4.4.2

We consider a decomposition similar to that of Christgau et al. [2023b], but with their non-vanishing processes combined. To this end, define the ρ -oracle residual estimates

$$\tilde{G}_t := \frac{\hat{E}_t}{\rho(t)} = \frac{Y_t(X_t - \hat{\Pi}_t)}{\rho(t)}, \quad t \in [0, 1].$$

Then note that

$$\sqrt{n_2} \cdot \tilde{\gamma} = U + R^{(1)} + R^{(2)} + R^{(3)} + R^{(4)}, \quad (4.A.3)$$

where the processes U , $R^{(1)}$, $R^{(2)}$, $R^{(3)}$, and $R^{(4)}$ are given by

$$\begin{aligned} U_t &= \frac{1}{\sqrt{n_2}} \sum_{i \in \mathcal{I}_2} \int_0^t G_s^{(i)} dM_s^{(i)}, \\ R_t^{(1)} &= \frac{1}{\sqrt{n_2}} \sum_{i \in \mathcal{I}_2} \int_0^t G_s^{(i)} (\lambda_s^{(i)} - \hat{\lambda}_s^{(i)}) ds, \\ R_t^{(2)} &= \frac{1}{\sqrt{n_2}} \sum_{i \in \mathcal{I}_2} \int_0^t (\tilde{G}_s^{(i)} - G_s^{(i)}) dM_s^{(i)}, \\ R_t^{(3)} &= \frac{1}{\sqrt{n_2}} \sum_{i \in \mathcal{I}_2} \int_0^t (\tilde{G}_s^{(i)} - G_s^{(i)}) (\lambda_s^{(i)} - \hat{\lambda}_s^{(i)}) ds, \\ R_t^{(4)} &= \frac{1}{\sqrt{n_2}} \sum_{i \in \mathcal{I}_2} \int_0^t (\tilde{G}_s^{(i)} - G_s^{(i)}) (\lambda_s^{(i)} - \lambda_s^{(i)}) ds. \end{aligned}$$

In view of this decomposition, it suffices to show that

$$\sup_{t \in [0, 1]} |R_t^{(\ell)}| \xrightarrow{P} 0$$

as $n \rightarrow \infty$ for $\ell = 1, \dots, 4$. To this end, note that Assumption 4.4.1 ensures that Assumption 4.1 in Christgau et al. [2023b] holds with:

$$\max\{|G_t|, |\tilde{G}_t|\} \leq c_\rho^{-1} C_E, \quad \text{and} \quad \max\{\lambda_t, \hat{\lambda}_t\} \leq C_h.$$

Assumption 4.4.2 can be translated directly into Assumption 4.2 in Christgau et al. [2023b]. We therefore conclude that $R^{(1)}$, $R^{(2)}$, and $R^{(3)}$ converge (uniformly) to zero in probability by Proposition 4.4 in Christgau et al. [2023b]. For the last remainder process we have that

$$\sup_{t \in [0, 1]} |R_t^{(4)}| \leq c_\rho^{-1} \sup_{t \in [0, 1]} \left| \frac{1}{\sqrt{n_2}} \sum_{i \in \mathcal{I}_2} \int_0^t (\hat{E}_s^{(i)} - E_s^{(i)}) (\lambda_s^{(i)} - \lambda_s^{(i)}) ds \right| \xrightarrow{P} 0,$$

where the convergence in probability is established in Lemma A.10 in Christgau et al. [2023b]. \square

4.A.4 Proof of Proposition 4.4.4

For n sufficiently large it holds that $c_{\hat{\rho}} < c_{\rho}$, in which case Assumption 4.4.1(iii) implies that $|\hat{\rho}(t) - \rho(t)| \leq |\tilde{\rho}(t) - \rho(t)|$ for all $t \in [0, 1]$. Hence it suffices to establish the bound for the empirical estimator $\tilde{\rho}$ from (4.3.16). Note also that

$$\tilde{\rho}(t) = \frac{1}{n_2} \sum_{i \in \mathcal{I}_2} (\hat{E}_t^{(i)})^2,$$

and consider the following decomposition

$$\sqrt{n_2} \cdot (\tilde{\rho}(t) - \rho(t)) = \sqrt{n_2} \cdot (\bar{\rho} - \rho) + 2r_1(t) + r_2(t)$$

where

$$\begin{aligned} \bar{\rho}(t) &:= \frac{1}{n_2} \sum_{i \in \mathcal{I}_2} (E_t^{(i)})^2, \\ r_1(t) &:= \frac{1}{\sqrt{n_2}} \sum_{i \in \mathcal{I}_2} E_t^{(i)} (\hat{E}_t^{(i)} - E_t^{(i)}), \\ r_2(t) &:= \frac{1}{\sqrt{n_2}} \sum_{i \in \mathcal{I}_2} (\hat{E}_t^{(i)} - E_t^{(i)})^2. \end{aligned}$$

We first note that the terms in r_1 are conditionally i.i.d. given $\mathcal{D}_1 := \sigma(\mathbf{W}^{(i)} : i \in \mathcal{I}_1)$ with conditional mean zero. It follows that

$$\begin{aligned} \mathbb{E}[r_1(t)^2] &= \mathbb{E}[\text{Var}(r_1(t) | \mathcal{D}_1)] \\ &= \mathbb{E}\left[\frac{1}{n_2} \sum_{i \in \mathcal{I}_2} \text{Var}(E_t^{(i)} (\hat{E}_t^{(i)} - E_t^{(i)}) | \mathcal{D}_1)\right] \\ &= \mathbb{E}[(E_t^{(i)})^2 (\hat{E}_t^{(i)} - E_t^{(i)})^2] \leq C_E^2 \mathbb{E}[(\Pi_t - \hat{\Pi}_t)^2], \end{aligned}$$

where we have used that $\hat{E}_t^{(i)} - E_t^{(i)} = Y_t(X_t^{(i)} - \hat{\Pi}_t^{(i)} - X_t^{(i)} + \Pi_t^{(i)}) = Y_t(\Pi_t^{(i)} - \hat{\Pi}_t^{(i)})$. By integrating over time we obtain

$$\|r_1\|_2 \leq C_E \|\hat{\Pi} - \Pi\|_2 \rightarrow 0. \quad (4.A.4)$$

For the other remainder, Cauchy-Schwarz yields

$$\mathbb{E}[r_2(t)^2] = \frac{1}{n_2} \sum_{i, j \in \mathcal{I}_2} \mathbb{E}[(\hat{E}_t^{(i)} - E_t^{(i)})^2 (\hat{E}_t^{(j)} - E_t^{(j)})^2] \leq n_2 \mathbb{E}[(\hat{E}_t - E_t)^4].$$

By Assumption 4.4.3 we conclude that

$$\|r_2\|_2^2 \leq n_2 \|\hat{E} - E\|_4^4 = n_2 \|\hat{\Pi} - \Pi\|_4^4 \rightarrow 0, \quad n \rightarrow \infty. \quad (4.A.5)$$

4 Assumption-lean Aalen regression

For the oracle sum, $\bar{\rho}$, we note that $\mathbb{E}[\bar{\rho}(t)] = \rho(t)$, and that $\bar{\rho}(t)$ is a sum of i.i.d. random variables with the same distribution as $E_t^2 \in [0, C_E^2]$. Invoking Popoviciu's inequality on variances [Popoviciu, 1935], we obtain

$$\int_0^1 \mathbb{E}[(\bar{\rho}(t) - \rho(t))^2] dt = \int_0^1 \text{Var}(\bar{\rho}(t)) dt = \frac{1}{n_2} \int_0^1 \text{Var}(E_t^2) dt \leq \frac{C_E^4}{4n_2}. \quad (4.A.6)$$

Combining (4.A.4), (4.A.5), and (4.A.6) via Minkowski's inequality we obtain (4.4.24). \square

4.A.5 Proof of Theorem 4.4.5

We first note that

$$\frac{1}{\hat{\rho}(s)} - \frac{1}{\rho(s)} = \frac{\rho(s) - \hat{\rho}(s)}{\rho(s)\hat{\rho}(s)} = \underbrace{\frac{\rho(s) - \hat{\rho}(s)}{\rho(s)^2}}_{=:\zeta_s} + \underbrace{\frac{(\rho(s) - \hat{\rho}(s))^2}{\hat{\rho}(s)\rho(s)^2}}_{=:\xi_s}.$$

Now consider a decomposition of B similar to that of $\sqrt{n_2} \cdot \tilde{\gamma}$,

$$B = \tilde{B}_t + \sum_{\ell=1}^5 \tilde{R}^{(\ell)}$$

where

$$\begin{aligned} \tilde{B}_t &= \frac{1}{\sqrt{n_2}} \sum_{i \in \mathcal{I}_2} \int_0^t \zeta_s E_s^{(i)} (\lambda_s^{(i)} - \hat{\lambda}_s^{(i)}) ds, \\ \tilde{R}_t^{(1)} &= \frac{1}{\sqrt{n_2}} \sum_{i \in \mathcal{I}_2} \int_0^t \zeta_s (\hat{E}_s^{(i)} - E_s^{(i)}) (\lambda_s^{(i)} - \hat{\lambda}_s^{(i)}) ds \\ \tilde{R}_t^{(2)} &= \frac{1}{\sqrt{n_2}} \sum_{i \in \mathcal{I}_2} \int_0^t \zeta_s \hat{E}_s^{(i)} dM_s^{(i)}, \\ \tilde{R}_t^{(3)} &= \frac{1}{\sqrt{n_2}} \sum_{i \in \mathcal{I}_2} \int_0^t \zeta_s \hat{E}_s^{(i)} (\lambda_s^{(i)} - \hat{\lambda}_s^{(i)}) ds, \\ \tilde{R}_t^{(4)} &= \frac{1}{\sqrt{n_2}} \sum_{i \in \mathcal{I}_2} \int_0^t \xi_s \hat{E}_s^{(i)} dM_s^{(i)} \\ \tilde{R}_t^{(5)} &= \frac{1}{\sqrt{n_2}} \sum_{i \in \mathcal{I}_2} \int_0^t \xi_s \hat{E}_s^{(i)} (\lambda_s^{(i)} - \hat{\lambda}_s^{(i)}) ds \end{aligned}$$

The decomposition may seem peculiar at first glance: some of the remainder processes could be decomposed further into terms that would be more manageable. However, crude bounds on these error terms will suffice since ζ and ξ also decay at an order of $n_2^{-1/2}$. Theorem 4.4.5 follows from combining Proposition 4.A.2 and Proposition 4.A.3, stated below. \square

Proposition 4.A.2. *Under Assumptions 4.4.1, 4.4.2, and 4.4.3, it holds that*

$$\sup_{t \in [0,1]} |\tilde{R}_t^{(\ell)}| \xrightarrow{P} 0 \quad (4.A.7)$$

for $\ell \in \{1, 2, 3, 4, 5\}$.

Proof. Let $\langle \cdot, \cdot \rangle_2$ denote the inner product on $L^2([0, 1] \times \Omega, dt \otimes \mathbb{P})$. The triangle inequality and Cauchy-Schwarz can be used to bound the expectation of the first remainder term as follows:

$$\begin{aligned} \mathbb{E} \left[\sup_{t \in [0,1]} |\tilde{R}_t^{(1)}| \right] &\leq \frac{1}{\sqrt{n_2}} \sum_{i \in \mathcal{I}_2} \langle |\zeta|, |\hat{E}^{(i)} - E^{(i)}| \cdot |\boldsymbol{\lambda}^{(i)} - \lambda^{(i)}| \rangle_2 \\ &\leq \sqrt{n_2} \cdot \|\zeta\|_2 \cdot \|(\hat{E} - E) \cdot (\boldsymbol{\lambda} - \lambda)\|_2 \\ &\leq c_\rho^{-2} C_{\mathbf{h}} \sqrt{n_2} \cdot \|\hat{\rho} - \rho\|_2 \cdot \|\hat{E} - E\|_2, \end{aligned} \quad (4.A.8)$$

where we have used Assumption 4.4.1 for the last inequality. By a similar argument we obtain

$$\mathbb{E} \left[\sup_{t \in [0,1]} |\tilde{R}_t^{(3)}| \right] \leq c_\rho^{-2} C_E \sqrt{n_2} \cdot \|\hat{\rho} - \rho\|_2 \cdot \|\hat{\lambda} - \lambda\|_2. \quad (4.A.9)$$

Now, combining Assumption 4.4.2 with Proposition 4.4.4, we see that the right-hand sides of (4.A.8) and (4.A.9) vanish as $n \rightarrow \infty$. Hence we conclude that the desired convergence, (4.A.7), holds for $\ell = 1$ and $\ell = 3$.

For the second and fourth remainder, let $\mathcal{D}_1 := \sigma(\mathbf{W}_i : i \in \mathcal{I}_1)$ denote the information contained in the training data. Let further $\mathbf{G}_t = \mathcal{D}_1 \vee \sigma(\mathcal{G}_t^{(i)} : i \in \mathcal{I}_2)$ denote the filtration containing the histories of the processes $(N^{(i)}, X^{(i)}, Z^{(i)})_{i \in \mathcal{I}_2}$ together with the information \mathcal{D}_1 . Then note that:

- since the observations are independent, the compensated $\mathcal{G}_t^{(i)}$ -martingale $\mathbf{M}_t^{(i)}$ is in fact also a \mathbf{G}_t -martingale for each $i \in \mathcal{I}_2$.
- each term in the empirical estimator $\tilde{\rho}$, defined in (4.3.16), is \mathbf{G}_t -predictable, and hence $\hat{\rho}(t) = \max\{c_{\tilde{\rho}}, \tilde{\rho}(t)\}$ is \mathbf{G}_t -predictable. As a consequence, both ζ_t and ξ_t are also \mathbf{G}_t -predictable.
- Assumption 4.4.1 implies that $|\zeta_t| \leq c_\rho^{-2} C_E^2$ and that $|\xi_t| \leq c_\rho^{-1} c_\rho^{-2} C_E^4$. Hence both $\int_0^1 \zeta_t^2 \boldsymbol{\lambda}_t dt$ and $\int_0^1 \xi_t^2 \boldsymbol{\lambda}_t dt$ have finite expectations.

We conclude that the terms in $\tilde{R}^{(2)}$ are mean-zero \mathbf{G}_t -martingales, cf. Lemma A.1 in Christgau et al. [2023b], and hence $\tilde{R}^{(2)}$ is also a \mathbf{G}_t -martingale. By the same argument, $\tilde{R}^{(4)}$ is a \mathbf{G}_t -martingale. Therefore we can apply Doob's submartingale inequality, which yields that

$$\mathbb{P} \left(\sup_{t \in [0,1]} |\tilde{R}_t^{(\ell)}| > \varepsilon \right) \leq \frac{1}{\varepsilon^2} \mathbb{E} [(\tilde{R}_1^{(\ell)})^2] \quad (4.A.10)$$

4 Assumption-lean Aalen regression

for $\ell \in \{2, 4\}$. Now using the calculus of the quadratic characteristic for compensated counting processes (see, for example, Proposition II.4.1 in [Andersen et al., 1993]) we obtain

$$\begin{aligned} \mathbb{E}\left[(\tilde{R}_1^{(2)})^2\right] &= \mathbb{E}\left[\langle \tilde{R}_1^{(2)} \rangle\right] = \mathbb{E}\left[\int_0^1 \zeta_t^2 \hat{E}_t^2 \lambda_t dt\right] \\ &= C_h C_E^2 \int_0^1 \frac{(\rho(t) - \hat{\rho}(t))^2}{\rho(t)^4} dt \leq c_\rho^{-4} C_h C_E^2 \|\rho - \hat{\rho}\|_2^2 \rightarrow 0. \end{aligned}$$

Similarly, observe that

$$\begin{aligned} \mathbb{E}\left[(\tilde{R}_1^{(4)})^2\right] &= \mathbb{E}\left[\langle \tilde{R}_1^{(4)} \rangle\right] = \mathbb{E}\left[\int_0^1 \xi_t^2 \hat{E}_t^2 \lambda_t dt\right] \\ &\leq C_h C_E^2 \int_0^1 \frac{(\rho(t) - \hat{\rho}(t))^4}{\hat{\rho}(t)^2 \rho(t)^4} dt \leq c_\rho^{-2} c_\rho^{-4} C_h C_E^6 \|\rho - \hat{\rho}\|_2^2 \rightarrow 0, \end{aligned}$$

where the convergence follows from Proposition 4.4.4 and the fact that $c_\rho^{-2} = o(n)$. Combined with (4.A.10), we conclude that (4.A.7) holds for $\ell = 2$ and $\ell = 4$.

For the final term, a direct bound yields

$$\mathbb{E}\left[\sup_{t \in [0,1]} |\tilde{R}_t^{(5)}|\right] \leq C_E C_h \sqrt{n_2} \|\xi\|_1 \leq c_\rho^{-2} c_\rho^{-1} C_E C_h \sqrt{n_2} \|\rho - \hat{\rho}\|_2^2 \rightarrow 0.$$

Thus we have established (4.A.7) for $\ell \in \{1, 2, 3, 4, 5\}$. \square

Proposition 4.A.3. *Under Assumptions 4.4.1 and 4.4.3, it holds that*

$$\sup_{t \in [0,1]} |\tilde{B}_t - V_t| \xrightarrow{P} 0,$$

where $V = (V_t)$ is the process given by

$$V_t = \frac{1}{\sqrt{n_2}} \sum_{i \in \mathcal{I}_2} \int_0^t \frac{(\rho(s) - (E_s^{(i)})^2)}{\rho(s)^2} \mathbb{E}[E_s(\lambda_s - \lambda_s)] ds, \quad t \in [0, 1].$$

Proof. The results follows from combining Lemmas 4.A.4 and 4.A.5. \square

Lemma 4.A.4. *Under Assumptions 4.4.1 and 4.4.3, it holds that $\sup_{t \in [0,1]} |\tilde{B}_t - \hat{V}_t| \xrightarrow{P} 0$, where $\hat{V} = (\hat{V}_t)$ is the process given by*

$$\hat{V}_t = \sqrt{n_2} \int_0^t \zeta_s \mathbb{E}[E_s(\lambda_s - \lambda_s)] ds, \quad t \in [0, 1].$$

Proof. Letting $\|\cdot\|_{[0,1]}$ denote the norm on $L^2([0,1], dt)$, an application of Cauchy-Schwarz yields that

$$\begin{aligned} \sup_t |\tilde{B}_t - \hat{V}_t| &= \sqrt{n_2} \int_0^1 |\zeta_s| \left| \frac{1}{n_2} \sum_{i \in \mathcal{I}_2} E_s^{(i)} (\lambda_s^{(i)} - \lambda_s^{(i)}) - \mathbb{E}[E_s(\lambda_s - \lambda_s)] \right| ds \\ &= \left\| \frac{1}{n_2} \sum_{i \in \mathcal{I}_2} E^{(i)} (\lambda^{(i)} - \lambda^{(i)}) - \mathbb{E}[E(\lambda - \lambda)] \right\|_{[0,1]} \cdot \sqrt{n_2} \|\zeta\|_{[0,1]} \quad (4.A.11) \end{aligned}$$

Since $\mathbb{E}\|E(\boldsymbol{\lambda} - \lambda)\|_{[0,1]} \leq C_E C_h$, the law of large numbers – in the space $L^2([0, 1])$ – gives that the first factor of (4.A.11) converges to zero almost surely, and in particular also in probability. See for example Ledoux and Talagrand [1991, Cor. 7.10] for a general formulation of the law of large numbers on separable Banach spaces.

For the second factor, we note that

$$\sqrt{n_2}\|\zeta\|_{[0,1]} \leq c_\rho^{-2}\sqrt{n_2}\|\hat{\rho} - \rho\|_{[0,1]}.$$

The right-hand side is bounded in expectation by (4.4.24), and in particular it is bounded in probability. Combined we conclude that the right-hand side of (4.A.11) converges to zero in probability. \square

Lemma 4.A.5. *Under Assumptions 4.4.1 and 4.4.3, it holds that*

$$\mathbb{E}\left[\sup_{t \in [0,1]} |\hat{V}_t - V_t|\right] \xrightarrow{P} 0$$

Proof. Let $\tilde{V} = (\tilde{V}_t)$ be the process given by

$$\tilde{V}_t := \sqrt{n_2} \int_0^t \varsigma_s \mathbb{E}[E_s(\boldsymbol{\lambda}_s - \lambda_s)] ds, \quad \text{where } \varsigma_s := \frac{\rho(s) - \tilde{\rho}(s)}{\rho(s)^2}.$$

In other words, ς_s corresponds to ζ_s , but where the clipped estimate, $\hat{\rho}$, is replaced by the empirical estimate $\tilde{\rho}$ from (4.3.16). As such, we have equality of the events $(\varsigma_s \neq \zeta_s) = (\tilde{\rho}(s) < c_{\hat{\rho}})$. Let n be large enough such that $c_{\hat{\rho}} < c_\rho < \rho(s)$ for all $s \in [0, 1]$, and then we observe that

$$\begin{aligned} \sup_{t \in [0,1]} |\hat{V}_t - \tilde{V}_t| &\leq C_E C_h \sqrt{n_2} \int_0^1 |\zeta_s - \varsigma_s| ds \\ &\leq c_{\hat{\rho}} c_\rho^{-2} C_E C_h \sqrt{n_2} \int_0^1 \mathbb{1}(\tilde{\rho}(s) < c_{\hat{\rho}}) ds \\ &\leq c_{\hat{\rho}} c_\rho^{-2} C_E C_h \sqrt{n_2} \int_0^1 \mathbb{1}((\tilde{\rho}(s) - \rho(s))^2 < (\rho(s) - c_{\hat{\rho}})^2) ds. \end{aligned}$$

Applying Tonelli's theorem and using Markov's inequality to the above we obtain

$$\begin{aligned} \mathbb{E}\left[\sup_{t \in [0,1]} |\hat{V}_t - \tilde{V}_t|\right] &\lesssim \sqrt{n_2} \int_0^1 \mathbb{P}((\tilde{\rho}(s) - \rho(s))^2 < (\rho(s) - c_{\hat{\rho}})^2) ds \\ &\leq \sqrt{n_2} \int_0^1 \frac{\mathbb{E}[(\tilde{\rho}(s) - \rho(s))^2]}{(\rho(s) - c_{\hat{\rho}})^2} ds \\ &\leq (c_\rho - c_{\hat{\rho}})^{-2} \sqrt{n_2} \|\tilde{\rho} - \rho\|_2^2 \rightarrow 0, \end{aligned}$$

where ' \lesssim ' is an inequality up to the constant $c_{\hat{\rho}} c_\rho^{-2} C_E C_h$, and where the last convergence follows from Proposition 4.4.4.

4 Assumption-lean Aalen regression

It now suffices to show that $\mathbb{E}\left[\sup_{t \in [0,1]} |\tilde{V}_t - V_t|\right] \xrightarrow{P} 0$. Using the decomposition $\tilde{\rho} = \bar{\rho} + r_1 + r_2$ from the proof of Proposition 4.4.4, we first see that

$$V_t = \sqrt{n_2} \int_0^t \frac{\rho(s) - \bar{\rho}(s)}{\rho(s)^2} \mathbb{E}[E_s(\boldsymbol{\lambda}_s - \lambda_s)] ds,$$

so that

$$\sup_{t \in [0,1]} |\tilde{V}_t - V_t| \leq \frac{C_E C_{\mathbf{h}}}{c_{\rho}^2} \sqrt{n_2} \int_0^1 |r_1(s)| + |r_2(s)| ds.$$

Using the intermediate results (4.A.4) and (4.A.5) from the proof of Proposition 4.4.4, it follows that

$$\mathbb{E}\left[\sup_{t \in [0,1]} |\tilde{V}_t - V_t|\right] = \frac{C_E C_{\mathbf{h}}}{c_{\rho}^2} \sqrt{n_2} (\|r_1\|_1 + \|r_2\|_1) \rightarrow 0.$$

□

4.A.6 Proof of Theorem 4.4.6

We first establish asymptotic Gaussianity of the oracle process.

Theorem 4.A.6. *Let $S = (S_t)_{t \in [0,1]}$ be the process given by $S_t = U_t + V_t - \sqrt{n_2} \gamma_t$, and let $\Gamma = (\Gamma_t)_{t \in [0,1]}$ be the process from Theorem 4.4.6.*

Under Assumption 4.4.1, it holds that $S \xrightarrow{d} \Gamma$ with respect to the uniform topology as $n \rightarrow \infty$.

Proof. For any $0 \leq x < y \leq 1$, define the integrals

$$I_{xy} := \int_x^y G_z dM_z, \quad \text{and} \quad J_{xy} := \int_x^y E_z G_z d\gamma_z.$$

The process S_t is a sum of i.i.d. processes with same distribution as $I_{0t} - J_{0t}$, and note that this process has mean zero: it holds that $\mathbb{E}[I_{0t}] = \gamma_t$ by the definition of the LCM, and that $\mathbb{E}[J_{0t}] = \gamma_t$ since $\mathbb{E}[E_s G_s] = 1$. Thus we can apply the central limit theorem in Skorokhod space derived in van der Vaart and Wellner [2023, Example 1.14.24]. In view of this CLT, it suffices to show that there exist continuous, strictly increasing functions $F_1, F_2: [0, 1] \rightarrow [0, \infty)$ and constants $\epsilon_1, \epsilon_2 > 0$, such that for every $0 \leq s < t < u \leq 1$,

$$\mathbb{E}[(I_{st} + J_{st})^2] \leq |F_1(t) - F_1(s)|^{1/2 + \epsilon_1}, \quad (4.A.12)$$

$$\mathbb{E}[(I_{st} + J_{st})^2 (I_{tu} + J_{tu})^2] \leq |F_2(u) - F_2(s)|^{1 + \epsilon_2}. \quad (4.A.13)$$

To this end, let $0 \leq s < t < u \leq 1$ be given. In the following use $x \lesssim y$ to denote that the inequality $x \leq K \cdot y$ holds for a constant $K > 0$ independent of (s, t, u) . In fact, in each case the constant will be a polynomial over the constants $C_{\mathbf{h}}, C_E$, and c_{ρ}^{-1} from Assumption 4.4.1.

We first note that

$$J_{st} \lesssim t - s, \quad \text{and} \quad J_{tu} \lesssim u - t, \quad (4.A.14)$$

under Assumption 4.4.1. Using the quadratic characteristic for compensated counting processes, we also have

$$\begin{aligned} \mathbb{E}[I_{st}^2] &\lesssim \mathbb{E}\left[\left(\int_s^t G_r dM_r\right)^2\right] + \mathbb{E}\left[\left(\int_s^t G_r(\lambda_r - \lambda_r) dr\right)^2\right] \\ &= \mathbb{E}\left[\int_s^t G_r^2 \lambda dr\right] + \mathbb{E}\left[\left(\int_s^t G_r(\lambda_r - \lambda_r) dr\right)^2\right] \\ &\lesssim (t - s) + (t - s)^2 \leq 2(t - s) \end{aligned}$$

Thus we conclude that

$$\mathbb{E}[(I_{st} + J_{st})^2] \lesssim \mathbb{E}[I_{st}^2] + \mathbb{E}[J_{st}^2] \lesssim t - s.$$

This shows that the condition (4.A.12) holds with $\epsilon_1 = \frac{1}{2}$ and $F_1(r) = C_1 r$ for a suitably large constant $C_1 > 0$.

For the second condition we have that

$$\mathbb{E}[(I_{st} + J_{st})^2(I_{tu} + J_{tu})^2] \leq 4(\mathbb{E}[I_{st}^2 I_{tu}^2] + \mathbb{E}[I_{st}^2 J_{tu}^2] + \mathbb{E}[J_{st}^2 I_{tu}^2] + \mathbb{E}[J_{st}^2 J_{tu}^2]).$$

By same argument as $\mathbb{E}[I_{st}^2] \lesssim t - s$, we also have that $\mathbb{E}[I_{tu}^2] \lesssim t - u$. Combining this with the inequality (4.A.14) we obtain

$$\mathbb{E}[I_{st}^2 J_{tu}^2] + \mathbb{E}[J_{st}^2 I_{tu}^2] + \mathbb{E}[J_{st}^2 J_{tu}^2] \lesssim (t - s)(u - t)$$

For the term $\mathbb{E}[I_{st}^2 I_{tu}^2]$, we note that if $T \in (s, t]$ then M is constant on the interval $(t, u]$, in which case $I_{tu} = 0$. Therefore

$$\mathbb{E}[I_{st}^2 I_{tu}^2] = \mathbb{E}\left[\left(\int_s^t G_r \lambda_r dr\right)^2 I_{tu}^2\right] \lesssim (t - s)^2 \mathbb{E}[I_{tu}^2] \lesssim (t - s)^2(u - t)$$

Combined we conclude that

$$\mathbb{E}[(I_{st} + J_{st})^2(I_{tu} + J_{tu})^2] \lesssim (t - s)(u - t) \leq (u - s)^2.$$

This shows that the second condition (4.A.13) is satisfied with $\epsilon = 1$ and $F_2(r) = C_2 r$ for a suitably large constant $C_2 > 0$. \square

Returning to the proof of Theorem 4.4.6, the limit distribution of the ACM estimator is now a simple consequence of the continuous mapping theorem.⁵ Using the limits established in Proposition 4.4.2, Theorem 4.4.5, and Theorem 4.A.6, we conclude that

$$\sqrt{n_2}(\hat{\gamma} - \gamma) = \underbrace{S}_{\xrightarrow{d} \Gamma} + \underbrace{(\sqrt{n_2} \cdot \tilde{\gamma} - U)}_{\xrightarrow{p} 0} + \underbrace{(B - V)}_{\xrightarrow{p} 0} \xrightarrow{d} \Gamma \quad (4.A.15)$$

⁵Applied to the map $\mathbb{D}^3 \ni (x, y, z) \mapsto x + y + z \in \mathbb{D}$, where \mathbb{D} is the space of càglàd functions on $[0, 1]$ endowed with the uniform norm.

4 Assumption-lean Aalen regression

in the uniform topology as $n \rightarrow \infty$.

Turning to the cross-fitted ACM estimator, let $\hat{\gamma}^k$ denote ACM estimator fitted with $\mathcal{I}_1 = [n] \setminus J_k$ and $\mathcal{I}_2 = J_k$ for each $k \in [K]$. Let also

$$S_t^k = \frac{1}{\sqrt{|J_k|}} \sum_{i \in J_k} s_t^{(i)}, \quad \text{where} \quad s_t^{(i)} := \int_0^t G_s^{(i)} dM_s^{(i)} + \int_0^t E_s^{(i)} G_s^{(i)} d\gamma_s,$$

denote the corresponding oracle processes.

In the following computation, we use $x \sim y$ as a shorthand for $\sup_{t \in [0,1]} |x_t - y_t| \xrightarrow{P} 0$. Using that folds are assumed to have sizes $|J_k| \in \{\lfloor n/K \rfloor, \lceil n/K \rceil\}$, Proposition 4.4.2 and Theorem 4.4.5 applied to each fold yields that

$$\sqrt{n}(\tilde{\gamma} - \gamma) \sim \frac{\sqrt{n}}{K} \sum_{k=1}^K \frac{S^k}{\sqrt{|J_k|}} = \frac{\sqrt{n}}{K} \sum_{k=1}^K \frac{1}{|J_k|} \sum_{i \in J_k} s^{(i)} \sim \frac{1}{\sqrt{n}} \sum_{i \in [n]} s^{(i)}$$

where last equivalence follows from Lemma 4.A.7. By the continuous mapping theorem, applied as in (4.A.15), it suffices to establish that $\frac{1}{\sqrt{n}} \sum_{i \in [n]} s^{(i)} \xrightarrow{d} \Gamma$. However, this is simply the CLT established in Theorem 4.A.6 with $\mathcal{I}_2 = [n]$. \square

The following lemma is a Banach space generalization of one of the standard steps for proving validity of cross-fitting, cf. Lemma S1 in Lundborg and Pfister [2023].

Lemma 4.A.7. *Let W be a random element in a Banach space $(\mathbb{D}, \|\cdot\|)$, satisfying that $\mathbb{E}[\|W\|] < \infty$, and let $(W_i)_{i \in \mathbb{N}}$ be i.i.d. copies of W . For a fixed $K \in \mathbb{N}$, let $[n] = J_1 \cup \dots \cup J_K$ be a disjoint partition of $[n]$ with $|J_k| \in \{\lfloor n/K \rfloor, \lceil n/K \rceil\}$ for each $k \in [K]$. Finally, let $S = \frac{1}{n} \sum_{i \in [n]} W_i$ and $S_k = \frac{1}{|J_k|} \sum_{i \in J_k} W_i$ for each $k \in [K]$.*

Then, for any $r < 1$,

$$n^r \cdot \left(\frac{1}{K} \sum_{k=1}^K S_k - S \right) \xrightarrow{L^1} 0$$

as $n \rightarrow \infty$.

Proof. Note that $\frac{1}{n+K} < \frac{1}{K|J_k|} < \frac{1}{n-K}$ and hence $|\frac{1}{K|J_k|} - \frac{1}{n}| \leq \frac{1}{n-K} - \frac{1}{n} = \frac{K}{(n-K)n}$. Since

$$\frac{1}{K} \sum_{k=1}^K S_k - S = \sum_{k=1}^K \sum_{i \in J_k} \left(\frac{1}{K|J_k|} - \frac{1}{n} \right) W_i,$$

we conclude that

$$n^r \cdot \mathbb{E} \left\| \frac{1}{K} \sum_{k=1}^K S_k - S \right\| \leq \frac{Kn^r}{n-K} \mathbb{E} \|W\| \rightarrow 0$$

as $n \rightarrow \infty$. \square

4.B. Multiplicative hazards models

The predictable projection, $\pi_t(\bar{Z}_t) = \mathbb{E}[X_t | T \geq t, \bar{Z}_t]$, is generally difficult to compute – even with baseline covariates Z and a binary baseline exposure X , where $\pi_t(Z) = \mathbb{P}(X = 1 | T \geq t, Z)$. However, for a multiplicative hazards model, and under a strengthened censoring assumption, we give a fairly explicit representation of $\pi_t(Z)$ in Proposition 4.B.1 below. Example 4.6.2 elaborates further on the ACM in the Cox proportional hazards model.

Proposition 4.B.1 (Predictable projection for binary exposure). *Suppose that $X \in \{0, 1\}$ is a binary baseline exposure, that Z is a baseline covariate and that the model is a nonparametric time-varying proportional hazards model:*

$$\mathbf{h}_t(X, Z) = (X\theta_t + 1)\varphi_t(Z).$$

If $C \perp\!\!\!\perp (T^*, X) | Z$, then the predictable projection is given by

$$\pi_t(Z) = \frac{\pi_0(Z)}{\pi_0(Z) + (1 - \pi_0(Z))e^{\mathbf{l}_t(Z)}},$$

where $\pi_0(Z) = \mathbb{P}(X = 1 | Z)$ is the (baseline) propensity score and

$$\mathbf{l}_t(Z) := \int_0^t \theta_s \varphi_s(Z) ds.$$

The predictable projection can be expressed as the logistic model:

$$\text{logit}(\pi_t(Z)) = \text{logit}(\pi_0(Z)) - \mathbf{l}_t(Z).$$

Proof. Recall that $Y_t = \mathbb{1}(T \geq t) = \mathbb{1}(C \geq t)\mathbb{1}(T^* \geq t)$. Using the definition of $\pi_t(Z)$ and the conditional independence assumption regarding censoring, we find that

$$\begin{aligned} \pi_t(Z) &= \mathbb{E}[X | T \geq t, Z] = \frac{\mathbb{E}[XY_t | Z]}{\mathbb{P}(T \geq t | Z)} \\ &= \frac{\mathbb{E}[X\mathbb{1}(T^* \geq t) | Z]\mathbb{P}(C \geq t | Z)}{\mathbb{P}(T^* \geq t | Z)\mathbb{P}(C \geq t | Z)} \\ &= \frac{\mathbb{E}[X\mathbb{1}(T^* \geq t) | Z]}{\mathbb{P}(T^* \geq t | Z)}. \end{aligned}$$

Using the correspondence between survival time and cumulative hazard we obtain that

$$\begin{aligned} \mathbb{E}[X\mathbb{1}(T^* \geq t) | Z] &= \mathbb{E}[X \cdot \mathbb{P}(T^* \geq t | X, Z) | Z] \\ &= \mathbb{E}\left[X \exp\left(-X \int_0^t \theta_s \varphi_s(Z) ds - \int_0^t \varphi_s(Z) ds\right) | Z\right] \\ &= \mathbb{P}(X = 1 | Z) \exp\left(-\int_0^t \theta_s \varphi_s(Z) ds - \int_0^t \varphi_s(Z) ds\right) \\ &= \pi_0(Z) e^{-\mathbf{l}_t(Z) - \mathbf{l}_t(Z)}, \end{aligned}$$

4 Assumption-lean Aalen regression

where $l_t(z) := \int_0^t \theta_s \varphi_s(z) ds$ and $l_t(z) := \int_0^t \varphi_s(z) ds$. Similarly the denominator is given by

$$\begin{aligned} \mathbb{P}(T^* \geq t | Z) &= \mathbb{E} \left[\exp \left(-X \int_0^t \theta_s \varphi_s(Z) ds - \int_0^t b_s \varphi_s(Z) ds \right) | Z \right] \\ &= ((1 - \pi_0(Z)) + e^{-l_t(Z)} \pi_0(Z)) e^{-l_t(Z)}. \end{aligned}$$

From this we obtain that

$$\begin{aligned} \pi_t(Z) &= \frac{\pi_0(Z) e^{-l_t(Z)}}{(1 - \pi_0(Z)) + e^{-l_t(Z)} \pi_0(Z)} \\ &= \frac{\pi_0(Z)}{\pi_0(Z) + (1 - \pi_0(Z)) e^{l_t(Z)}} \\ &= \frac{1}{1 + e^{-\text{logit}(\pi_0(Z))} e^{l_t(Z)}} = \text{expit}(\text{logit}(\pi_0(Z)) - l_t(Z)). \end{aligned}$$

□

Example 4.6.2 (ACM in a Cox proportional hazards model). With baseline covariates $Z \in \mathbb{R}^d$, in addition to a binary exposure $X \in \{0, 1\}$, the Cox proportional hazards model asserts that:

$$h_t(X, Z) = h^0(t) \exp(\beta_X X + \beta_Z^\top Z) = (X(e^{\beta_X} - 1) + 1) h^0(t) \exp(\beta_Z^\top Z). \quad (4.B.1)$$

Letting $\phi(Z) = (e^{\beta_X} - 1) e^{\beta_Z^\top Z}$ we see that this model is a special case of Example 4.1.1 with $\theta_t(Z) = h^0(t) \phi(Z)$. It follows by (4.1.4) that the ACM can be expressed as

$$\begin{aligned} \gamma_t &= \int_0^t \mathbb{E} [h^0(s) \phi(Z) w_s^{\text{AC}}(Z)] ds \\ &= \mathbb{E} \left[\phi(Z) \underbrace{\int_0^t h^0(s) w_s^{\text{AC}}(Z) ds}_{=: W_t^{\text{AC}}(Z)} \right] = \mathbb{E} [\phi(Z) W_t^{\text{AC}}(Z)] = \mathbb{E} \left[e^{\beta_Z^\top Z} W_t^{\text{AC}}(Z) \right] (e^{\beta_X} - 1). \end{aligned}$$

Thus $\gamma_t = c(t)(e^{\beta_X} - 1)$ is proportional to $e^{\beta_X} - 1$ with a t -dependent proportionality constant $c(t) = \mathbb{E} \left[e^{\beta_Z^\top Z} W_t^{\text{AC}}(Z) \right] \geq 0$.

Using the innovation theorem we see that

$$h_t(Z) = (\pi_t(Z)(e^{\beta_X} - 1) + 1) h^0(t) \exp(\beta_Z^\top Z),$$

and we could also arrive at the representation of the ACM directly from the identity (4.2.13). Indeed, $h_s(X, Z) - h_s(Z) = (e^{\beta_X} - 1)(X - \pi_s(Z)) h^0(s) \exp(\beta_Z^\top Z)$, and (4.2.13) gives

$$\gamma_t = (e^{\beta_X} - 1) \int_0^t h^0(s) \mathbb{E} \left[e^{\beta_Z^\top Z} \frac{Y_s(X - \pi_s(Z))^2}{\rho(s)} \right] ds$$

The proportionality constant must, of course, be $c(t)$, but this can also be verified directly using the same argument as in Example (4.2.6). We may note that $h_t(\cdot)$ is not generally of the form of a Cox proportional hazards model, which illustrates the fact that the Cox model is not closed under marginalization.

To compute the proportionality constant we need the weights

$$w_t^{\text{AC}}(Z) = \frac{Y_t \pi_t(Z)(1 - \pi_t(Z))}{\mathbb{E}[Y_t \pi_t(Z)(1 - \pi_t(Z))]},$$

and to this end we need to compute the predictable projection $\pi_t(Z)$. Under the censoring assumption $C \perp\!\!\!\perp (T^*, X) \mid Z$, it follows from Proposition 4.B.1 that

$$\text{logit}(\pi_t(Z)) = \text{logit}(\pi_0(Z)) - \phi(Z)H^0(t), \quad (4.B.2)$$

where $H^0(t) = \int_0^t h^0(s)ds$. This shows that $\pi_t(Z)$ evolves in time with a logistic growth rate precomposed with H^0 . Formula (4.B.2) makes it straightforward to compute $\pi_t(z)(1 - \pi_t(z))$ for any given t and z (provided that we can compute $\pi_0(z)$), but it does not appear to yield a simple analytic expression of neither the weights $w_t^{\text{AC}}(Z)$ nor the proportionality constant $c(t)$. It is, however, useful for Monte Carlo computation of the ACM.



Bibliography

- O. Aalen. A model for nonparametric regression analysis of counting processes. In *Mathematical Statistics and Probability Theory: Proceedings, Sixth International Conference, Wisła (Poland), 1978*, pages 1–25. Springer, 1980.
- O. O. Aalen. Dynamic modelling and causality. *Scandinavian Actuarial Journal*, pages 177–190, 1987.
- O. O. Aalen. A linear regression model for the analysis of life times. *Statistics in Medicine*, 8(8):907–925, 1989.
- O. O. Aalen, K. Røysland, J. M. Gran, and B. Ledergerber. Causality, mediation and time: a dynamic viewpoint. *Journal of the Royal Statistical Society. Series A (Statistics in Society)*, 175(4):831–861, 2012.
- M. Achab, E. Bacry, S. Gaïffas, I. Mastromatteo, and J.-F. Muzy. Uncovering causality from multivariate Hawkes integrated cumulants. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70, pages 1–10. PMLR, 06–11 Aug 2017.
- R. J. Adler and J. E. Taylor. *Random Fields and Geometry*, volume 80 of *Springer Monographs in Mathematics*. Springer, 2007.
- P. K. Andersen, Ø. Borgan, R. D. Gill, and N. Keiding. *Statistical models based on counting processes*. Springer Series in Statistics. Springer-Verlag, New York, 1993.
- E. Bacry, M. Bompain, P. Deegan, S. Gaïffas, and S. V. Poulsen. tick: a Python library for statistical learning, with an emphasis on Hawkes processes and time-dependent models. *Journal of Machine Learning Research*, 18(214):1–5, 2018.
- I. Baldé, Y. A. Yang, and G. Lefebvre. Reader reaction to “Outcome-adaptive lasso: Variable selection for causal inference” by Shortreed and Ertefaie (2017). *Biometrics*, 79(1):514–520, 2023.
- A. Bender, D. Rügamer, F. Scheipl, and B. Bischl. A general machine learning framework for survival analysis. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 158–173. Springer, 2020.
- V. Bengs and H. Holzmann. Uniform approximation in classical weak convergence theory. *arXiv preprint arXiv:1903.09864*, 2019.

Bibliography

- D. Benkeser and M. van der Laan. The highly adaptive lasso estimator. In *2016 IEEE international conference on data science and advanced analytics (DSAA)*, pages 689–696. IEEE, 2016.
- D. Benkeser, W. Cai, and M. van der Laan. A nonparametric super-efficient estimator of the average treatment effect. *Statistical Science*, 35:484–495, 08 2020.
- R. Berk, A. Buja, L. Brown, E. George, A. K. Kuchibhotla, W. Su, and L. Zhao. Assumption lean regression. *The American Statistician*, 2019.
- A. Bianchi, A. M. Christgau, and J. S. Pedersen. On generators and relations of the rational cohomology of Hilbert schemes. *Journal of Algebraic Combinatorics*, 57(3): 829–857, 2023.
- P. J. Bickel, C. A. Klaassen, Y. Ritov, and J. A. Wellner. *Efficient and Adaptive Estimation for Semiparametric Models*. Springer Book Archive. Springer-Verlag New York, 1 edition, 1998. ISBN 978-0-387-98473-5.
- P. Billingsley. *Convergence of probability measures*. John Wiley & Sons, 2013.
- P. Billingsley. *Probability and measure*. John Wiley & Sons, 2017.
- S. M. Bischofberger, M. Hiabu, E. Mammen, and J. P. Nielsen. Smooth backfitting for additive hazard rates. *arXiv preprint arXiv:2302.09510*, 2023.
- C. S. Bojer and J. P. Meldgaard. Kaggle forecasting competitions: An overlooked learning opportunity. *International Journal of Forecasting*, 37(2):587–603, 2021.
- S. Boucheron, G. Lugosi, and P. Massart. *Concentration inequalities: A nonasymptotic theory of independence*. Oxford University Press, 2013.
- O. Bousquet and A. Elisseeff. Stability and generalization. *Journal of Machine Learning Research*, 2:499–526, 2002.
- L. Breiman. Statistical modeling: The two cultures (with comments and a rejoinder by the author). *Statistical science*, 16(3):199–231, 2001.
- P. Brémaud. *Point processes and queues*. Springer-Verlag, New York, 1981.
- N. E. Breslow and N. E. Day. Statistical methods in cancer research. Volume II—The design and analysis of cohort studies. *IARC Scientific Publications*, 82:1–406, 1987.
- R. Cai, S. Wu, J. Qiao, Z. Hao, K. Zhang, and X. Zhang. THPs: Topological Hawkes processes for learning causal structure on event sequences. *IEEE Transactions on Neural Networks and Learning Systems*, pages 1–15, 2022.
- T. T. Cai and M. Yuan. Minimax and adaptive prediction for functional linear regression. *Journal of the American Statistical Association*, 107(499):1201–1216, 2012.

- Q. Chen, V. Syrgkanis, and M. Austern. Debiased machine learning without sample-splitting for stable estimators. *Advances in Neural Information Processing Systems*, 35:3096–3109, 2022.
- V. Chernozhukov, D. Chetverikov, M. Demirer, E. Duffo, C. Hansen, W. Newey, and J. Robins. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68, 01 2018. ISSN 1368-4221.
- M. Chickering, D. Heckerman, and C. Meek. Large-sample learning of bayesian networks is NP-hard. *Journal of Machine Learning Research*, 5:1287–1330, 2004.
- A. M. Christgau and N. R. Hansen. Assumption-lean Aalen regression. In preparation, 2024+.
- A. M. Christgau and N. R. Hansen. Efficient adjustment for complex covariates: Gaining efficiency with DOPE. *arXiv preprint arXiv:2402.12980*, 2024.
- A. M. Christgau, A. Arnaudon, and S. Sommer. Moment evolution equations and moment matching for stochastic image EPDiff. *Journal of Mathematical Imaging and Vision*, 65(4):563–576, 2023a.
- A. M. Christgau, L. Petersen, and N. R. Hansen. Nonparametric conditional local independence testing. *Annals of Statistics*, 51(5):2116–2144, 2023b.
- A. M. Christgau, L. Petersen, and N. R. Hansen. Supplement to: Nonparametric conditional local independence testing. *Annals of Statistics*, 51(5), 2023c. doi: 10.1214/23-AOS2323SUPPA.
- S. N. Cohen and R. J. Elliott. *Stochastic calculus and applications*, volume 2. Springer, 2015.
- D. Commenges and A. Gégout-Petit. A general dynamical statistical model with causal interpretation. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 71(3):719–736, 2009.
- C. S. Cox, J. J. Feldman, C. D. Golden, M. A. Lane, J. H. Madans, M. E. Mussolino, and S. T. Rothwell. Plan and operation of the NHANES I epidemiologic follow-up study, 1992. *Vital and health statistics.*, 12 1997.
- D. R. Cox. Regression models and life-tables. *Journal of the Royal Statistical Society: Series B (Methodological)*, 34(2):187–202, 1972.
- R. Daniel, J. Zhang, and D. Farewell. Making apples from oranges: comparing non-collapsible effect estimators and their standard errors after adjustment for different covariate sets. *Biometrical Journal*, 63(3):528–557, 2021.
- C. Davidson-Pilon. lifelines: survival analysis in Python. *Journal of Open Source Software*, 4(40):1317, 2019. doi: 10.21105/joss.01317.

Bibliography

- M. Delecroix, W. Härdle, and M. Hristache. Efficient estimation in conditional single-index regression. *Journal of Multivariate Analysis*, 86(2):213–226, 2003.
- J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. *arXiv:1810.04805*, 2018.
- V. Didelez. Graphical models for composable finite Markov processes. *Scandinavian Journal of Statistics*, 34(1):169–185, 2006.
- V. Didelez. Graphical models for marked point processes based on local independence. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 70(1):245–264, 2008.
- V. Didelez. Causal reasoning for events in continuous time: A decision-theoretic approach. In *Proceedings of the UAI 2015 Workshop on Advances in Causal Inference*, 2015.
- V. Didelez and M. J. Stensrud. On the logic of collapsibility for causal effect measures. *Biometrical Journal*, 64(2):235–242, 2022.
- O. Dukes, T. Martinussen, E. J. Tchetgen Tchetgen, and S. Vansteelandt. On doubly robust estimation of the hazard difference. *Biometrics*, 75(1):100–109, 2019.
- M. Émery and W. Schachermayer. On Vershik’s standardness criterion and Tsirelson’s notion of cosiness. In *Séminaire de Probabilités XXXV*, pages 265–305. Springer, 2001.
- E. X. Fang, Y. Ning, and H. Liu. Testing and confidence intervals for high dimensional proportional hazards models. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, pages 1415–1437, 2017.
- T. R. Fleming and D. P. Harrington. *Counting processes and survival analysis*, volume 169. John Wiley & Sons, 2011.
- P. Forré and J. M. Mooij. A mathematical introduction to causality. Lecture notes, 2023.
- S. S. Franklin, S. A. Khan, N. D. Wong, M. G. Larson, and D. Levy. Is pulse pressure useful in predicting risk for coronary heart disease? The Framingham heart study. *Circulation*, 100(4):354–360, 1999.
- J. H. Friedman. Greedy function approximation: a gradient boosting machine. *Annals of Statistics*, pages 1189–1232, 2001.
- N. Gnecco, J. Peters, S. Engelke, and N. Pfister. Boosted control functions. *arXiv:2310.05805*, 2023.
- C. W. J. Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 37(3):424–438, 1969.

- K. Greenewald, K. Shanmugam, and D. Katz. High-dimensional feature selection for sample efficient treatment effect estimation. In *International Conference on Artificial Intelligence and Statistics*, pages 2224–2232. PMLR, 2021.
- F. R. Guo, A. R. Lundborg, and Q. Zhao. Confounder selection: Objectives and approaches. *arXiv:2208.13871*, 2022.
- L. Györfi, M. Kohler, A. Krzyżak, and H. Walk. *A Distribution-Free Theory of Nonparametric Regression*. Springer Series in Statistics. Springer New York, NY, 1 edition, 2002. ISBN 978-0-387-22442-8. doi: 10.1007/b97848.
- J. Hahn. On the role of the propensity score in efficient semiparametric estimation of average treatment effects. *Econometrica*, pages 315–331, 1998.
- B. B. Hansen. The prognostic analogue of the propensity score. *Biometrika*, 95(2): 481–488, 2008.
- M. Hardt, B. Recht, and Y. Singer. Train faster, generalize better: Stability of stochastic gradient descent. In *International conference on Machine Learning*, pages 1225–1234. PMLR, 2016.
- J. Harezlak, B. A. Coull, N. M. Laird, S. R. Magari, and D. C. Christiani. Penalized solutions to functional regression problems. *Computational statistics & data analysis*, 51(10):4911–4925, 2007.
- T. Hastie, R. Tibshirani, J. H. Friedman, and J. H. Friedman. *The elements of statistical learning: data mining, inference, and prediction*, volume 2. Springer, 2009.
- L. Henckel, E. Perković, and M. H. Maathuis. Graphical criteria for efficient total effect estimation via adjustment in causal linear models. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 84(2):579–599, 2022.
- M. A. Hernán. The hazards of hazard ratios. *Epidemiology*, 21(1):13–15, 2010.
- M. Hiabu, E. Mammen, M. D. Martínez-Miranda, and J. P. Nielsen. Smooth backfitting of proportional hazards with multiplicative components. *Journal of the American Statistical Association*, 116(536):1983–1993, 2021.
- O. Hines, O. Dukes, K. Diaz-Ordaz, and S. Vansteelandt. Demystifying statistical learning based on efficient influence functions. *The American Statistician*, 76(3):292–304, 2022.
- O. Hines, K. Diaz-Ordaz, and S. Vansteelandt. Optimally weighted average derivative effects. *arXiv preprint arXiv:2308.05456*, 2023.
- T. Homan, S. Bordes, and E. Cichowski. Physiology, pulse pressure. In: StatPearls [Internet]. Treasure Island (FL): StatPearls Publishing, 2024. Updated 2023 Jul 10. Retrived from: <https://www.ncbi.nlm.nih.gov/books/NBK482408/>.

Bibliography

- J. Hou, J. Bradic, and R. Xu. Treatment effect estimation under additive hazards models with high-dimensional confounding. *Journal of the American Statistical Association*, 118(541):327–342, 2023.
- J. Huang. Efficient estimation of the partly linear additive Cox model. *Annals of Statistics*, 27(5):1536–1563, 1999.
- H. Ichimura. Semiparametric least squares (SLS) and weighted SLS estimation of single-index models. *Journal of Econometrics*, 58(1-2):71–120, 1993.
- C.-R. Jiang and J.-L. Wang. Functional single index models for longitudinal data. *Annals of Statistics*, 39(1):362–388, 2011.
- C. Ju, S. Gruber, S. D. Lendle, A. Chambaz, J. M. Franklin, R. Wyss, S. Schneeweiss, and M. J. van der Laan. Scalable collaborative targeted learning for high-dimensional data. *Statistical methods in medical research*, 28(2):532–554, 2019.
- C. Ju, D. Benkeser, and M. J. van der Laan. Robust inference on the average treatment effect using the outcome highly adaptive lasso. *Biometrics*, 76(1):109–118, 2020.
- O. Kallenberg. *Foundations of modern probability*, volume 3. Springer, 2021.
- M. Kasy. Uniformity and the delta method. *Journal of Econometric Methods*, 8(1), 2019.
- E. H. Kennedy. Semiparametric doubly robust targeted double machine learning: a review. *arXiv preprint arXiv:2203.06469*, 2022.
- H. Klyne and R. D. Shah. Average partial effect estimation using double machine learning. *arXiv preprint arXiv:2308.09207*, 2023.
- A. Kolmogoroff. *Grundbegriffe der Wahrscheinlichkeitsrechnung*. Ergebnisse der Mathematik und Ihrer Grenzgebiete. 1. Folge. Springer Berlin, Heidelberg, 1 edition, 1933. ISBN 978-3-642-49596-0.
- Ø. Kravdal. The attractiveness of an additive hazard model: An example from medical demography. *European Journal of Population / Revue européenne de Démographie*, 13(1):33–47, 1997. ISSN 1572-9885. doi: 10.1023/A:1005768425808.
- M. Ledoux and M. Talagrand. *Probability in Banach Spaces: Isoperimetry and Processes*, volume 23. Springer Science & Business Media, 1991.
- D. K. Lee, N. Chen, and H. Ishwaran. Boosted nonparametric hazards with time-dependent covariates. *Annals of Statistics*, 49(4):2101, 2021.
- M. Y. Lee S. McDaniel and R. Chappell. Analysis and design of clinical trials using additive hazards survival endpoints. *Statistics in Biopharmaceutical Research*, 11(3): 274–282, 2019. doi: 10.1080/19466315.2019.1575278.

- D. Y. Lin and Z. Ying. Semiparametric analysis of the additive risk model. *Biometrika*, 81(1):61–71, 03 1994. ISSN 0006-3444. doi: 10.1093/biomet/81.1.61.
- J. J. Lok. Statistical modeling of causal effects in continuous time. *Annals of Statistics*, 36(3):1464–1507, 06 2008.
- J. J. Lok. Mimicking counterfactual outcomes to estimate causal effects. *Annals of Statistics*, 45(2):461, 2017.
- S. M. Lundberg, G. Erion, H. Chen, A. DeGrave, J. M. Prutkin, B. Nair, R. Katz, J. Himmelfarb, N. Bansal, and S.-I. Lee. From local explanations to global understanding with explainable AI for trees. *Nature machine intelligence*, 2(1):56–67, 2020.
- A. R. Lundborg. Modern methods for variable significance testing. *Apollo - University of Cambridge Repository*, 2023. doi: 10.17863/CAM.93556.
- A. R. Lundborg and N. Pfister. Perturbation-based analysis of compositional data. *arXiv:2311.18501*, 2023.
- A. R. Lundborg, I. Kim, R. D. Shah, and R. J. Samworth. The projected covariance measure for assumption-lean variable significance testing. *arXiv preprint arXiv:2211.02039*, 2022a.
- A. R. Lundborg, R. D. Shah, and J. Peters. Conditional independence testing in Hilbert spaces with applications to functional data analysis. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 84(5):1821–1850, 2022b. doi: 10.1111/rssb.12544.
- A. Maity. Nonparametric functional concurrent regression models. *Wiley Interdisciplinary Reviews: Computational Statistics*, 9(2):e1394, 2017.
- N. Malfait and J. O. Ramsay. The historical functional linear model. *The Canadian Journal of Statistics*, 31(2):115–128, 2003.
- T. Manrique. *Functional linear regression models: application to high-throughput plant phenotyping functional data*. PhD thesis, Université de Montpellier, 2016.
- T. Manrique, C. Crambes, and N. Hilgert. Ridge regression for the functional concurrent model. *Electronic Journal of Statistics*, 12(1):985–1018, 2018.
- T. Martinussen and T. H. Scheike. *Dynamic regression models for survival data*, volume 1. Springer, 2006.
- T. Martinussen and S. Vansteelandt. On collapsibility and confounding bias in Cox and Aalen regression models. *Lifetime Data Analysis*, 19(3):279–296, 2013. ISSN 1380-7870. doi: 10.1007/s10985-013-9242-z.
- T. Martinussen, S. Vansteelandt, and P. K. Andersen. Subtleties in the interpretation of hazard contrasts. *Lifetime Data Analysis*, 26:833–855, 2020.

Bibliography

- I. W. McKeague and P. D. Sasieni. A partly parametric additive risk model. *Biometrika*, 81(3):501–514, 1994.
- S. W. Mogensen and N. R. Hansen. Markov equivalence of marginalized local independence graphs. *Annals of Statistics*, 48(1):539–559, 2020.
- S. W. Mogensen and N. R. Hansen. Graphical modeling of stochastic processes driven by correlated noise. *Bernoulli*, 28(4):3023–3050, 2022.
- S. W. Mogensen, D. Malinsky, and N. R. Hansen. Causal learning for partially observed stochastic dynamical systems. In *Proceedings of the 34th conference on Uncertainty in Artificial Intelligence*, pages 350–360, 2018.
- W. K. Newey. Uniform convergence in probability and stochastic equicontinuity. *Econometrica*, 59(4):1161–1167, 1991.
- W. K. Newey. The asymptotic variance of semiparametric estimators. *Econometrica: Journal of the Econometric Society*, pages 1349–1382, 1994.
- M. Neykov, S. Balakrishnan, and L. Wasserman. Minimax optimal conditional independence testing. *Annals of Statistics*, 49(4):2151–2177, 2021.
- Z. Niu, A. Chakraborty, O. Dukes, and E. Katsevich. Reconciling model-x and doubly robust approaches to conditional independence testing. *arXiv preprint arXiv:2211.14698*, 2022.
- A. Pakbin, X. Wang, B. J. Mortazavi, and D. K. Lee. BoXHED2.0: Scalable boosting of dynamic survival analysis. *arXiv preprint arXiv:2103.12591*, 2021.
- S. Parkinson, G. Ongie, and R. Willett. Linear neural network layers promote learning single-and multiple-index models. *arXiv:2305.15598*, 2023.
- A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.
- J. Pearl. *Causality*. Cambridge University Press, 2009.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- E. Perković, J. Textor, M. Kalisch, and M. H. Maathuis. Complete graphical characterization and construction of adjustment sets in markov equivalence classes of ancestral graphs. *Journal of Machine Learning Research*, 18(220):1–62, 2018.

- J. Peters, D. Janzing, and B. Schölkopf. *Elements of causal inference: foundations and learning algorithms*. The MIT Press, 2017.
- L. Petersen and N. R. Hansen. Testing conditional independence via quantile regression based partial copulas. *Journal of Machine Learning Research*, 22(70):1–47, 2021.
- J. Pfanzagl and W. Wefelmeyer. Contributions to a general asymptotic statistical theory. *Statistics & Risk Modeling*, 3(3-4):379–388, 1985.
- D. Pollard. *Convergence of stochastic processes*. Springer Series in Statistics. Springer-Verlag, New York, 1984.
- S. Pölsterl. scikit-survival: A library for time-to-event analysis built on top of scikit-learn. *Journal of Machine Learning Research*, 21(212):1–6, 2020.
- T. Popoviciu. Sur les équations algébriques ayant toutes leurs racines réelles. *Mathematica*, 9(129-145):20, 1935.
- J. L. Powell, J. H. Stock, and T. M. Stoker. Semiparametric estimation of index coefficients. *Econometrica: Journal of the Econometric Society*, pages 1403–1430, 1989.
- R. Rebolledo. Central limit theorems for local martingales. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 51(3):269–286, 1980.
- D. Revuz and M. Yor. *Continuous martingales and Brownian motion*, volume 293. Springer Science & Business Media, 2013.
- J. M. Robins and A. Rotnitzky. Semiparametric efficiency in multivariate regression models with missing data. *Journal of the American Statistical Association*, 90(429):122–129, 1995.
- J. M. Robins, A. Rotnitzky, and L. P. Zhao. Estimation of regression coefficients when some regressors are not always observed. *Journal of the American statistical Association*, 89(427):846–866, 1994.
- P. M. Robinson. Root-n-consistent semiparametric regression. *Econometrica: Journal of the Econometric Society*, pages 931–954, 1988.
- L. C. G. Rogers and D. Williams. *Diffusions, Markov processes, and martingales*, volume 2. Cambridge University Press, Cambridge, 2000.
- P. R. Rosenbaum and D. B. Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1):41–55, 1983.
- A. Rotnitzky and E. Smucler. Efficient adjustment sets for population average causal treatment effect estimation in graphical models. *Journal of Machine Learning Research*, 21(188):1–86, 2020.

Bibliography

- K. Røysland, P. Ryalen, M. Nygård, and V. Didelez. Graphical criteria for the identification of marginal causal effects in continuous-time survival and event-history analyses. *arXiv preprint arXiv:2202.02311*, 2022.
- H. C. Rytgaard, T. A. Gerds, and M. J. van der Laan. Continuous-time targeted minimum loss-based estimation of intervention-specific mean outcomes. *Annals of Statistics*, 50(5):2469–2491, 2022.
- H. C. W. Rytgaard, F. Eriksson, and M. van der Laan. Estimation of time-specific intervention effects on continuously distributed time-to-event outcomes by targeted maximum likelihood estimation. *arXiv:2106.11009*, 2021.
- P. Sasieni. Information bounds for the conditional hazard ratio in a nested family of regression models. *Journal of the Royal Statistical Society: Series B (Methodological)*, 54(2):617–635, 1992.
- saz. Calculating the expectation of the supremum of absolute value of a Brownian motion. Mathematics Stack Exchange, 2019. Retrieved (2019-06-06) from: <https://math.stackexchange.com/q/3252132>.
- C. Scheidegger, J. Hörrmann, and P. Bühlmann. The weighted generalised covariance measure. *Journal of Machine Learning Research*, 23(273):1–68, 2022.
- R. L. Schilling. *Measures, integrals and martingales*. Cambridge University Press, 2017.
- R. L. Schilling and L. Partzsch. *Brownian Motion: An Introduction to Stochastic Processes*. De Gruyter, 2012. ISBN 9783110278989. doi: doi:10.1515/9783110278989.
- B. Schölkopf, F. Locatello, S. Bauer, N. R. Ke, N. Kalchbrenner, A. Goyal, and Y. Bengio. Toward causal representation learning. *Proceedings of the IEEE*, 109(5):612–634, 2021.
- T. Schweder. Composable Markov processes. *Journal of Applied Probability*, 7(2):400–410, 1970.
- D. Şentürk and H.-G. Müller. Functional varying coefficient models for longitudinal data. *Journal of the American Statistical Association*, 105(491):1256–1264, 2010.
- R. D. Shah and J. Peters. The hardness of conditional independence testing and the generalised covariance measure. *Annals of Statistics*, 48(3):1514–1538, 2020.
- C. Shi, D. Blei, and V. Veitch. Adapting neural networks for the estimation of treatment effects. *Advances in neural information processing systems*, 32, 2019.
- S. M. Shortreed and A. Ertefaie. Outcome-adaptive lasso: variable selection for causal inference. *Biometrics*, 73(4):1111–1122, 2017.
- E. Smucler, A. Rotnitzky, and J. M. Robins. A unifying approach for doubly-robust l1-regularized estimation of causal contrasts. *arXiv:1904.03737*, 2019.

- J. A. Soloff, R. F. Barber, and R. Willett. Stability via resampling: statistical problems beyond the real line. *arXiv preprint arXiv:2405.09511*, 2024a.
- J. A. Soloff, R. F. Barber, and R. Willett. Bagging provides assumption-free stability. *Journal of Machine Learning Research*, 25(131):1–35, 2024b.
- M. J. Stensrud and M. A. Hernán. Why test for proportional hazards? *Jama*, 323(14):1401–1402, 2020.
- A. B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer Series in Statistics. Springer New York, NY, 1 edition, 2009. ISBN 978-0-387-79052-7. doi: 10.1007/b13794.
- C. Uhler, G. Raskutti, P. Bühlmann, and B. Yu. Geometry of the faithfulness assumption in causal inference. *Annals of Statistics*, pages 436–463, 2013.
- M. J. van der Laan and S. Gruber. Collaborative double robust targeted maximum likelihood estimation. *The international journal of biostatistics*, 6(1), 2010.
- M. J. van der Laan and S. Rose. *Targeted learning: causal inference for observational and experimental data*, volume 4. Springer, 2011.
- A. W. van der Vaart. *Asymptotic statistics*, volume 3. Cambridge university press, 2000.
- A. W. van der Vaart and J. A. Wellner. *Weak convergence and empirical processes*. Springer Series in Statistics. Springer-Verlag, New York, 1996.
- A. W. van der Vaart and J. A. Wellner. *Weak Convergence and Empirical Processes: With Applications to Statistics*. Springer Series in Statistics. Springer Cham, 2 edition, 2023. ISBN 978-3-031-29038-1.
- G. Van Rossum, F. L. Drake, et al. *Python reference manual*, volume 111. Centrum voor Wiskunde en Informatica Amsterdam, 1995.
- S. Vansteelandt and O. Dukes. Assumption-lean inference for generalised linear model parameters. *Journal of the Royal Statistical Society Series B: Statistical Methodology (with discussion)*, 84(3):657–685, 2022.
- S. Vansteelandt, O. Dukes, K. Van Lancker, and T. Martinussen. Assumption-lean Cox regression. *Journal of the American Statistical Association*, pages 1–10, 2022.
- V. Veitch, D. Sridhar, and D. Blei. Adapting text embeddings for causal inference. In *Conference on Uncertainty in Artificial Intelligence*, pages 919–928. PMLR, 2020.
- X. Wang, A. Pakbin, B. Mortazavi, H. Zhao, and D. Lee. BoXHED: Boosted eXact Hazard Estimator with Dynamic covariates. In *International Conference on Machine Learning*, pages 9973–9982. PMLR, 2020.
- M. T. Wells. Nonparametric kernel estimation in counting processes with explanatory variables. *Biometrika*, 81(4):795–801, 1994.

Bibliography

- S. Xiao, J. Yan, M. Farajtabar, L. Song, X. Yang, and H. Zha. Learning time series associated event sequences with recurrent point process networks. *IEEE Transactions on Neural Networks and Learning Systems*, 30(10):3124–3136, 2019.
- H. Xu, M. Farajtabar, and H. Zha. Learning Granger causality for Hawkes processes. In *Proceedings of The 33rd International Conference on Machine Learning*, volume 48, pages 1717–1726, 2016.
- F. Yao, H.-G. Müller, and J.-L. Wang. Functional linear regression analysis for longitudinal data. *Annals of Statistics*, 33(6):2873 – 2903, 2005.
- M. Yuan and T. T. Cai. A reproducing kernel Hilbert space approach to functional linear regression. *Annals of Statistics*, 38(6):3412–3444, 2010.
- Q. Zhong, J. Mueller, and J.-L. Wang. Deep learning for the partially linear Cox model. *Annals of Statistics*, 50(3):1348–1375, 2022.
- K. Zhou, H. Zha, and L. Song. Learning social infectivity in sparse low-rank networks using multi-dimensional Hawkes processes. In *Proceedings of the 16th International Conference on Artificial Intelligence and Statistics*, 2013.
- P. N. Zivich and A. Breskin. Machine learning for causal inference: on the use of cross-fit estimators. *Epidemiology (Cambridge, Mass.)*, 32(3):393, 2021.

