Time-Varying Causal Survival Learning

Anonymous submission

Abstract

This work bridges the gap between staggered adoption designs and survival analysis to estimate causal effects in settings with time-varying treatments, addressing a fundamental challenge in medical research exemplified by the Stanford Heart Transplant study. In medical interventions, particularly organ transplantation, the timing of treatment varies significantly across patients due to factors such as donor availability and patient readiness, introducing potential bias in treatment effect estimation if not properly accounted for. We identify conditions under which staggered adoption assumptions can justify the use of survival analysis techniques for causal inference with time-varying treatments. By establishing this connection, we enable the use of existing survival analysis methods while maintaining causal interpretability. Furthermore, we enhance estimation performance by incorporating double machine learning methods, improving efficiency when handling complex relationships between patient characteristics and survival outcomes. Through both simulation studies and application to heart transplant data, our approach demonstrates superior performance compared to traditional methods, reducing bias and offering theoretical guarantees for improved efficiency in survival analysis settings.

Introduction

In healthcare, understanding the causal effect of medical interventions on patient survival is crucial. Heart transplantation is a compelling example, as demonstrated by the Stanford Heart Transplant study (Crowley and Hu 1977; Zhu et al. 2021), where patients with end-stage heart failure undergo surgery to replace their failing hearts with healthy donor hearts. While this procedure is likely to extend patients' lives on average, researchers are particularly interested in how treatment effects vary with patient characteristics and surgical details. Understanding these heterogeneous treatment effects—how the impact varies as a function of patient characteristics—is essential for improving patient selection criteria, optimizing intervention timing, and ultimately enhancing survival outcomes in transplant medicine (Trulock et al. 2007; Kilic et al. 2021; DeFilippis et al. 2022)

However, measuring the causal effect of such medical interventions is not straightforward. After a patient is listed as a candidate for heart transplant, they must wait for an available donor heart before undergoing the procedure (Almond et al. 2009). This means patients who receive heart transplants experience both control time (waiting period) and treatment time (post-transplant period). Furthermore, the treatment timing is random and highly variable. The waiting time for a heart transplant ranges from a few days to more than a year (Evans et al. 1986; O'Connell et al. 1992). This situation differs from traditional causal inference where treatment and control groups are determined at the study's outset and remain fixed throughout.

Staggered adoption designs in econometrics (Athey and Imbens 2022) provide a way to estimate causal effects with random treatment times. These designs compare outcomes between treated and yet-to-be-treated units, using later-treated units as controls for earlier-treated ones. While these designs work well for repeatedly measured continuous outcomes, they cannot handle time-to-event outcomes that are observed only once—either at the event occurrence (such as mortality) or at study end. This limitation stems from staggered adoption's reliance on multiple outcomes per unit, whereas survival settings provide only a single endpoint.

To address this challenge, we extend staggered adoption designs to time-to-event outcomes by integrating survival analysis techniques (Cox 1972; Klein and Moeschberger 1997; Fleming and Harrington 2005; Kalbfleisch and Prentice 2011). Using hazard functions to model instantaneous event probabilities allows us to characterize outcomes continuously over time, spanning both control and treatment periods. Our key contribution lies in establishing conditions under which these hazard-based models enable valid causal inference.

In addition to handling random treatment timing, we must address the complex, non-linear relationships between covariates and outcomes that often arise in real-world applications (Hastie et al. 2009). To handle these complexities, we employ double machine learning (DML) techniques (Chernozhukov et al. 2018; Künzel et al. 2019; Nie and Wager 2021; Gao and Hastie 2021) in our estimation procedure. DML provides a powerful framework for improving estimation efficiency, allowing us to flexibly model non-linear relationships while maintaining robustness to potential model misspecifications.

The contributions of this paper are twofold. First, we bridge the gap between staggered adoption designs and survival analysis by identifying conditions under which time-varying treatment effects can be estimated in a survival

framework. Specifically, we show how key assumptions from staggered adoption designs can be adapted to justify the use of existing survival analysis techniques for causal inference with time-varying treatments. Second, we propose an estimator that addresses the complexity of real-world data, enhancing performance through Double Machine Learning (DML) techniques to ensure unbiased and efficient estimation of treatment effects, thereby advancing survival analysis methods for handling time-varying treatments.

Literature Review

Prior work combining causal inference with survival analysis has primarily focused on static treatments or simple time-varying confounders. While (Robins, Rotnitzky, and Zhao 1992) and (Hern'an 2010) established foundational frameworks, and (Li and Greene 2015) developed doubly robust estimators, these approaches don't fully address random treatment timing. (Vansteelandt and Joffe 2014)'s work on time-varying treatments considers only scheduled interventions that can be reversed, unlike our setting with stochastic, irreversible treatments.

In econometrics, (Athey and Imbens 2022) and (Goodman-Bacon 2021) developed methods for handling random treatment timing in panel data, while (Sun and Abraham 2021) highlighted biases from ignored treatment effect heterogeneity. However, these approaches require repeated outcome measurements. Our work bridges this gap by adapting (Shaikh and Toulis 2021)'s Cox model framework, originally designed for continuous outcomes, to handle time-to-event data with staggered adoption patterns.

The organization of the paper is as follows: In Section 2, we formalize the notation, introduce assumptions for the causal framework, and present the statistical problem. In Section 3, we review existing methods for handling time-varying treatments in survival analysis. In Section 4, we introduce our double machine learning framework for robust estimation of heterogeneous treatment effects. In Section 5, we present simulation results demonstrating the performance of our method. In Section 6, we analyze the Stanford Heart Transplant dataset to evaluate treatment effect heterogeneity. Section 7 concludes with a discussion of our findings and limitations.

Problem Set Up

Notation and data

Let capital letters denote random variables and lowercase letters denote their realizations. Consider $i = 1, ..., N$ units. Let $T_i \in [0, \infty)$ denote the time until an event of interest occurs, such as time until mortality in transplant studies. Each unit i has a set of potential outcomes for T_i, denoted as $\{T_i(a) \in [0, \infty]\}$, where a represents the date (or time) when a binary treatment is first adopted by the unit. We refer to this as the adoption date, consistent with the terminology used in the staggered adoption literature (e.g., (Athey and Imbens 2022)). A unit can adopt the treatment at any of the time point $a \in [0, \infty)$, or not adopt the treatment at all during the time of observation, which we denote as $a = \infty$. We take a super-population perspective of $T_i(a)$, i.e., $T_i(a) \sim P$ are i.i.d. for some probability distribution P, the choice of which is discussed below. We observe for each unit in the population the adoption date $A_i \in [0, \infty]$. The observed event time of interest is denoted as T_i.

We also observe pre-treatment covariates $X_i \in \mathbb{R}^p$. We adapt the following standard causal assumptions from (Rubin 1974):

Assumption 1 (Stable Unit Treatment Value Assumption, SUTVA): Each unit's potential outcome is determined solely by its own treatment assignment, with no interference between units and uniform treatment versions. For each unit i,

$$T_i = T_i(A_i)$$

where the observed outcome equals the potential outcome under the assigned treatment. This assumption tells how potential outcomes map to observed outcomes.

Assumption 2 (Unconfoundedness): The treatment assignment is unconfounded, conditional on covariates X_i. Formally,

$$A_i \perp\!\!\!\perp T_i(a) \mid X_i$$

meaning that the treatment assignment A_i is independent of the potential outcome $T_i(a)$, given the covariates X_i.

Assumption 3 (Overlap): The probability of receiving treatment at time t, conditional on covariates, is strictly between 0 and 1 for all units. Specifically,

$$P(A_i \leq t \mid X_i) = a_t(X_i) \in [\epsilon, 1 - \epsilon] \quad \text{for some } \epsilon > 0.$$

This assumption ensures that each unit has a non-zero probability of receiving either treatment or control.

Distributional Assumption and Introduction of Hazard

For the triplet of covariates, treatment, and outcome (X_i, A_i, T_i), we impose the following general distributional assumptions:

$$\begin{aligned} X_i &\sim_{i.i.d.} f_X \\ A_i \mid X_i &\sim k(\cdot \mid X_i) \\ T_i(a) \mid X_i &\sim f(\cdot \mid a, X_i) \end{aligned}$$

Here, f_X represents the marginal density of the covariates on $\mathbb{R}^p$, without any additional parametric assumptions. The functions $k(\cdot \mid x)$ and $f(\cdot \mid a, X_i)$ denote conditional densities on $[0, \infty]$, corresponding to the treatment and outcome, respectively.

We also need to account for censoring. Censoring occurs when the event of interest—in this case, patient mortality—is not observed for all units within the study period. The censoring time, denoted as $C_i \in [0, \infty]$, represents the time at which the unit's data becomes unavailable for observation.

Censoring can arise for several reasons: a patient may be lost to follow-up, the study period may end before mortality is observed, or administrative reasons may prevent further observation. For censored units, the exact event time T_i is unknown; we only know that it exceeds the censoring time

C_i. To handle this, we introduce a binary indicator variable Δ_i, where $\Delta_i = 0$ indicates censored data and $\Delta_i = 1$ indicates fully observed event times. Thus, the observed data consists of covariates X_i, treatment adoption time A_i, observed time $U_i = T_i \wedge C_i$ (the minimum of event time and censoring time), and censoring indicator Δ_i. We can represent each unit's data as the tuple $(X_i, A_i, U_i, \Delta_i)$.

When censoring is present, we cannot directly estimate the distribution of event times. Instead of using the probability density function (pdf) $f(\cdot \mid a, X_i)$, we parameterize the distribution using the hazard function $h(\cdot \mid a, X_i)$ (Cox 1972; Kalbfleisch and Prentice 2011). The hazard function $h(t)$ at time t represents the instantaneous rate of event occurrence:

$$h(t) = \lim_{\Delta t \to 0} \frac{\Pr(t \leq T < t + \Delta t \mid T \geq t)}{\Delta t}.$$

The hazard function is particularly useful for censored data because it characterizes the risk of an event at time t, given survival up to that time.

Note that for each fixed x, the number of counterfactual hazard functions $h(\cdot \mid a, X_i)$ is infinite, as a is continuous on $[0, \infty]$. To address this complexity, we introduce two "exclusion" assumptions that simplify the model.

The first assumption states that the exact future transplant date doesn't affect current outcomes (Abbring and Van den Berg 2005; Abbring 2008):

Assumption 4: No Anticipation
For all i and for all adoption dates a such that $t < a$,

$$h(t \mid a, X_i) = h(t \mid \infty, X_i)$$

This reduces the infinite set of potential distributions for $t < a$ to a single one by assuming that before treatment adoption, the outcome event follows the baseline (or control) hazard f_0. In practical terms, this means that future treatment adoption does not influence current outcomes.

The second assumption asserts that conditional on treatment adoption, the magnitude of the adoption time does not matter for potential outcomes, but o This assumption is more restrictive but likely holds when a unit's characteristics X and event time T have a stable relationship that does not . However, this assumption might not hold when (X) and (T) have a dynamic relationship—for example, when the effectiveness of the transplant depends on how long the patient has had it, or when patient characteristics change significantly over time post-transplant.

Assumption 5: Invariance to History
For all i and for all adoption dates a such that $t \geq a$,

$$h(t \mid a, X_i) = h(t \mid 0, X_i)$$

Together, assumption 4 and 5 enable us to simplify to only two hazard functions—one for the control, $h_0(t \mid X_i) = h(t \mid \infty, X_i)$, and one for the treated, $h_1(t \mid X_i) = h(t \mid 0, X_i)$. These assumptions have been widely adopted in the Difference-in-Differences (DID) literature. For a comprehensive list of related works, see Section 3.2 in (Athey and Imbens 2022).

We can now represent the hazard for the potential outcome, $h(t|a, X_i)$ by the following:

$$h(t|a, X_i) = h_0(t|X_i) \cdot \left(\frac{h_1(t|X_i)}{h_0(t|X_i)}\right)^{w_t} \tag{1}$$

where $w_t := 1(t \geq a)$ is a binary function that indicates whether treatment has begun at time t.

We model the hazard using the proportional hazards model (Cox 1972):

$$\begin{aligned} h_0(t \mid x) &= \lambda(t) \exp(\eta_0(x)) \\ h_1(t \mid x) &= \lambda(t) \exp(\eta_1(x)) \end{aligned} \tag{2}$$

where $\lambda(t)$ represents the baseline hazard function, which captures the underlying risk of event at time t when all covariates are at their reference levels. The functions $\eta_0(x)$ and $\eta_1(x)$ can be either linear or non-linear functions of the covariates. A key advantage of this model is its ability to decouple the time component, $\lambda(t)$, from the covariate-dependent components, $\eta_0(x)$ and $\eta_1(x)$. This leads to a simplified treatment effect definition where the time component cancels out:

$$\begin{aligned} \tau(x) &= \log\left(\frac{h_1(t \mid x)}{h_0(t \mid x)}\right) \\ &= \log\left(h_1(t \mid x)\right) - \log\left(h_0(t \mid x)\right) \\ &= \eta_1(x) - \eta_0(x) \end{aligned} \tag{3}$$

We refer to this as the heterogeneous log hazard ratio (HLHR), which measures the treatment effect on the hazard rate for a patient with covariate profile x.

In this paper, we assume that the treatment effect, $\tau(x)$, follows a linear parametric form. While more flexible specifications are possible, we focus on this linear specification for several reasons. First, in the context of heart transplant studies, key patient characteristics like age, medical history, and physiological measures often have approximately linear relationships with treatment outcomes (Choudhry et al. 2019). Second, this specification mirrors the successful progression in the causal inference literature, where initial work on heterogeneous treatment effects for continuous outcomes began with linear models before expanding to more complex specifications (Imai and Ratkovic 2013; Kennedy 2023). Specifically, for some $\beta \in \mathbb{R}^p$, we model $\tau(x)$ as:

$$\tau(x) = \beta^T x \tag{4}$$

As a result, the hazard function can be expressed as:

$$h(t \mid a, x) = \lambda(t) \exp\left(\eta_0(x) + w_t \cdot \tau(x)\right) \tag{5}$$

where $\lambda(t)$ is the baseline hazard and $w_t = 1(t \geq a)$ is an indicator function that denotes whether the treatment has been adopted by time t.

Review of Handling Time-Varying Treatment in Survival Models

In this section, we review the incorporation of time-varying variables in survival models, as discussed in (Fisher and Lin 1999; Kalbfleisch and Prentice 2011).

Review of Partial Likelihood and Ordinary Cox Regression

We begin by revisiting the ordinary Cox regression model to guide the reader through the derivation of maximum partial likelihood. In this subsection, we assume treatment is fixed from the start. Under this assumption, the hazard function takes the form:

$$h(t|w,x) = \lambda(t)\exp(\eta_0(x) + w \cdot \tau(x)) \tag{6}$$

where $\eta_0(x)$ represents the control group log hazard as a function of covariates x, and w is the treatment indicator.

The partial likelihood (Cox 1972) is constructed by summing terms over the instances when an event (e.g., conversion) occurs, that is, when $\Delta_i = 1$ for a particular unit i. Let $\mathcal{R}_i = \{j : U_j \geq U_i\}$ denote the risk set for unit i, representing the set of individuals who have not yet experienced the event at time U_i. Furthermore, denote W_i as the treatment status for individual i. The log partial likelihood for the model is then expressed as follows:

$$\begin{aligned}\mathrm{pl}_n(\tau,\eta_0) :=& \log\left(\prod_{\Delta_i=1}\frac{h(T_i|W_i,X_i)}{\sum_{j\in\mathcal{R}_i}h(T_i|W_j,X_j)}\right)\\ =& \sum_{\Delta_i=1}\Bigg(\eta_0(X_i) + W_i\tau(X_i)\\ &- \log\left(\sum_{j\in\mathcal{R}_i}\exp(\eta_0(X_j) + W_j\tau(X_j))\right)\Bigg),\end{aligned} \tag{7}$$

To estimate the parameters η_0 and τ, we maximize this partial likelihood. This method is widely used in practice due to its robustness, as it does not require explicit specification of the baseline hazard function $\lambda(t)$, while still maintaining desirable statistical properties (Andersen et al. 1993).

Review of Handling Time-Varying Treatment

To accommodate time-varying treatment, we replace the fixed treatment indicator w with a time-varying indicator w_t:

$$h(t|a,x) = \lambda(t)\exp(\eta_0(x) + w_t\tau(x)) \tag{8}$$

where $w_t = 1(t \geq a)$ indicates whether the treatment has been initiated by time t.

To address the issue of time-varying covariates, we incorporate time variation into the partial likelihood framework (Fisher and Lin 1999; Kalbfleisch and Prentice 2011). This extension allows for the correct modeling of covariates that change over time.

We retain the risk set $\mathcal{R}_i$ as defined previously. Let $W_i(t) = 1(A_i < t)$ represent the treatment status of unit i at time t. The partial likelihood for this time-varying model is then expressed as:

$$\begin{aligned}\mathrm{pl}_n(\tau,\eta_0) :=& \frac{1}{n}\sum_{\Delta_i=1}\Bigg(\eta_0(X_i) + W_i(U_i)\tau(X_i)\\ &- \log\left(\sum_{j\in\mathcal{R}_i}\exp\left(\eta_0(X_j) + W_j(U_i)\tau(X_j)\right)\right)\Bigg),\end{aligned} \tag{9}$$

This approach of incorporating time-varying treatments within the partial likelihood framework maintains desirable statistical properties and produces consistent estimators (Kalbfleisch and Prentice 2011). Recent work by (Tay, Narasimhan, and Hastie 2023) has extended this framework to allow lasso fitting when η_0 is non-linear.

Double Machine Learning Estimator

We now focus on efficient estimation of the treatment effect function $\tau(x)$ in the time-varying Cox proportional hazards model specified in Equation 6. As mentioned above, a straightforward approach would be jointly estimating $\tau(x)$ and the baseline hazard η_0 using the pseudo-likelihood in Equation 9.

However, this direct approach faces significant challenges because the baseline hazard η_0 acts as a nuisance function, and its estimation can interfere with the consistent estimation of the treatment effect. Traditional outcome-based methods that rely on correctly specifying the outcome model (in this case, the hazard function) are particularly vulnerable in this setting. This vulnerability arises because misspecification of the baseline hazard can directly bias the treatment effect estimates through the partial likelihood structure—errors in estimating η_0 propagate non-linearly through the risk set calculations, leading to biased estimates of $\tau(x)$.

To address these challenges, we adopt a double machine learning (DML) framework (Chernozhukov et al. 2018; Künzel et al. 2019; Nie and Wager 2021; Gao and Hastie 2021). The key insight of DML is to introduce propensity score estimation alongside the outcome model, providing double robustness and improved convergence rates.

Under the DML framework, the treatment effect estimator achieves favorable convergence rates through the product of nuisance parameter estimation errors. More precisely, let $e_0(x) = a_t(x)$ be our time-varying propensity score as defined in Equation 10, and let $\eta_0(x)$ be our baseline log hazard function as defined in Equation 6. If we denote the L2 convergence rates of their estimators as $\|\hat{e} - e_0\|_2 = O_p(r_n^e)$ and $\|\hat{\eta}_0 - \eta_0\|_2 = O_p(r_n^\eta)$ respectively, then through careful orthogonalization of the score function, the treatment effect estimator satisfies:

$$\|\hat{\tau} - \tau_0\|_2 = O_p(r_n^e \cdot r_n^\eta + \frac{1}{\sqrt{n}})$$

This product structure is crucial: even if one nuisance component converges at a slower rate (e.g., $r_n^e = n^{-1/4}$), the treatment effect estimator can still achieve the optimal $\sqrt{n}$-rate of convergence as long as the other component converges sufficiently fast (e.g., $r_n^\eta = n^{-1/4}$). This property,

known as rate double robustness, makes the estimator robust to moderate misspecification of either nuisance component. For a comprehensive theoretical analysis of these convergence properties in the general framework of orthogonal statistical learning, we refer readers to (Foster and Syrgkanis 2023).

Time-Varying Causal Survival Learner (TV-CSL)

Building upon the work of (Gao and Hastie 2021), we propose TV-CSL (Time-Varying Causal Survival Learner) to handle time-varying treatments. The model is characterized by:

$$\begin{aligned} a_t(x) &= \frac{\int_0^t \mathbb{P}(\Delta = 1 \mid A = s, X) f(A = s|X) ds}{\int_0^\infty \mathbb{P}(\Delta = 1 \mid A = s, X) f(A = s|X) ds} \\ &= \mathbb{P}(A \leq t | \Delta = 1, X)., \\ \nu_t(x) &= \tau(x) \cdot a_t(x) + \eta_0(x) \end{aligned} \tag{10}$$

Here, $a_t(x) = \mathbb{P}(A \leq t \mid \Delta = 1, X)$ represents the probability of adoption by time t for non-censored data. This is analogous to the "treatment probability" or the propensity score at time t. When all data is not censored, $a_t(x) = \mathbb{P}(A \leq t \mid X)$. The full estimation procedure is presented in Algorithm 1.

Our work differs from (Gao and Hastie 2021), which developed a DML method for linear heterogeneous effects under a Cox model with treatment fixed at baseline, in two key aspects. First, the outcome models differ: their work uses the hazard form in Equation 6, while we use the time-varying form in Equation 8. Second, the propensity scores are distinct: their nuisance function maps covariates x to probabilities in $(0, 1)$, whereas our propensity score is a function of both x and t, representing the cumulative distribution of adoption time A conditional on X.

Theoretical justification

Similar to existing causal inference literature (Nie and Wager 2021; Künzel et al. 2019), we can derive theoretical results for the reduction of the learning rate.

Proposition 1 (Convergence Rate of Parameter Estimation)**.** *Let the model for $a_t(x)$ be denoted as $\gamma(x)$. Under the following regularity conditions:*

1. *The covariates X are bounded, the true parameter β_0 lies in a bounded region $\mathcal{B}$, and the nuisance functions $\gamma_0(x)$, $\eta_0(x)$ along with their estimators $\gamma_n(x)$, $\eta_n(x)$ are uniformly bounded;*
2. *The minimal eigenvalues of the score derivative $\nabla_\beta s(\gamma(x), \eta(x), \beta)$ [1] in $\mathcal{B}$ are lower bounded by some constant $C > 0$;*

If $\|\gamma_n(x) - \gamma_0(x)\|_2 = O(\alpha_n)$, $\|\eta_n(x) - \eta_0(x)\|_2 = O(\rho_n)$, and $\alpha_n \to 0$, $\rho_n \to 0$, then

$$\|\beta_n - \beta_0\|_2 = \tilde{O}\left(\alpha_n \rho_n + n^{-1/2}\right) \tag{11}$$

[1] See the Appendix for a definition of the score

Algorithm 1 Cox Model with Partial Likelihood for Time-Varying Treatment (Under No Censoring)

Input: Dataset $\{(X_i, T_i, \Delta_i, A_i)\}_{i=1}^n$, where X_i are covariates, T_i are survival times, Δ_i are event indicators, and A_i are treatment adoption dates
First Stage (Fold One):
  Estimate propensity score $a_t(x) = P(A_i \leq t | X_i = x, \Delta_i = 1)$
  Estimate nuisance function ν_t by:
    1. Maximizing the partial likelihood (Equation 9) to obtain $\hat{\eta}_0(x)$ and $\hat{\tau}(x)$
    2. Computing $\hat{\nu}_t(x) = \hat{\tau}(x) \cdot \hat{a}_t(x) + \hat{\eta}_0(x)$
Second Stage (Fold Two):
  Estimate treatment effect $\tau(x)$ by solving:

$$\hat{\beta} = \min_{\beta'} \frac{1}{n} \sum_{\Delta_i = 1} \Big[\hat{\nu}_{\tau_i}(X_i) + (W_i(\tau_i) - \hat{a}_{\tau_i}(X_i)) X_i^\top \beta' - \log\Big(\sum_{l \in \mathcal{R}_i} \exp(\hat{\nu}_{\tau_i}(X_l) + (W_l(\tau_i) - \hat{a}_{\tau_i}(X_l)) X_l^\top \beta') \Big) \Big]$$

where:
  $W_i(t) = \mathbf{1}(A_i < t)$ ▷ Treatment status at time t
  $\mathcal{R}_i = \{j : U_j \geq U_i\}$ ▷ Risk set for subject i
  $\tau_i = U_i$ for i where $\Delta_i = 1$ ▷ Event times
Output: Estimated treatment effect function $\hat{\tau}(x) = x^T \hat{\beta}$

Proof: See Appendix.

Proposition 1 states that for $\hat{\tau}(x) = x^\top \hat{\beta}$ to reach a certain level of accuracy, the conditions on $\hat{a}_t(x)$ and $\hat{\eta}_0(x)$ are relatively loose. Specifically, if we want the estimate of β to achieve $n^{-1/2}$ convergence, we only need the product of the convergence rates of the outcome model $\eta_0(x)$ and treatment model $a_t(x)$ to be $n^{-1/2}$. For example, $\eta_0(x)$ could converge at rate $n^{-1/4}$ and $a_t(x)$ at rate $n^{-1/4}$. In comparison, outcome-based methods that do not use the treatment model $a_t(x)$ can only achieve $n^{-1/4}$ convergence rate.

Simulation Study

We evaluate our method's performance through simulation studies comparing TV-CSL against existing approaches.

Simulation Design

Let $X_i = (X_{i1}, X_{i2}, X_{i3})^\top \in \mathbb{R}^3$ denote baseline covariates generated from $X_i \sim \mathcal{N}(\mathbf{0}, I_3)$. Treatment times follow $A_i \mid X_i \sim \text{Exp}(X_{i2} + X_{i3})$. Survival times are generated through the hazard function:

$$h(t \mid a, x) = t \cdot \exp(\eta_0(x) + 1(a \leq t)\tau(x))$$

where $\eta_0(x) = -\frac{1}{2} \cdot \varsigma(X_1) \cdot \varsigma(X_{10})$ with scaled sigmoid $\varsigma(x) = \frac{2}{1+e^{-12(x-\frac{1}{2})}}$, and treatment effect $\tau(x) = x_1 + x_2 + x_3$. We implement random censoring with $C_i = \min(20, \tilde{C}_i)$, where $\tilde{C}_i \sim \text{Exp}(0.1)$, yielding 75% non-censored observations. The simulation was repeated 100 times for each scenario, with sample sizes $n \in \{200, 500, 1000, 2000\}$.

Methods for Comparison

We evaluate two approaches for estimating heterogeneous treatment effects:

S-Lasso Method: Employs a single regression combining baseline risk η_0 and treatment effect τ through additive specification. For both components, we consider:

- Linear specification: $\eta_0(x) = \beta_1 X_1 + \beta_2 X_2 + \beta_3 X_3$ and $\tau(x) = \omega_0 W + \sum_{j=1}^{3} \omega_j (W \cdot X_j)$
- Complex specification: Includes natural splines, squared terms, and all pairwise interactions

The combined model is fit using Lasso regularization with cross-validated penalty parameter.

TV-CSL Method: Implements doubly robust estimation through cross-fitting in two stages:

- First stage estimates treatment model $a_t(x) = P(A \leq t \mid X = x)$ under both correct (all covariates) and misspecified (single covariate) settings
- Second stage estimates baseline outcome model using Lasso with specifications matching S-Lasso

Performance is evaluated using Mean Squared Error (MSE):

$$\text{MSE} = \frac{1}{n} \sum_{i=1}^{n} (\hat{\tau}(X_i) - \tau(X_i))^2$$

Results

Our simulation results demonstrate the relative performance of TV-CSL and S-Lasso under various specifications, focusing on estimation accuracy and robustness to model misspecification. We examine two key aspects: the impact of the treatment model specification and the performance under complex treatment effect specifications.

Effect of Treatment Model Specification We first examine how the treatment model specification affects estimation quality. To isolate this effect, we maintain a correctly specified HTE model to ensure optimal conditions for both methods.

The results are shown in Figure 1. Both methods achieve lower MSE with the complex baseline outcome specification, which aligns with the true data-generating process where the baseline hazard (η_0) follows a non-linear pattern.

Comparing TV-CSL and S-Lasso, we observe that TV-CSL consistently outperforms S-Lasso regardless of treatment model specification, with the performance advantage being more pronounced when the treatment model is correctly specified. This aligns with our theoretical findings that the convergence rate depends on the product of errors in the nuisance estimators.

Notably, TV-CSL maintains its advantage over S-Lasso even when the treatment model is misspecified. This robustness can be attributed to the simplicity of our treatment model, where minor misspecifications have limited impact on the overall estimation error. The double machine learning framework effectively mitigates the impact of treatment model misspecification, allowing TV-CSL to maintain robust performance.

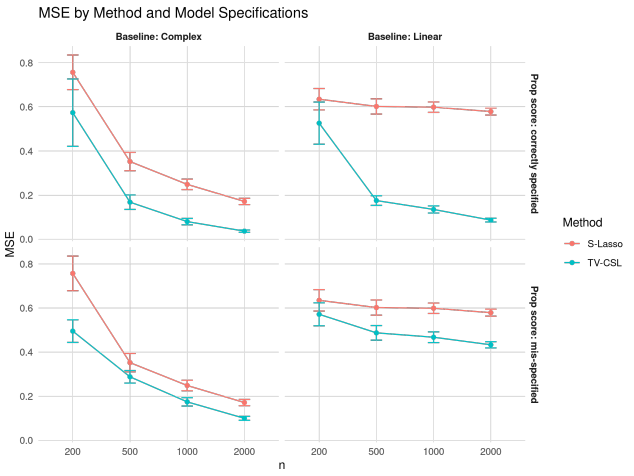

Figure 1: MSE Comparison between S-Lasso and TV-CSL Methods by Sample Size. Panels show combinations of baseline hazard (η_0) and propensity score specifications. Error bars: ± 1.96 Monte Carlo SE.

Performance Under Complex Treatment Effect Specifications While our previous analysis focused on a linear (correctly specified) HTE model, we now evaluate the performance when using a complex model to estimate the HTE. Figure 2 presents these results.

For both methods, holding the outcome model fixed, the use of complex HTE specifications leads to higher MSE, though the magnitude of this increase varies between methods. This increased error can be attributed to the additional complexity in estimating the treatment effect model.

TV-CSL demonstrates superior performance relative to S-Lasso under two conditions: First, when the HTE model is correctly specified as linear, TV-CSL consistently outperforms S-Lasso across all sample sizes. This advantage stems from the double machine learning framework's ability to reduce the impact of nuisance parameter estimation errors. Second, for complex HTE specifications, TV-CSL's performance shows strong sample size dependency. While maintaining comparable performance at smaller sample sizes, TV-CSL outperforms S-Lasso at larger sample sizes.

Data Analysis – Stanford Heart Transplant Dataset

We analyze data from the Stanford Heart Transplant Program, which tracks survival times of 103 patients from program acceptance through transplantation. The dataset includes patient age at enrollment, previous surgery status (binary), and enrollment year (measured from study initiation in 1967). The analysis focuses on heterogeneous treatment effects (HTE) by examining interactions between transplant status and patient characteristics.

Time-Varying Treatment Effects

We first examine the impact of incorporating time-varying information when evaluating transplant effects compared to

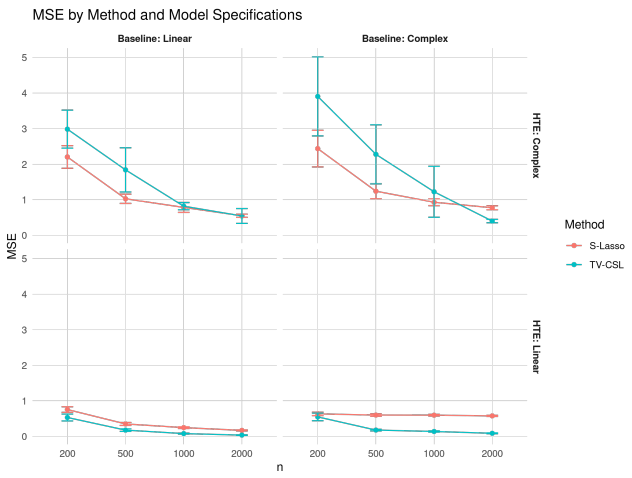

Figure 2: MSE Comparison between S-Lasso and TV-CSL Methods by Sample Size. Panels show combinations of baseline hazard (η_0) and HTE specifications. Error bars: ± 1.96 Monte Carlo SE.

Variable	Ignore Treatment Time		Include Treatment Time	
	Coef (SE)	P-value	Coef (SE)	P-value
Trt	-1.504 (0.292)	0.00	0.117 (0.340)	0.73
Age $\times$ Trt	-0.259 (0.285)	0.36	0.286 (0.254)	0.26
Surgery $\times$ Trt	-2.191 (0.778)	0.00	-0.557 (0.777)	0.47
Year $\times$ Trt	0.206 (0.261)	0.43	0.421 (0.260)	0.11

Table 1: Comparison of Heterogeneous Treatment Effects Between Models With and Without Time-Varying Information

not incorporating it. We compare two Cox proportional hazards models: one treating transplant as a fixed treatment covariate and another incorporating the time-varying nature of waiting time.

Table 1 shows differences between the two models. In the model ignoring treatment time, the surgery-treatment interaction shows a significant effect (coef $= -2.191$, p-value < 0.01), suggesting a transplant benefit for patients who have had previous surgery. However, this effect disappears in the model that includes treatment time (coef $= -0.557$, p-value $= 0.47$). This indicates bias in the model that ignores treatment timing. The treatment variable itself also shows this difference. Other interactions remain non-significant in both models.

Machine Learning Effects

We next compare single-fit S-lasso with our TV-CSL method for predicting treatment effects. We excluded the surgery variable due to its extreme class imbalance (76% true-to-false ratio), as this imbalance would be further exacerbated after splitting the data into two datasets in cross-fitting. Following simulation insights, we use a linear model for η_0 given the small sample size (n=103). The treatment model is specified as $A|X \sim Exp(\alpha_0 + \alpha_1 Years)$, reflecting that waiting times primarily depend on enrollment year rather than patient characteristics.

	Baseline Outcome Model	
Method	complex	linear
S-lasso	0.386	0.492
TV-CSL	1.220	1.150

Table 2: MSE by Method and Baseline Outcome Model

Results in Table 2 show that with this small sample size, single-fit methods outperform TV-CSL across all specifications. These findings align with our simulation results regarding sample size sensitivity and suggest that practitioners should prefer simpler models for smaller datasets.

Conclusion

In this paper, we propose a novel framework for estimating causal effects of time-varying treatments on time-to-event outcomes by extending the staggered adoption framework from econometrics to a survival analysis setting. Our approach leverages the Cox proportional hazards model and incorporates double machine learning (DML) to address complexities in real-world data, such as nonlinear covariate relationships and high-dimensional settings. Through simulations, we demonstrate that our estimator effectively reduces bias and improves efficiency compared to traditional methods, particularly in cases with significant treatment effect heterogeneity.

Our proposed estimator advances the capabilities of causal inference in survival analysis, providing a robust approach for analyzing staggered treatment adoption with time-varying interventions in diverse applied contexts.

Limitations and Future Work

Our study has several limitations that warrant discussion. A key limitation is our assumption of linear treatment effect heterogeneity. Although this specification allows substantial methodological progress and provides interpretable results for medical decision-making, treatment effects in complex medical interventions may exhibit non-linear patterns across patient characteristics. Future work could extend our framework to accommodate more flexible specifications of $\tau(x)$ using non-parametric or semi-parametric approaches.

Another limitation is that in reality, there is an instantaneous increase in risk immediately after heart transplant. This occurs because some patients may experience severe rejection reactions when receiving a new heart (Lipkova et al. 2022). After surviving this critical period, the patient's risk typically decreases. In future work, rather than modeling a single hazard function post-transplant as we did, we should consider two distinct hazard functions: one capturing the elevated risk immediately after transplant, and another reflecting the lower risk level that follows successful adaptation.

References

Abbring, J. H. 2008. The event-history approach to program evaluation. In *Modelling and Evaluating Treatment*

Effects in Econometrics, 33–55. Emerald Group Publishing Limited.

Abbring, J. H.; and Van den Berg, G. J. 2005. Social experiments and instrumental variables with duration outcomes.

Almond, C. S.; Thiagarajan, R. R.; Piercey, G. E.; Gauvreau, K.; Blume, E. D.; Bastardi, H. J.; Fynn-Thompson, F.; and Singh, T. 2009. Waiting list mortality among children listed for heart transplantation in the United States. *Circulation*, 119(5): 717–727.

Andersen, P. K.; Borgan, Ø.; Gill, R. D.; and Keiding, N. 1993. *Statistical Models Based on Counting Processes*. Springer Series in Statistics. New York, NY: Springer US. ISBN 978-0-387-94519-4, 978-1-4612-4348-9.

Athey, S.; and Imbens, G. W. 2022. Design-based analysis in difference-in-differences settings with staggered adoption. *Journal of Econometrics*, 226(1): 62–79.

Chernozhukov, V.; Chetverikov, D.; Demirer, M.; Duflo, E.; Hansen, C.; Newey, W.; and Robins, J. 2018. Double/debiased machine learning for treatment and structural parameters.

Choudhry, S.; Wang, Y.; Denfield, S. W.; Cabrera, A. G.; Price, J. F.; Tunuguntla, H. P.; Dharnidharka, V. R.; Canter, C. E.; and Dreyer, W. J. 2019. A recipient risk prediction tool for short-term mortality after pediatric heart transplantation. *Transplantation*, 103(11): 2434–2439.

Cox, D. R. 1972. Regression models and life-tables. *Journal of the Royal Statistical Society: Series B (Methodological)*, 34(2): 187–202.

Crowley, J.; and Hu, M. 1977. Covariance analysis of heart transplant survival data. *Journal of the American Statistical Association*, 72(357): 27–36.

DeFilippis, E. M.; Khush, K. K.; Farr, M. A.; Fiedler, A.; Kilic, A.; and Givertz, M. M. 2022. Evolving characteristics of heart transplantation donors and recipients: JACC focus seminar. *Journal of the American College of Cardiology*, 79(11): 1108–1123.

Evans, R. W.; Manninen, D. L.; Garrison, L. P.; and Maier, A. M. 1986. Donor availability as the primary determinant of the future of heart transplantation. *Jama*, 255(14): 1892–1898.

Fisher, L. D.; and Lin, D. Y. 1999. Time-Dependent Covariates in the Cox Proportional-Hazards Regression Model. *Annual Review of Public Health*, 20(1): 145–157.

Fleming, T. R.; and Harrington, D. P. 2005. *Counting processes and survival analysis*. Wiley series in probability and statistics. New York: Wiley-Interscience. ISBN 978-0-471-76988-0.

Foster, D. J.; and Syrgkanis, V. 2023. Orthogonal Statistical Learning. ArXiv:1901.09036 [cs, econ, math, stat].

Gao, Z.; and Hastie, T. 2021. Estimating heterogeneous treatment effects for general responses. *arXiv preprint arXiv:2103.04277*.

Goodman-Bacon, A. 2021. Difference-in-differences with variation in treatment timing. *Journal of econometrics*, 225(2): 254–277.

Hastie, T.; Tibshirani, R.; Friedman, J. H.; and Friedman, J. H. 2009. *The elements of statistical learning: data mining, inference, and prediction*, volume 2. Springer.

Hern'an, M. A. 2010. The hazards of hazard ratios. *Epidemiology*, 21(1): 13–15.

Imai, K.; and Ratkovic, M. 2013. Estimating treatment effect heterogeneity in randomized program evaluation.

Kalbfleisch, J. D.; and Prentice, R. L. 2011. *The statistical analysis of failure time data*. John Wiley & Sons.

Kennedy, E. H. 2023. Towards optimal doubly robust estimation of heterogeneous causal effects. ArXiv:2004.14497 [math, stat].

Kilic, A.; Mathier, M. A.; Hickey, G. W.; Sultan, I.; Morell, V. O.; Mulukutla, S. R.; and Keebler, M. E. 2021. Evolving trends in adult heart transplant with the 2018 heart allocation policy change. *JAMA cardiology*, 6(2): 159–167.

Klein, J. P.; and Moeschberger, M. L. 1997. *Survival Analysis*. Statistics for Biology and Health. New York, NY: Springer New York. ISBN 978-1-4757-2730-2 978-1-4757-2728-9.

Künzel, S. R.; Sekhon, J. S.; Bickel, P. J.; and Yu, B. 2019. Metalearners for estimating heterogeneous treatment effects using machine learning. *Proceedings of the National Academy of Sciences*, 116(10): 4156–4165.

Li, L.; and Greene, T. 2015. Evaluating the effect of health care interventions using doubly robust estimation. *Statistics in Medicine*, 34(3): 539–559.

Lipkova, J.; Chen, T. Y.; Lu, M. Y.; Chen, R. J.; Shady, M.; Williams, M.; Wang, J.; Noor, Z.; Mitchell, R. N.; Turan, M.; et al. 2022. Deep learning-enabled assessment of cardiac allograft rejection from endomyocardial biopsies. *Nature medicine*, 28(3): 575–582.

Nie, X.; and Wager, S. 2021. Quasi-oracle estimation of heterogeneous treatment effects. *Biometrika*, 108(2): 299–319.

O'Connell, J. B.; Bourge, R. C.; Costanzo-Nordin, M. R.; Driscoll, D. J.; Morgan, J. P.; Rose, E.; and Uretsky, B. 1992. Cardiac transplantation: recipient selection, donor procurement, and medical follow-up. A statement for health professionals from the Committee on Cardiac Transplantation of the Council on Clinical Cardiology, American Heart Association. *Circulation*, 86(3): 1061–1079.

Robins, J. M.; Rotnitzky, A.; and Zhao, L. P. 1992. Estimation of regression coefficients when some regressors are not always observed. *Journal of the American Statistical Association*, 87(420): 1013–1022.

Rubin, D. B. 1974. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of educational Psychology*, 66(5): 688.

Shaikh, A. M.; and Toulis, P. 2021. Randomization tests in observational studies with staggered adoption of treatment. *Journal of the American Statistical Association*, 116(536): 1835–1848.

Sun, L.; and Abraham, S. 2021. Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of econometrics*, 225(2): 175–199.

Tay, J. K.; Narasimhan, B.; and Hastie, T. 2023. Elastic Net Regularization Paths for All Generalized Linear Models. *Journal of Statistical Software*, 106(1): 1–31.

Trulock, E. P.; Christie, J. D.; Edwards, L. B.; Boucek, M. M.; Aurora, P.; Taylor, D. O.; Dobbels, F.; Rahmel, A. O.; Keck, B. M.; and Hertz, M. I. 2007. Registry of the International Society for Heart and Lung Transplantation: twenty-fourth official adult lung and heart–lung transplantation report—2007. *The Journal of heart and lung transplantation*, 26(8): 782–795.

Van der Vaart, A. W. 2000. *Asymptotic statistics*, volume 3. Cambridge university press.

Vansteelandt, S.; and Joffe, M. 2014. Structural Nested Models and G-estimation: The Partially Realized Promise. *Statistical Science*, 29(4).

Zhu, Y.; Lingala, B.; Baiocchi, M.; Toro Arana, V.; Williams, K. M.; Shudo, Y.; Oyer, P. E.; and Woo, Y. J. 2021. The Stanford experience of heart transplantation over five decades. *European Heart Journal*, 42(48): 4934–4943.

Derivation of the TV-CSL Estimator

The main technique involves expanding the log-likelihood of the data around the true parameter to obtain a score that approximates the true parameter at the fastest possible rate. See the next section for a derivation of the likelihood expansion. A key result from that section is the estimating equation resulting from the optimal score:

$$0 = \mathbb{E}\left[(1(A \leq t) - a_t(X))X\left(\Delta - \Lambda\left(U\right)e^{\eta_0^*(X)+\tau 1(A\leq t)}\right)\right] \tag{12}$$

Use the tower property to take $E[\cdot|A, X]$ inside, we obtain

$$\begin{aligned} 0 = \mathbb{E}\Bigg[&(1(A \leq t) - a_t(X))X \\ &\left(1 - e^{\eta_A^*(X)-\eta_{a_t}(X)}\right)\mathbb{P}(\Delta \mid A, X)\Bigg] \end{aligned} \tag{13}$$

where

$$\begin{aligned} \mathbb{P}(\Delta \mid A, X) &= \int_0^A \left(1 - \exp(-\Lambda(c)e^{\eta_0^*(X)})\right) f_C(c \mid A, X)dc \\ &\quad + \int_A^\infty \Bigg(1 - \exp(-\Lambda(A)e^{\eta_0^*(X)} - \\ &\quad (\Lambda(c) - \Lambda(A))e^{\eta_0^*(X)+\tau})\Bigg) f_C(c \mid A, X)dc \\ &= \int_0^A \left(1 - \exp(-\Lambda(c)e^{\eta_0^*(X)})\right) f_C(c \mid A, X)dc \\ &\quad + \int_A^\infty \Bigg(1 - \exp([-\Lambda(A)(1 - e^{\tau}) - \\ &\quad \Lambda(c)e^{\tau}]e^{\eta_0^*(X)}\Bigg) f_C(c \mid A, X)dc \end{aligned}$$

Note that $\eta_W^*(X) - \eta_W(X) = \nu^*(X) - \nu(X)$. Taking $E[\cdot|X]$ inside, we obtain

$$\begin{aligned} 0 = \mathbb{E}[\{&\int_0^t (1 - a_t(X))\mathbb{P}(\Delta \mid A = s, X)f(A = s|X)ds + \\ &\int_t^\infty (0 - a_t(X))\mathbb{P}(\Delta \mid A = s, X)f(A = s|X)ds\} \\ &X \cdot \left(1 - e^{\nu^*(X)-\nu(X)}\right)] \end{aligned} \tag{14}$$

Hence we need

$$\begin{aligned} &(1 - a_t(X)) \int_0^t \mathbb{P}(\Delta = 1 \mid A = s, X)f(A = s|X)ds \\ &- a_t(X) \int_t^\infty \mathbb{P}(\Delta = 1 \mid A = s, X)f(A = s|X)ds = 0 \\ &\frac{a_t(X)}{1 - a_t(X)} = \frac{\int_0^t \mathbb{P}(\Delta = 1 \mid A = s, X)f(A = s|X)ds}{\int_t^\infty \mathbb{P}(\Delta = 1 \mid A = s, X)f(A = s|X)ds} \end{aligned} \tag{15}$$

Here, $\int_0^\infty \mathbb{P}(\Delta = 1 \mid A = s, X)f(A = s|X)ds = P(\Delta = 1 \mid X)$ is the marginal censoring probability.

Hence, $a_t(X) = \frac{\int_0^t \mathbb{P}(\Delta = 1 \mid A = s, X)f(A = s|X)ds}{\int_0^\infty \mathbb{P}(\Delta = 1 \mid A = s, X)f(A = s|X)ds} = \frac{\int_0^t \mathbb{P}(\Delta = 1, A = s|X)ds}{\mathbb{P}(\Delta = 1 \mid X)} = \int_0^t \mathbb{P}(A = s|\Delta = 1, X)ds = \mathbb{P}(A \leq t|\Delta = 1, X)$. When $\mathbb{P}(\Delta = 1 \mid A = s, X) = 1$ for all s, i.e., when all observations are not-censored, then $a_t(X) = \mathbb{P}(A \leq t|X)$ is the treated probability at time t for a unit with covariate X.

Key Lemma for Deriving the Score Function

Here we derive the key lemma that shows how the estimator is obtained. The key step is to calculate the expansion of the

likelihood. Let $\ell(Y;\eta')$ denote the log-likelihood of the exponential family. Lemma 1 of (Gao and Hastie 2021) states that for arbitrary η', the likelihood of Y satisfies

$$\ell(Y;\eta') = \ell(Y;\eta) - \frac{1}{2}\psi''(\eta)\left(r+\eta-\eta'\right)^2 + \frac{1}{2}\psi''(\eta)r^2 + O\left(\|\eta-\eta'\|_2^3\right),$$

where $r := (Y-\psi'(\eta))/\psi''(\eta)$.

Further

$$\ell(Y;\eta') = \ell(Y;\eta) - \frac{1}{2}\psi''(\eta)\left(r+\eta-\eta'\right)^2 + \frac{1}{2}\psi''(\eta)r^2 \tag{16}$$

The key insight is that we parametrize $\eta_w(x) = \nu(x) + (w-a(x))\tau$, where $\nu(x) = a(x)\tau + \eta_0(x)$. Instead of parametrizing $\eta_w(x)$ using $\eta_0(x)$ and τ, we re-parametrize it by adding and subtracting $a(x)\tau$ to obtain double robustness.

When comparing η and η', we keep $\nu(x)$ and $a(x)$ fixed, choosing $\eta' = \nu(x) + (w-a(x))\tau'$ and $\eta = \nu(x) + (w-a(x))\tau$. This implies that $\eta' - \eta = (w-a(x))(\tau'-\tau)$. We can then apply this to Equation 16:

$$\begin{aligned}&\ell(Y;\eta') - \ell(Y;\eta)\\ =&-\frac{1}{2}\psi''(\eta)\left(r+\eta-\eta'\right)^2 + \frac{1}{2}\psi''(\eta)r^2\\ =&-\frac{1}{2}\psi''(\eta)\left(r-(w-a(x))(\tau'-\tau)\right)^2 + \frac{1}{2}\psi''(\eta)r^2\end{aligned} \tag{17}$$

Move the negative sign, we have

$$\begin{aligned}&\ell(Y;\eta) - \ell(Y;\eta')\\ =&\frac{1}{2}\psi''(\eta)\left(r+(w-a(x))(\tau-\tau')\right)^2 + \frac{1}{2}\psi''(\eta)r^2\end{aligned} \tag{18}$$

We take the expectation, and maximize the LHS by differentiating τ, and set the derivactive to zero (this is to find the tangent direction)

$$0 = E\left[(w-a(x))\psi''(\eta)\left(r+(w-a(x))(\tau-\tau')\right)\right] \tag{19}$$

This says

$$\begin{aligned}E\left[(w-a(x))\psi''(\eta)r\right] =& E\left[(w-a(x))^2\psi''(\eta)\right](\tau'-\tau)\\ (\tau'-\tau) =& E\left[(w-a(x))\psi''(\eta)r\right]\\ &/E\left[(w-a(x))^2\psi''(\eta)\right]\end{aligned} \tag{20}$$

We also want $\tau' = \tau$. This is because we are taking one-dimentional efficient scores so we finally need tangency. The condition is that the numerator is zero, i.e.

$$E\left[(w-a(x))\psi''(\eta)r\right] \quad \text{for all } \eta \tag{21}$$

Plugin $r := (Y-\psi'(\eta))/\psi''(\eta)$, and taking x, w as random variables, we have

$$\begin{aligned}&E\left[(W-a(X))\left(Y-\psi'(\eta)\right)\right]\\ =&E\left[(W-a(X))\left(Y-(\psi'(W\eta_1(X)+(1-W)\eta_0(X)))\right)\right]\\ =&E\Bigg[(e(X)-a(X))\Bigg(Y-(e(X)\psi'(\eta_1(X))\\ &+(1-e(X))\psi'(\eta_0(X))\Bigg)\Bigg]\end{aligned}$$

Proof of Proposition 1

Proof: We write $\gamma_n(x) = \gamma(x) + \alpha_n\xi_n(x)$, where $\mathbb{E}\left[\xi_n^2(X)\right] = 1$ is a unit directional vector, and α_n is the distance from $\gamma_n(x)$ to $\gamma(x)$. Similarly, we can write $\eta_n(x) = \eta(x) + \rho_n\zeta_n(x)$, where $\mathbb{E}\left[\zeta_n^2(X)\right] = 1$. By the assumption of proposition 1, $\alpha_n \to 0$, $\rho_n \to 0$.

The score function for the partial likelihood of the i-th sample is:

$$\begin{aligned}&S_i(\gamma,\eta,\beta) = S_i(\gamma,\nu,\beta) \quad \text{since } \nu = \eta + \tau\cdot a\\ :=&s\left(\gamma(X_i),\nu(X_i),\beta\right)\\ =&\frac{\partial}{\partial\beta}\Big[\nu(X_i) + (W_i - a(X_i))X_i^\top\beta\\ &-\log\Big(\sum_{l\in\mathcal{R}_i}\exp(\nu(X_l) + (W_l - \hat{a}(X_l))X_l^\top\beta)\Big)\Big]\\ =&Z_i - \frac{\sum_{l\in\mathcal{R}_i} Z_l\exp(\nu(X_l)+Z_l^\top\beta)}{\sum_{l\in\mathcal{R}_i}\exp(\nu(X_l)+Z_l^\top\beta)}\end{aligned}$$

where $Z_i := (W_i - a(X_i))X_i$.

Denote the expected score as $s(\gamma,\nu,\beta) = E[S_i(\gamma,\nu,\beta)]$ and define the empirical score $s_n(\gamma,\nu,\beta) = \frac{1}{n}\sum_{i=1}^n S_i(\gamma,\nu,\beta)$. For simplicity, we write $s_n(\gamma_n,\nu_n,\beta_n)$ as $s_n(\alpha_n,\rho_n,\beta_n)$

We first show β_n is consistent under $s_n(0,0,\beta)$. Taylor's expansion of $s_n(\alpha_n,\rho_n,\beta_n)$ at $\alpha_n = \rho_n = 0$ is

$$\begin{aligned}&s_n(\alpha_n,\rho_n,\beta_n)\\ =&s_n(0,0,\beta_n) + \nabla_\alpha s_n(\alpha_\varepsilon,\rho_\varepsilon,\beta_n)\alpha_n + \nabla_\rho s_n(\alpha_\varepsilon,\rho_\varepsilon,\beta_n)\rho_n\end{aligned}$$

where $\alpha_\varepsilon \in [0,\alpha_n], \rho_\varepsilon \in [0,\rho_n]$.

Note that $s(0,0,\beta_0) = 0$ (See (Fleming and Harrington 2005), Chapter 8 for a proof). Thus $s_n(0,0,\beta_n) = s_n(0,0,\beta_n) - 0 = s_n(0,0,\beta_n) - s(0,0,\beta_0)$. We now argue $s_n(0,0,\beta_n) - s(0,0,\beta_0) = \nabla_\beta s_n(0,0,\beta_\varepsilon)(\beta_n-\beta_0) + s_n(0,0,\beta_0)$ where $\beta_\varepsilon \in [\beta_n,\beta]$.

We make the following decomposition:

$$\begin{aligned}&s_n(0,0,\beta_n) - s(0,0,\beta_0)\\ =&(s_n(0,0,\beta_0) - s(0,0,\beta_0)) + (s(0,0,\beta_n) - s(0,0,\beta_0)) +\\ &[(s_n(0,0,\beta_0) - s(0,0,\beta_0)) - (s_n(0,0,\beta_n) - s(0,0,\beta_n))]\end{aligned}$$

The last term in the bracket is an empirical process term. Given that our score function $s(0,0,\beta)$ is a Donsker class, and β_n is consistent, the empirical process term is $o_P(n^{-1/2})$ (Lemma 19.24 of (Van der Vaart 2000)).

Furthermore, by mean value theorem,

$$s(0,0,\beta_n) = s(0,0,\beta_0) + \nabla_\beta s(0,0,\beta_\varepsilon)(\beta_n-\beta_0)$$

for some $\beta_\varepsilon \in [\beta_n,\beta]$.

Thus,

$$\begin{aligned}s_n(0,0,\beta_n) =& s_n(0,0,\beta_n) - s(0,0,\beta_0)\\ =&\nabla_\beta s_n(0,0,\beta_\varepsilon)(\beta_n-\beta_0) + s_n(0,0,\beta_0)\end{aligned} \tag{22}$$

Furthermore, by central limit theorem (CLT),

$$s_n(0,0,\beta_0) = s(0,0,\beta_0) + O_p\left(n^{-1/2}\right) = O_p\left(n^{-1/2}\right)$$

Notice that $\nabla_\alpha s_n(\alpha_n, \rho_n, \beta_n), \nabla_\rho s_n(\alpha_n, \rho_n, \beta_n)$ are bounded, i.e., they are both $O_P(1)$

$$
\begin{aligned}
&0 = s_n(\alpha_n, \rho_n, \beta_n) = s_n(\alpha_n, \rho_n, \beta_n) = s_n(0, 0, \beta_n) + \\
&\nabla_\alpha s_n(\alpha_\varepsilon, \rho_\varepsilon, \beta_n)\alpha_n + \nabla_\rho s_n(\alpha_\varepsilon, \rho_\varepsilon, \beta_n)\rho_n \\
=&\nabla_\beta s_n(0, 0, \beta_\varepsilon)(\beta_n - \beta_0) + O_p\left(n^{-1/2} + \alpha_n + \rho_n\right)
\end{aligned}
$$

Then, because the minimum eigenvalue of $\nabla_\beta s(0, 0, \beta_\varepsilon)$ and is lower bounded, the above turns to :

$$
\beta_n - \beta_0 = \left(\nabla_\beta s_n(0, 0, \beta_\varepsilon)\right)^{-1} O_p\left(n^{-1/2} + \alpha_n + \rho_n\right) = o_p(1)
$$

Therefore, β_n is consistent.

We now prove the rate result. To do this, we make a second order Taylor's expansion of $s_n(\alpha_n, \rho_n, \beta_n)$ at $\alpha_n = \rho_n = 0$:

$$
\begin{aligned}
&s_n(\alpha_n, \rho_n, \beta_n) \\
=&s_n(0, 0, \beta_n) + \nabla_\alpha s_n(0, 0, \beta_n)\alpha_n + \nabla_\rho s_n(0, 0, \beta_n)\rho_n \\
&+ \frac{1}{2}\nabla_\alpha^2 s_n(\alpha_\varepsilon, \rho_\varepsilon, \beta_n)\alpha_n^2 + \frac{1}{2}\nabla_\rho^2 s_n(\alpha_\varepsilon, \rho_\varepsilon, \beta_n)\rho_n^2 \\
&+ \nabla_{\alpha\rho} s_n(\alpha_\varepsilon, \rho_\varepsilon, \beta_n)\alpha_n\rho_n \\
=&s_n(0, 0, \beta_0) + \nabla_\beta s_n(0, 0, \beta_\varepsilon)(\beta_n - \beta_0) + \nabla_\alpha s_n(0, 0, \beta_n)\alpha_n \\
&+ \nabla_\rho s_n(0, 0, \beta_n)\rho_n + \frac{1}{2}\nabla_\alpha^2 s_n(\alpha_\varepsilon, \rho_\varepsilon, \beta_n)\alpha_n^2 \\
&+ \frac{1}{2}\nabla_\rho^2 s_n(\alpha_\varepsilon, \rho_\varepsilon, \beta_n)\rho_n^2 + \nabla_{\alpha\rho} s_n(\alpha_\varepsilon, \rho_\varepsilon, \beta_n)\alpha_n\rho_n
\end{aligned}
$$

where $\beta_\varepsilon \in [\beta_n, \beta_0], \alpha_\varepsilon \in [0, \alpha_n], \rho_\varepsilon \in [0, \rho_n]$.

The first order Taylor's expansion of $\nabla_\alpha s_n(0, 0, \beta_n)$ at β_0 is:

$$
\begin{aligned}
&\nabla_\alpha s_n(0, 0, \beta_n) \\
&= \nabla_\alpha s_n(0, 0, \beta_0) + \nabla_{\alpha\beta} s_n(0, 0, \beta_\varepsilon)(\beta_n - \beta_0) \\
&= \nabla_\alpha s(0, 0, \beta_0) + O_p\left(n^{-1/2}\right) + \nabla_{\alpha\beta} s_n(0, 0, \beta_\varepsilon)(\beta_n - \beta_0) \\
&= O_p\left(n^{-1/2}\right) + \nabla_{\alpha\beta} s_n(0, 0, \beta_\varepsilon)(\beta_n - \beta_0)
\end{aligned}
$$

where we apply the CLT in the second equation and use $\nabla_\alpha s(0, 0, \beta_0) = 0$ in the last equation due to the Neyman orthogonality of the score function for the partial likelihood (see Appendix in (Gao and Hastie 2021) for a proof). A similar analysis holds for $\nabla_\rho s_n(0, 0, \beta_n)$.

Combining these results:

$$
\begin{aligned}
&s_n(\alpha_n, \rho_n, \beta_n) = O_p\left(n^{-1/2}\right) + \nabla_\beta s_n(0, 0, \beta_\varepsilon)(\beta_n - \beta_0) \\
&+ O_p\left(n^{-1/2}(\alpha_n + \rho_n)\right) + \nabla_{\alpha\beta} s_n(0, 0, \beta_\varepsilon)(\beta_n - \beta_0)\alpha_n \\
&+ \nabla_{\rho\beta} s_n(0, 0, \beta_\varepsilon)(\beta_n - \beta_0)\rho_n + \frac{1}{2}\nabla_\alpha^2 s_n(\alpha_\varepsilon, \rho_\varepsilon, \beta_n)\alpha_n^2 \\
&+ \frac{1}{2}\nabla_\rho^2 s_n(\alpha_\varepsilon, \rho_\varepsilon, \beta_n)\rho_n^2 + \nabla_{\alpha\rho} s_n(\alpha_\varepsilon, \rho_\varepsilon, \beta_n)\alpha_n\rho_n \\
=&\left(\nabla_\beta s_n(0, 0, \beta_\varepsilon) + O_p(\alpha_n + \rho_n)\right)(\beta_n - \beta_0) \\
&+ O_p\left(\alpha_n^2 + \rho_n^2 + \alpha_n\rho_n + n^{-1/2}\right)
\end{aligned}
$$

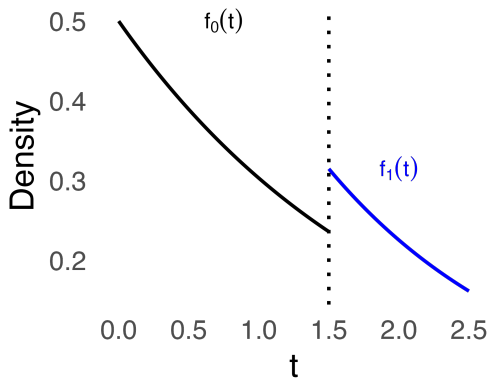

Figure 3: Comparison of treatment and control PDFs

where we use the boundedness of the second derivatives. Since the minimal eigenvalue of $\nabla_\beta s_n(0, 0, \beta_\varepsilon)$ is uniformly lower bounded by $C/2$, we have:

$$
\beta_n - \beta_0 = O_p\left(n^{-1/2} + \alpha_n^2 + \rho_n^2 + \alpha_n\rho_n\right)
$$

This completes the proof.

Parametrizing the PDF for the Piecewise Cox Model

Given two probability density functions (PDFs) $f^{co}(t|a, X_i)$ and $f^{tx}(t|a, X_i)$, there are two ways to parametrize the desired piecewise PDF to show treatment effect. Both parametrizations integrate to one and are illustrated in Figure 3.

Parametrization 1:

$$
f(t|a, X_i) = \begin{cases} f^{co}(t|a, X_i) & \text{for } t < a \\ f^{tx}(t|a, X_i) \cdot \frac{1 - F^{co}(a|a, X_i)}{1 - F^{tx}(a|a, X_i)} & \text{for } t \geq a \end{cases} \tag{23}
$$

Parametrization 2:

$$
f(t|a, X_i) = \begin{cases} f^{co}(t|a, X_i) & \text{for } t < a \\ f^{tx}(t - a|a, X_i) \cdot [1 - F^{co}(a|a, X_i)] & \text{for } t \geq a \end{cases} \tag{24}
$$

Question: Are these parametrizations equivalent? If not, which one is preferable?

Answer: No, they are not equivalent in general. Let's convert Parametrizations 1 and 2 into hazard functions. We only need to compare the expressions for $t \geq a$. For clarity of notation, we omit the conditioning on a, X_i in the derivations below.

Parametrization 1: For $t \geq a$:

$$
h(t) = \frac{f^{tx}(t) \cdot \frac{1 - F^{co}(a)}{1 - F^{tx}(a)}}{1 - \int_0^t f(s)ds}
$$

The denominator:

$$\begin{aligned}
1 - \int_0^t f(s)ds &= 1 - \left(\int_0^a f^{co}(s)ds + \int_a^t f^{tx}(s)ds \cdot \frac{1 - F^{co}(a)}{1 - F^{tx}(a)} \right) \\
&= 1 - \left(F^{co}(a) + (F^{tx}(t) - F^{tx}(a)) \cdot \frac{1 - F^{co}(a)}{1 - F^{tx}(a)} \right) \\
&= (1 - F^{co}(a)) - (F^{tx}(t) - F^{tx}(a)) \cdot \frac{1 - F^{co}(a)}{1 - F^{tx}(a)} \\
&= (1 - F^{co}(a)) \left(1 - (F^{tx}(t) - F^{tx}(a)) \cdot \frac{1}{1 - F^{tx}(a)} \right) \\
&= (1 - F^{co}(a)) \frac{1 - F^{tx}(t)}{1 - F^{tx}(a)} \\
&= (1 - F^{tx}(t)) \frac{1 - F^{co}(a)}{1 - F^{tx}(a)}
\end{aligned}$$

Therefore:

$$\begin{aligned}
h(t) &= \frac{f^{tx}(t) \cdot \frac{1 - F^{co}(a)}{1 - F^{tx}(a)}}{(1 - F^{tx}(t)) \frac{1 - F^{co}(a)}{1 - F^{tx}(a)}} \\
&= \frac{f^{tx}(t)}{1 - F^{tx}(t)} \\
&= h^{tx}(t)
\end{aligned}$$

Parametrization 2: For $t \geq a$:

$$h(t) = \frac{f^{tx}(t - a) \cdot (1 - F^{co}(a))}{1 - \int_0^t f(s)ds}$$

The denominator:

$$\begin{aligned}
&1 - \int_0^t f(s)ds \\
&= 1 - \left(\int_0^a f^{co}(s)ds + \int_a^t f^{tx}(s - a)ds \cdot (1 - F^{co}(a)) \right) \\
&= 1 - F^{co}(a) - \int_0^{t-a} f^{tx}(s)ds \cdot (1 - F^{co}(a)) \\
&= (1 - F^{co}(a)) \cdot (1 - F^{tx}(t - a))
\end{aligned}$$

Therefore:

$$\begin{aligned}
h(t) &= \frac{f^{tx}(t - a) \cdot (1 - F^{co}(a))}{(1 - F^{co}(a)) \cdot (1 - F^{tx}(t - a))} \\
&= \frac{f^{tx}(t - a)}{1 - F^{tx}(t - a)} \\
&= h^{tx}(t - a)
\end{aligned}$$

Discussion Parametrization 1 yields $h(t) = h^{tx}(t)$, while Parametrization 2 yields $h(t) = h^{tx}(t-a)$. These are equivalent only when h^{tx} is constant. For example, if $h^{tx}(s) = s$, the parametrizations differ. Therefore, the equivalence PDF parametrization from the paper's hazard model corresponds to Parametrization 1.