A New Causal Approach to Account for Treatment Switching in Randomized Experiments under a Structural Cumulative Survival Model

Andrew Ying and Eric J. Tchetgen Tchetgen

Department of Statistics, the Wharton School, University of Pennsylvania

Abstract

Background: Treatment switching in a randomized controlled trial is said to occur when a patient randomized to one treatment arm switches to another treatment arm during follow-up. This can occur at the point of disease progression, whereby patients in the control arm may be offered the experimental treatment. It is widely known that failure to account for treatment switching can seriously dilute the estimated effect of treatment on overall survival. In this paper, we aim to account for the potential impact of treatment switching in a re-analysis evaluating the treatment effect of Nucleoside Reverse Transcriptase Inhibitors (NRTIs) on a safety outcome (time to first severe or worse sign or symptom) in participants receiving a new antiretroviral regimen that either included or omitted NRTIs in the Optimized Treatment That Includes or Omits NRTIs (OPTIONS) trial.

Methods: We propose an estimator of a treatment causal effect under a structural cumulative survival model (SCSM) that leverages randomization as an instrumental variable to account for selective treatment switching. Unlike Robins’ accelerated failure time model often used to address treatment switching, the proposed approach avoids the need for artificial censoring for estimation. We establish that the proposed estimator is uniformly consistent and asymptotically Gaussian under standard regularity conditions. A consistent variance estimator is also given and a simple resampling approach provides uniform confidence bands for the causal difference comparing treatment groups over time on the cumulative intensity scale. We develop an R package named “ivsacim” implementing all proposed methods, freely available to download from R CRAN. We examine the finite performance of estimator via extensive simulations.

Results: 357 participants in the OPTIONS trial were randomly assigned at baseline to add-NRTIs or omit-NRTIs treatment group; 93% subsequently completed a

48-week visit. Using the proposed methods, we found statistically significant evidence against the sharp null hypothesis of no treatment effect on the safety outcome (P value 0.034) and our SCSM estimator revealed an increased risk for a safety outcome in participants receiving a new antiretroviral regimen that included NRTIs when compared to participants receiving a regimen that omitted NRTIs. In fact, under an SCSM encoding a constant additive hazards model, we estimated a hazards difference equal to 0.0039 (95% CI 0.0002, 0.0075) over the 48-week follow-up.

Conclusions: Treatment-experienced patients with HIV infection starting a new optimized regimen will experience a higher risk of severe or worse sign or symptom. Previous analyses concluded that treatment-experienced patients with HIV infection starting a new optimized regimen can safely omit NRTIs without compromising virologic efficacy. Our analysis suggests that adding NRTIs is not only unnecessary to achieve optimal outcomes but may increase the risk for a safety outcome.

Keywords: treatment switching, treatment crossover, G-estimation, instrumental variable.

1 Background

1.1 Treatment Switching

Treatment switching (also called contamination or crossover) in a randomized controlled trial (RCT) is said to occur when a patient randomized to one treatment arm changes to another arm during the course of follow-up. This can occur at the point of disease progression, where patients in the control arm are switched to the experimental treatment in the hope of improving their prognostic. This can happen for a variety of reasons: a clinician may find switching treatment arms might be the best option for his/her patient; it may sometimes be pre-specified in the trial protocol as part of a dynamic treatment strategy; or driven by emerging evidence (internally or externally to the specific RCT) of the active treatment's benefit may have broken the original trial equipoise. For example, Demetri et al. (2006) reported the results of a randomized controlled trial into the use of Sunitinib for the treatment of advanced gastrointestinal stromal tumours in patients for whom conventional

therapy (Imatinib) had failed because of resistance or intolerance. Early trial results showed a strong benefit in favour of the new treatment in terms of time to tumour progression. It then led to change to an open label protocol and made Sunitinib available to the patients in the placebo arm, which led to the vast majority of eligible patients in the placebo arm switching to receive Sunitinib. This makes interpretation of data used in subsequent analyses difficult. More examples can be found in Cuzick, Edwards, and Segnan (1997), Motzer et al. (2008), and Morden et al. (2011).

The estimated treatment effect on overall survival can be substantially diluted after patients switch without appropriate handling (Bowden et al. 2016). Treatment switching can also bias the intent-to-treat effect by contaminating the treatment arm. Simple methods include censoring patients at the time-point where switching occurs, excluding switching patients from the analysis altogether, or modeling treatment as a time-varying covariate. Although commonly used, these methods are prone to selection bias because patients who switch treatment tend to have a different overall survival prognosis than patients who remain on their originally assigned treatment.

More rigorous approaches designed to account for this selection bias are available, including structural nested model (SNM) (Robins and Greenland 1994; Yamaguchi and Ohashi 2004), inverse probability of censoring weights (IPCW) (Robins and Finkelstein 2000), and 2-stage adjustment methods (Latimer, K. Abrams, et al. 2017) among those currently considered most promising. Simulation studies have shown that these methods tend to produce more accurate estimates of the switching-adjusted estimand than simple adjustment methods or a standard intention to treat (ITT) analysis. See Jimenez et al. (2017), Latimer, K. R. Abrams, et al. (2018), and Sullivan et al. (2020) for a systematic review of their applications. However, their performance can be compromised when the key underlying assumption of “no unmeasured confounding” is violated.

1.2 Instrumental Variable

To accommodate unmeasured confounding, a prevailing approach resorts to "so-called" instrumental variables (IV). Instrumental variables have a long tradition in econometrics (Angrist and Krueger 2001; Wooldridge 2010). They have received increasing attention in epidemiology due to a revival of Mendelian randomization studies (Davey Smith and Ebrahim 2003; Katan 2004).

With time-to-event data, Tchetgen Tchetgen, Walter, et al. (2015) demonstrated the validity of two-stage estimation approaches in additive hazard models for event times when the exposure obeys a particular location shift model. Two-stage estimation approaches (Li, Fine, and Brookhart 2015; Ying, Xu, and Murphy 2019) usually require a first-stage exposure distribution model that can be unnecessarily restrictive. Recently, Martinussen et al. (2017) leveraged IVs using G-estimation to examine time-varying effects of a point exposure on the survival function, in the presence of unmeasured confounding. By instead modeling the structural cumulative survival function and the conditional mean of the instrumental variable given covariates, they managed to handle arbitrary exposures and IVs. However, all aforementioned IV methods in survival analysis that accommodate unmeasured confounding can only handle point exposure and are not readily available for treatment switching, which formally entails a time-varying treatment.

The rank-preserving structural failure time model (RPSFTM) of Robins and Tsiatis (1991) also offers a solution to treatment switching leveraging randomization as an IV. The RPSFTM is usually estimated by inverting a logrank test statistic. Unfortunately, logrank-test-based estimation can be inefficient (Schoenfeld 1981; Robins and Greenland 1994). Bowden et al. (2016) proposed a weighted logrank test in an attempt to recover some information. In either approach, a major complication with estimation of RPSFTM is the need for artificial censoring to address administrative censoring, which further aggravates

efficiency loss and renders estimation computationally challenging (Robins and Tsiatis 1991; M. M. Joffe 2001; M. M. Joffe, Yang, and Feldman 2012; Vansteelandt, M. Joffe, et al. 2014).

In this paper, we accommodate treatment switching under a structural cumulative survival model (SCSM) and demonstrate estimation and inference leveraging randomization as an IV, without requiring artificial censoring. We further develop an asymptotic framework for inference based on a proposed recursive estimator of the SCSM which effectively extends the G-estimator under a structural cumulative failure time model for point exposure to the time-varying exposure setting (Martinussen et al. 2017). We also propose inferential tools to investigate exposure effects that vary over time. All proposed estimators and tools are implemented in the freely available R package “ivsacim” (Ying 2021).

We illustrate the proposed approach in an analysis that aims to evaluate the treatment effect of Nucleoside Reverse Transcriptase Inhibitors (NRTIs) on a key safety outcome (time to first severe or worse sign or symptom) in participants receiving a new antiretroviral regimen that omitted or added NRTIs in the Optimized Treatment That Includes or Omits NRTIs (OPTIONS) trial (Tashima, Mollan, et al. 2015; Tashima, L. M. Smeaton, et al. 2015), where treatment switching is present due to possible discontinuation of the NRTIs assignment.

1.3 The OPTIONS Trial

The OPTIONS trial was a multi-center, open-label, prospective, randomized, controlled study evaluating the benefits and risks of omitting versus adding NRTIs to a new optimized antiretroviral regimen (Tashima, Mollan, et al. 2015). The study population consists of HIV-infected patients for whom a PI-based regimen has failed and who have triple-class experience (NNRTIs, NRTIs, and PIs) and viral resistance. Study participants were recruited from 62 outpatient medical clinics into the trial centers across the United States from March

2008 through May 2011, with follow-up through 48 weeks (31 May 2012). The study population included HIV-1–infected persons who were at least 16 years of age; had a plasma HIV RNA level of 1000 copies/mL or more; had received a PI-based antiretroviral regimen; had previously used or had evidence of resistance to NRTIs and NNRTIs; and had acceptable laboratory values, including a calculated 2 creatinine clearance of 50 mL/min/1.73 m. Persons were ineligible if they had active hepatitis B infection, were pregnant or breast-feeding, or were using prohibited medications. A key criterion for randomization was that an individualized regimen with a cPSS greater than 2.0 could be constructed using study antiretroviral medications, excluding NRTIs. A cPSS (0 [not susceptible] to 1 [susceptible]) was calculated or assigned for each drug in a potential regimen based on the participant’s prior drug exposure, virus susceptibility, and tropism result (Tashima, L. M. Smeaton, et al. 2015; Tashima, Mollan, et al. 2015).

Participants were randomly assigned either to omit or to add NRTIs after choosing an optimized regimen and an NRTI regimen. Before randomization, a cPSS was calculated for each participant for 20 different optimized regimens. One or more optimized regimens with a cPSS greater than 2.0 and NRTI regimens were recommended by the study team and sent to sites for selection before randomization.

Treatment switching occurred in this trial due to potential discontinuation of NRTI assignment. This occurred when a participant in the omit-NRTIs group started any NRTI or when a participant in the add-NRTIs group failed to initiate or permanently discontinued all NRTIs (event time was the scheduled week during which the event was recorded).

Previously Tashima, L. M. Smeaton, et al. (2015) examined treatment efficacy by comparing time to virologic failure between groups. However, they handled treatment switching by re-defining their primary outcome as a composite outcome reflecting regimen failure, defined as virologic failure or change in NRTI group assignment (whichever comes first), evaluated through 48 weeks. This ad hoc approach was not necessarily unreasonable for

the primary outcome as only 10 of 357 subjects switched their treatment prior to virologic failure or censoring. However, in assessing treatment safety in this trial, the primary safety outcome was defined as time to first severe or worse sign or symptom; for which the higher number 17 of 357 subjects switched treatment during follow-up. In fact, they conducted simple safety analyses by performing a stratified log-rank tests for randomized treatment. They concluded that “time to first severe or worse sign or symptom did not significantly differ between groups ($P = 0.149$)”. However, their safety analysis did not formally account for treatment switching and therefore may be biased towards the null due to contamination of the treatment arms, thus failing to reveal safety concerns. We later show that correctly accounting for treatment switching using our proposed methods reveals a concerning safety signal as evidenced by a statistically significant causal effect of treatment actually taken on safety outcomes, with exposure to an add-NRTIs treatment experiencing a higher rate of the safety outcome than the exposure to an omit-NRTIs treatment.

2 Methods

2.1 Notation

Data for subjects $i = 1, \dots, n$ are treated as independent and identically distributed. Define

- $\tilde{T}_i$, safety outcome corresponding to time to first severe or worse sign or symptom;
- C_i , potential censoring time;
- $T_i = \min(\tilde{T}_i, C_i)$ denotes a subject’s censored event time;
- $\Delta_i = \mathbb{1}(\tilde{T}_i \leq C_i)$ denotes a subject’s observed event indicator;

- We introduce the counting process notation. We write $N_i(t) = \mathbb{1}(T_i \leq t, \Delta_i = 1)$ as the observed counting process, $\tilde{N}_i(t) = \mathbb{1}(\tilde{T}_i \leq t)$ the counting process of time to first severe or worse sign or symptom. We also define $Y_i(t) = \mathbb{1}(T_i \geq t)$ and $\tilde{Y}_i(t) = \mathbb{1}(\tilde{T}_i \geq t)$ the associated at-risk processes;
- We assume that recorded data on treatment do not change except at discrete times $\{t_1, \dots, t_M\}$. Thus the time varying treatment $D_i(t_m) = 1$ if subject i is treated or exposed at time t_m , 0 otherwise, write $\bar{D}_i(t_m) = \{D_i(t_l) : 0 < l \leq m\}$. For any $t_m > T_i$, $D_i(t_m)$ is not observed. We define $D_i(t_m) = 0$ for $t_m > T_i$, so that the whole treatment process is well defined for each subject even after the outcome event has occurred;
- Z_i denotes the instrumental variable corresponding to baseline randomization in the OPTIONS trial;
- L_i denotes baseline (pre-randomization) covariates which we allow for.

We may drop the subscript i when there is no confusion for simplicity. We also introduce the counterfactual outcomes framework for time-varying treatments (Robins 1986; Robins 1987),

- $\tilde{T}(\bar{d}(t_m), 0)$, the potential time to event had possibly contrary to fact, the subject followed the treatment regime $\bar{d}(t_m)$ up to time t_m and the control treatment thereafter. We make the consistency assumption that $T = \tilde{T}(\bar{d}(t_m), 0)$ with probability one for individuals with observed $\bar{D}(t_m) = d(t_m)$ and $D(t_l) = 0$, for $l > m$. We further assume that intervening on an exposure can only affect survival after the time of that exposure, in other words, the event $\tilde{T}(\bar{D}(t_{m-1}), 0) \geq t_m$ occurs if and only if the event $\tilde{T}(\bar{D}(t_l), 0) \geq t_m$ also occurs for all $l \geq m$. It follows that $\{\tilde{T} \geq t\}$ and $\{\tilde{T}(\bar{D}(t_m), 0) \geq t\}$ are the same events for $t \in [t_m, t_{m+1})$;

- We write $\tilde{N}_{\bar{d}(t_m),0}(t)$ and $\tilde{Y}_{\bar{d}(t_m),0}(t)$ as the associated potential counting process and potential at risk process;

2.2 Model

We assume a structural cumulative survival model,

$$\begin{aligned} & \frac{\mathbb{P}\left[\tilde{T}(\bar{D}(t_m), 0) > t | \bar{D}(t_m), Z, L, \tilde{T} \geq t_m\right]}{\mathbb{P}\left[\tilde{T}(\bar{D}(t_{m-1}), 0) > t | \bar{D}(t_m), Z, L, \tilde{T} \geq t_m\right]} \\ &= \exp\left\{-\int_{t_m}^{t \wedge t_{m+1}} D(t_m) dB_D(s)\right\}, \end{aligned} \quad (1)$$

for any $t \geq t_m$.

This model may be interpreted as encoding for individuals still at risk for the outcome at time t_m with covariates, IV and treatment history $L, Z, \bar{D}(t_m)$, the ratio of survival probabilities of remaining event free at time $t \geq t_m$ upon receiving one final blip of treatment at time $t_m - 1$ versus at time t_m . This ratio of conditional survival probabilities is modeled by the RHS of (1). As NRTIs can be safely omitted without compromising efficacy (Tashima, L. M. Smeaton, et al. 2015), $dB_D(t)$ here can be interpreted as a measure of safety risk for those who share the same history and are treated at time t_m and not thereafter, compared to individuals who are not treated after time t_{m-1} . A positive value of $dB_D(t)$ implies that NTRIs increase patients' safety risk. It provides direct evidence as to whether patients should switch to the control arm for safety reasons.

Remark 1. Seaman et al. (2020) investigated a general class of structural cumulative survival models (SCSMs), which in our setting and notation may be defined as

$$\begin{aligned} & \frac{\mathbb{P}\left[\tilde{T}(\bar{D}(t_m), 0) > t | \bar{D}(t_m), Z, L, \tilde{T} \geq t_m\right]}{\mathbb{P}\left[\tilde{T}(\bar{D}(t_{m-1}), 0) > t | \bar{D}(t_m), Z, L, \tilde{T} \geq t_m\right]} \\ &= \exp\left\{-D(t_m)\gamma_m(t; \bar{D}(t_{m-1}), L, Z)\right\}. \end{aligned}$$

This SCSM accommodates time varying effects as a function of both the survival time t under consideration and the time of a final treatment blip $t_m \leq t$; as well as effect heterogeneity as a function of a patient's observed history $(\bar{D}(t_{m-1}), L, Z)$. Model (1) is a special case of the model in the above display which posits (i) that the causal effect of a final blip of treatment at time t_m is short-lived in the sense that $\gamma_m(t; \bar{D}(t_{m-1}), L, Z) = \gamma_m(t_{m+1}; \bar{D}(t_{m-1}), L, Z)$ for all $t \geq t_{m+1}$; and (ii) does not depend on a patient's history, i.e. $\gamma_m(t; \bar{D}(t_{m-1}), L, Z)$ is a constant function of $(\bar{D}(t_{m-1}), L, Z)$. Together (i) and (ii) imply:

$$\gamma_m(t; L) = B_D(t \wedge t_{m+1}) - B_D(t_m).$$

Our simpler model specification delivers a convenient summary of the cumulative treatment effect comparing the marginal survival function under the always-treated treatment regime (i.e. the survival curve for $\tilde{T}(\bar{1})$) versus that of the never-treated treatment regime (i.e. the survival curve for $\tilde{T}(\bar{0})$) in terms of the function $B_D(t)$, our primary causal effect of interest. That is, under model (1), and the additional condition of no-current treatment value interaction (Robins and Greenland 1994),

$$\begin{aligned} & \frac{\mathbb{P} \left[\tilde{T}(\bar{D}(t_{m-1}), d(t_m), 0) > t \mid \bar{D}(t_{m-1}), d(t_m), Z, L, \tilde{T} \geq t_m \right]}{\mathbb{P} \left[\tilde{T}(\bar{D}(t_{m-1}), 0) > t \mid \bar{D}(t_{m-1}), d(t_m), Z, L, \tilde{T} \geq t_m \right]} \\ &= \frac{\mathbb{P} \left[\tilde{T}(\bar{D}(t_{m-1}), d(t_m), 0) > t \mid \bar{D}(t_{m-1}), D(t_m) = 0, Z, L, \tilde{T} \geq t_m \right]}{\mathbb{P} \left[\tilde{T}(\bar{D}(t_{m-1}), 0) > t \mid \bar{D}(t_{m-1}), D(t_m) = 0, Z, L, \tilde{T} \geq t_m \right]}, \end{aligned}$$

one can readily establish that

$$\frac{\mathbb{P} \left[\tilde{T}(\bar{1}) > t \right]}{\mathbb{P} \left[\tilde{T}(\bar{0}) > t \right]} = \exp \left\{ - \int_0^t 1 dB_D(s) \right\} = \exp \{ -B_D(t) \}.$$

Therefore our estimand $B_D(t)$ can be interpreted as the difference in the log-marginal cumulative intensity function comparing always-treated versus never-treated regimes up to

time t . The no-current treatment value interaction assumption essentially states that the instantaneous causal effect of one final blip of treatment at time t_m among individuals who were treated at time t_m is equal to that among individuals who were not treated at time t_m conditional on past history.

In comparison, Seaman et al. (2020) were interested in the more ambitious goal of learning about treatment effects under all possible treatment regimes under a sequential ignorability condition, i.e., that the time-varying treatment mechanism is not subject to unmeasured confounding. This more ambitious goal may not be identified in the OPTIONS study for two reasons: first, unlike Seaman et al. (2020), we do not make an assumption that treatment switching is unconfounded, thus allowing for both time-fixed and time-varying unmeasured confounding of the treatment-outcome relationship, thus making identification of complex treatment effects far more challenging; second, as patients who switched treatment arms did so at most once in the OPTIONS study, investigating the causal effect of more complex treatment regimes may not be supported by the observed sample. It is worth noting that other possible SCSM variants may be specified, including one which postulates a constant effect for each dose of treatment over time, while allowing for the magnitude of the treatment effect to depend on the timing of the final blip of treatment:

$$\gamma_m(t; L) = \beta_{D,m}(t - t_m).$$

A special case of this model may further impose that the treatment effects are constant as a function of the timing of the final treatment blip:

$$\gamma_m(t; L) = \beta_D(t - t_m). \tag{2}$$

2.3 Assumptions

Our proposed identification strategy leverages the randomization process as an instrumental variable satisfying three key standard IV assumptions:

Assumption 1 (IV relevance). *The instrument is associated with the exposure at t_m for individuals still at risk for the event time for all t_m ; specifically,*

$$Z \not\perp D(t_m) \mid \tilde{T} \geq t_m, \bar{D}(t_{m-1}), L,$$

for all t_m .

IV relevance requires that for subjects who remain at risk for the outcome event at time t_m , the instrument remains predictive of current treatment status even after conditioning on treatment and covariate history. This is a reasonable assumption in OPTIONS, given that individuals randomized to the active arm are more likely than the control arm to be treated over time, even upon conditioning on their history.

Assumption 2 (IV independence). *The instrument is independent of the potential outcome under no treatment, conditional on baseline covariates ,*

$$Z \perp \tilde{T}(0) \mid L.$$

IV independence ensures that the IV itself is unconfounded (conditional on L). This assumption is clearly satisfied in the OPTIONS trial as Z is randomized treatment assignment.

Assumption 3 (Exclusion restriction). *The instrument has no direct causal effect on the outcome other than through exposure, namely,*

$$\tilde{T}(\bar{d}, z) = \tilde{T}(\bar{d}, z'),$$

for all values of $\bar{d}$, z and z' , where $\tilde{T}(\bar{d}, z)$ denotes the potential outcome had one intervened to set Z and $\bar{D} = \{D_{t_1}, \dots, D_{t_M}\}$ to z and $\bar{d}$, respectively.

The exclusion restriction rules out the possibility that randomization itself can impact safety via a pathway not involving treatment actually taken, a reasonable assumption in the

OPTIONS trial to the extent that participants did not modify their behavior as a result of randomization in a manner that might influence their safety through a pathway independent of the treatment regime ultimately followed throughout the trial. There is no a priori reason to suspect that randomized assignment to the intervention adding NRTIs to a treatment regime would result in violation of this underlying assumption.

In order to facilitate the exposition, we make a standard conditional independent censoring assumption.

Assumption 4 (Conditional independent censoring).

$$C \perp (\tilde{T}, \bar{D}(t), Z) \mid L.$$

Note that although not further pursued here, the above assumption can be relaxed substantially by only requiring that $C \perp \tilde{T} \mid \bar{D}(t_m), Z, L, T \geq t_m$, which however would require further adjustment for dependent censoring that can be achieved by standard inverse probability censoring weighting (Robins and Finkelstein 2000; Robins 2000).

2.4 Estimation

The crux of our approach for estimating $B_D(t)$, $t > 0$ is that once the exposure effect has been eliminated from the event time under the assumed SCSM, the “residualized outcome” would mimic $T(0)$ distributionally, and therefore should in principle satisfy Assumption 2. Building on this intuition, we formally establish that the parameter of interest $B_D(t)$ is in fact a solution to the following unbiased estimating equation:

$$\mathbb{E} \left\{ \underbrace{Z^c \exp \left(\int_0^{t-} D(s) dB_D(s) \right)}_{\text{Residualizes the at risk process}} \left[dN(t) - \underbrace{Y(t) D(t) dB_D(t)}_{\text{Residualizes the counting process}} \right] \right\} = 0, \quad (3)$$

where $Z^c = Z - \mathbb{E}(Z|L)$.

Proposition 1. *Under Assumptions 1, 2, 3, 4 and the model (1), the parameter of interest $B_D(t)$ in (1) is the unique solution to the population estimating equation (3), which admits a closed form solution*

$$B_D(t) = \int_0^t \frac{\mathbb{E} \left\{ Z^c \exp \left[\int_0^{s-} D(u) dB_D(u) \right] dN(s) \right\}}{\mathbb{E} \left\{ Z^c Y(s) \exp \left[\int_0^{s-} D(u) dB_D(u) \right] D(s) \right\}}, \quad (4)$$

provided that $\mathbb{E}\{Z^c Y(t) \exp[\int_0^{t-} D(s) dB_D(s)] D(t)\} \neq 0$ for all $t > 0$.

We note that IV relevance is necessary but not sufficient for $\mathbb{E}\{Z^c Y(t) \exp[\int_0^{t-} D(s) dB_D(s)] D(t)\} \neq 0$ for all $t > 0$. Nevertheless, in settings such as the OPTIONS study in which Z and $D(t)$ are positively associated conditional on $\tilde{T} \geq t_m, \bar{D}(t_{m-1}), L$, for all $t > t_m$ and m , its covariance with $Y(t) D(t) \exp \left[\int_0^{s-} D(u) dB_D(u) \right]$ can be expected to be strictly positive under IV relevance. The closed form (4) of $B_D(t)$ suggests a two-step estimator. Suppose that one can correctly specify a model $\mathbb{E}(Z|L; \theta)$ for $E(Z|L)$ indexed by the finite dimensional parameter θ . Note that in the OPTIONS trial the randomization probability $E(Z|L; \theta) = Pr(Z = 1)$ is known by design. Otherwise, assume that one can readily obtain a consistent estimator $\hat{\theta}$ for θ that is asymptotic linear with influence function $\epsilon_{2,i}(\theta)$, namely, $n^{1/2}(\hat{\theta} - \theta) = n^{-1/2} \sum_{i=1}^n \epsilon_{2,i}(\theta) + o_p(1)$. An empirical analog of (3) is thus:

$$\sum_{i=1}^n Z_i^c(\hat{\theta}) \exp \left(\int_0^{t-} D_i(s) d\hat{B}_D(s) \right) \left[dN_i(t) - Y_i(t) D_i(t) d\hat{B}_D(t) \right] = 0.$$

The corresponding estimator $\hat{B}_D(t, \hat{\theta})$ has explicit recursive form:

$$\sum_{i: T_i \leq t} \frac{\Delta_i Z_i^c(\hat{\theta}) \exp \left[\sum_{l: T_l < T_i, \Delta_l = 1} D_i(T_l) d\hat{B}_D(T_l, \hat{\theta}) \right]}{\sum_{j: T_j \geq T_i} Z_j^c(\hat{\theta}) \exp \left[\sum_{l: T_l < T_i, \Delta_l = 1} D_i(T_l) d\hat{B}_D(T_l, \hat{\theta}) \right] D_j(T_i)}.$$

Because of its recursive structure, and the key fact that the estimator only changes values at observed event times, we may evaluate it forward in time, with initial value $\hat{B}_D(0, \hat{\theta}) = 0$.

Under model (1), assuming the treatment effect function is a sufficiently smooth function of time, one may write, $B_D(t) = \int_0^t \beta_D(s)ds$, so that in the special case of a constant instantaneous treatment effect (2), $\beta_D(t) = \beta_D$ one may use the estimator

$$\hat{\beta}_D = \int_0^\tau w(t)d\hat{B}_D(t),$$

with $w(t) = \tilde{w}(t) / \int_0^\tau \tilde{w}(s)ds$, $\tilde{w}(t) = \sum_{i=1}^n R_i(t)$ and τ denoting time of end of study. In the Appendix, we establish that under sufficient regularity conditions, the estimator $\hat{B}_D(t, \hat{\theta})$ is uniformly consistent for $B_D(t)$ and $\{\sqrt{n}(\hat{B}_D(t, \hat{\theta}) - B_D(t)) : t\}$ is asymptotically equivalent to a certain Gaussian process. An explicit expression of the covariance function of the limiting Gaussian process is derived in the Appendix.

Theorem 1. *Under Assumptions 1, 2, 3, 4, the model (1), and given the technical Assumptions 5, 6, 7 listed in the Appendix, the estimator $\hat{B}_D(t, \hat{\theta})$ is uniformly consistent for $B_D(t)$ on $[0, \tau]$, namely,*

$$\sup_{t \in [0, \tau]} \left| \hat{B}_D(t, \hat{\theta}) - B_D(t) \right| \rightarrow 0 \quad a.s..$$

Theorem 2. *Under Assumptions 1, 2, 3, 4, the model (1), and given the technical Assumptions 5, 6, 7 listed in the Appendix, the normalized process $\sqrt{n}[\hat{B}_D(t, \hat{\theta}) - B_D(t)]$ converges weakly to a zero mean Gaussian process with variance function that can be consistently estimated by*

$$\frac{1}{n} \sum_{i=1}^n \hat{\epsilon}_i(t, \hat{\theta})^2,$$

with $\hat{\epsilon}_i(t, \hat{\theta})$ defined in the Appendix equation (53).

We have developed an R package named “ivscaim” (Ying 2021) available on R CRAN. In addition to implementing inferences based on the estimator $\hat{B}_D(t, \hat{\theta})$, our package also provides a goodness-of-fit test for the constant causal hazards difference model (2) ($H_0 : B_D(t) = \beta_D t$ for all t) and as well as for the causal null hypothesis ($H_0 : B_D(t) \equiv 0$ for all t), thus

effectively extending to the time-varying treatment similar test statistics for point treatment SCSMs developed by Martinussen et al. (2017).

2.5 Simulation

In order to investigate the finite sample performance of our proposed methods, we conducted a simulation study in which we generated $B = 1000$ data sets of i.i.d data with sample size $N = 800, 1600$. For simplicity, similar to the OPTIONS trial, we omit observed baseline covariates L , and generate a bivariate baseline variable U_i which confounds the relationship between time-varying treatment and the time to event outcome:

$$U_i = (U_{1,i}, U_{2,i})^\top \sim \mathcal{N} \left((1.5, 1.5)^\top, \begin{pmatrix} 1/4 & -1/6 \\ -1/6 & 1/4 \end{pmatrix} \right).$$

We simulate a scenario in which initial treatment assignment $D_i(0)$ is generated as an independent Bernoulli random variable with event probability $\mathbb{P}(D_i(0) = 1) = 0.5$, corresponding to randomized treatment assignment so that $Z_i = D_i(0)$. We also generate a potential treatment switching time W_i for each individual according to

$$\mathbb{P}(W_i > t | Z_i, U_i) = \exp(-0.05 \cdot t - 0.1 \cdot U_{1,i} \cdot t - 0.1 \cdot Z_i), \quad (5)$$

and discretize it into a grid with step size = 0.1. Subject i experiences treatment Z_i before W_i and is switched to $1 - Z_i$ right after W_i . Thus, Z_i and W_i determine a patient's entire treatment process. By (5) we also allow both directions of treatment switching, which fits our application setting. The potential time to event $\tilde{T}_i(\bar{d})$ is then generated according to

$$\mathbb{P}(\tilde{T}_i(\bar{d}) > t | U_i, Z_i) = \exp \left(-0.25 \cdot t - 0.1 \cdot \int_0^t d(s) ds - 0 \cdot Z_i \cdot t - 0.15 \cdot U_{2,i} \cdot t \right),$$

and the observed time to event T_i is thus generated via consistency. In the appendix, we confirm that under the proposed data generating mechanism, we have that

$$\frac{\mathbb{P}(\tilde{T}_i(\bar{D}_i(t_m), 0) > t | \bar{D}_i(t_m), Z_i, \tilde{T}_i > t_m)}{\mathbb{P}(\tilde{T}_i(\bar{D}(t_{m-1}), 0) > t | \bar{D}_i(t_m), Z_i, \tilde{T}_i > t_m)} = \exp \left(-0.1 \cdot \int_{t_m}^{t \wedge t_{m+1}} D(s) ds \right).$$

Independent censoring was then generated with overall rate of 22%. Treatment switching occurred at an approximate rate of 14%, which is chosen to fit the application setting.

Simulation results concerning $B_D(t)$ are given in Table 1, where (average) bias is reported at time points $t = 1, 2, 3$ along with corresponding coverage of 95% confidence intervals $\text{CP}(\hat{B}_D(t))$.

Simulation results confirm that the proposed estimator $\hat{B}_D(t)$ has small bias both at sample size 800 and 1600 at $t = 1, 2, 3$. When $N = 800$, estimated standard errors appear to be overstated resulting in larger coverage than the nominal level. This upward bias is especially large at $t = 3$. However, estimated standard errors match Monte Carlo standard errors as sample size goes up to 1600, and overall 95% confidence intervals attain the nominal level.

3 Results from IV analysis of OPTIONS trial

In the OPTIONS trial, there are 180 patients in the add-NRTIs treatment arm and 177 patients in the omit-NRTIs control arm. Among those, a total of 10 patients switched from the add-NRTIs treatment arm into the omit-NRTIs arm, while 17 patients in the omit-NRTIs group switched into the add-NRTIs arm, corresponding to 5.3% and 9.5% of patients switching treatment arm in the add- and omit-NRTIs groups, respectively. Of these, 9 out of 10 and 8 out of 17 treatment switching occurred before severe or worse sign or symptom in the add- and omit-NRTIs groups, respectively (Figure 2).

A total of 51 severe/worse sign/symptom in the add-NRTIs group and 35 in the omit-

NRTIs group occurred by 48 weeks, corresponding to 28.3% and 19.8% in the add- and omit-NRTIs groups, respectively (Table 2). This amounted to an intent-to-treat estimated difference (i.e. a difference in outcome rates between the randomized arms), of 8.6 percentage points [95% CI, -0.8 to 17.4 percentage points]. Tashima, L. M. Smeaton, et al. (2015) also conducted simple safety analyses by performing a stratified log-rank tests for randomized treatment and concluded that “time to first severe or worse sign or symptom did not significantly differ between groups ($P = 0.149$). These served as basis for Tashima, L. M. Smeaton, et al. (2015) concluding that there was no significant statistical evidence of a safety difference between the randomized groups. However, the above analysis fails to account for treatment switching during follow-up (Figure 1), and is therefore subject to bias towards the null hypothesis.

Using the proposed approach to formally account for treatment switching by leveraging randomized treatment assignment as an IV for treatment actually received which is likely confounded by unmeasured factors, we performed a test of the sharp null hypothesis of no individual causal effect, i.e. $B_D(t) = 0$, against which we found significant statistical evidence, P value 0.034. Our approach also delivered a nonparametric estimator $\hat{B}_D(t)$ along with 95% pointwise confidence bands displayed in Figure 3. From the figure, we clearly observe an excess hazard rate for experiencing the safety outcome severe/worse sign/symptom over time in the add-NRTI group compared to the omit-group. Under a constant hazards difference model, our approach estimated a hazards difference of 0.0039 (0.0002, 0.0075), P value 0.0391. Although there is significant evidence of a time-varying effect, as indicated by our goodness-of-fit test rejecting the constant effect model (P value 0.041) in favor of a time varying effect.

4 Discussion

Tashima, L. M. Smeaton, et al. (2015) concluded that treatment-experienced patients with HIV infection starting a new optimized regimen can safely omit NRTIs without compromising virologic efficacy. However, they failed to detect a significant safety issue due to a loss of power due to treatment switching. We were able to uncover a treatment difference in terms of safety which amounts to a causal effect, by leveraging randomization as an instrumental variable under a structural cumulative survival model. Our analysis suggests that adding NRTIs is not only unnecessary to achieve optimal outcomes but may in fact significantly increase the risk for a safety outcome.

Our IV analysis does not rely on unconfoundedness assumption on which standard adjustment such as inverse-probability weighting rely upon. In fact, similar to Robins g-estimation of an rank preserving structural failure time model, by leveraging randomization, our estimator can account for measured and unmeasured confounding factors driving the treatment switching mechanism. Importantly, unlike RPSFTMs, our approach obviates the need for artificial censoring for unbiasedness. In simulations studies, we found that inferences based on the proposed approach are reliable in moderate to large samples but may be somewhat conservative in small samples. We should note that as described in Section 2.2, our IV analysis does rely on three key assumptions: IV relevance, IV independence and exclusion restriction. In randomized settings such as the OPTIONS study, IV relevance and IV independence are generally expected to hold, however the exclusion restriction may be violated when, as in the OPTIONS trial, the randomized study is unblinded. This is because, one cannot rule out the possibility that patients or study physicians might modify a patient's treatment course as a result of their randomized treatment assignment in a manner that may in turn directly impact the outcome in view, thus inducing an unintended direct effect of the randomized treatment on the outcome. Such violation of the exclusion restriction can invali-

date an IV analysis. We do not believe this to be a serious issue in the OPTIONS trial given that as reported in Section 3, there was no significant empirical evidence of a statistically significant association between the randomized treatment assignment on the safety outcome. Nevertheless, it would be of interest to extend the GENIUS approach of Tchetgen Tchetgen, Sun, and Walter (2017), which is robust to violation of the exclusion restriction assumption, to the current setting in order to assess possible sensitivity of the proposed methodology to departure from this assumption. We plan to further investigate this possibility in future work.

References

- J. D. Angrist and A. B. Krueger (2001). Instrumental variables and the search for identification: from supply and demand to natural experiments. *Journal of Economic Perspectives* 15(4), pp. 69–85.
- J. Bowden, S. Seaman, X. Huang, and I. R. White (2016). Gaining power and precision by using model-based weights in the analysis of late stage cancer trials with substantial treatment switching. *Statistics in Medicine* 35(9), pp. 1423–1440.
- J. Cuzick, R. Edwards, and N. Segnan (1997). Adjusting for non-compliance and contamination in randomized clinical trials. *Statistics in Medicine* 16(9), pp. 1017–1029.
- G. Davey Smith and S. Ebrahim (2003). ‘Mendelian randomization’: can genetic epidemiology contribute to understanding environmental determinants of disease? *International Journal of Epidemiology* 32(1), pp. 1–22.
- G. D. Demetri, A. T. van Oosterom, C. R. Garrett, M. E. Blackstein, M. H. Shah, J. Verweij, G. McArthur, I. R. Judson, M. C. Heinrich, J. A. Morgan, et al. (2006). Efficacy and safety of sunitinib in patients with advanced gastrointestinal stromal tumour after failure of imatinib: a randomised controlled trial. *The Lancet* 368(9544), pp. 1329–1338.

- S. Jimenez, D. Lai, K. L. Gould, and B. R. Davis (2017). Evaluating the effects of treatment switching with randomization as an instrumental variable in a randomized controlled trial. *Communications in Statistics-Simulation and Computation* 46(6), pp. 4966–4990.
- M. M. Joffe (2001). Administrative and artificial censoring in censored regression models. *Statistics in Medicine* 20(15), pp. 2287–2304.
- M. M. Joffe, W. P. Yang, and H. Feldman (2012). G-estimation and artificial censoring: problems, challenges, and applications. *Biometrics* 68(1), pp. 275–286.
- M. B. Katan (2004). Apolipoprotein E isoforms, serum cholesterol, and cancer. *International Journal of Epidemiology* 33(1), pp. 9–9.
- N. R. Latimer, K. R. Abrams, P. C. Lambert, J. P. Morden, and M. J. Crowther (2018). Assessing methods for dealing with treatment switching in clinical trials: a follow-up simulation study. *Statistical Methods in Medical Research* 27(3), pp. 765–784.
- N. R. Latimer, K. Abrams, P. Lambert, M. Crowther, A. Wailoo, J. Morden, R. Akehurst, and M. Campbell (2017). Adjusting for treatment switching in randomised controlled trials—a simulation study and a simplified two-stage method. *Statistical Methods in Medical Research* 26(2), pp. 724–751.
- J. Li, J. Fine, and A. Brookhart (2015). Instrumental variable additive hazards models. *Biometrics* 71(1), pp. 122–130.
- D. Y. Lin, L.-J. Wei, I. Yang, and Z. Ying (2000). Semiparametric regression for the mean and rate functions of recurrent events. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 62(4), pp. 711–730.
- T. Martinussen, S. Vansteelandt, E. J. Tchetgen Tchetgen, and D. M. Zucker (2017). Instrumental variables estimation of exposure effects on a time-to-event endpoint using structural cumulative survival models. *Biometrics* 73(4), pp. 1140–1149.

- J. P. Morden, P. C. Lambert, N. Latimer, K. R. Abrams, and A. J. Wailoo (2011). Assessing methods for dealing with treatment switching in randomised controlled trials: a simulation study. *BMC Medical Research Methodology* 11(1), pp. 1–20.
- R. J. Motzer, B. Escudier, S. Oudard, T. E. Hutson, C. Porta, S. Bracarda, V. Grünwald, J. A. Thompson, R. A. Figlin, N. Hollaender, et al. (2008). Efficacy of everolimus in advanced renal cell carcinoma: a double-blind, randomised, placebo-controlled phase III trial. *The Lancet* 372(9637), pp. 449–456.
- J. M. Robins (1987). A graphical approach to the identification and estimation of causal parameters in mortality studies with sustained exposure periods. *Journal of Chronic Diseases* 40, 139S–161S.
- (1986). A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect. *Mathematical Modelling* 7(9-12), pp. 1393–1512.
- (2000). Marginal structural models versus structural nested models as tools for causal inference. *Statistical models in epidemiology, the environment, and clinical trials*. Springer, pp. 95–133.
- J. M. Robins and D. M. Finkelstein (2000). Correcting for noncompliance and dependent censoring in an AIDS clinical trial with inverse probability of censoring weighted (IPCW) log-rank tests. *Biometrics* 56(3), pp. 779–788.
- J. M. Robins and S. Greenland (1994). Adjusting for differential rates of prophylaxis therapy for PCP in high-versus low-dose AZT treatment arms in an AIDS randomized trial. *Journal of the American Statistical Association* 89(427), pp. 737–749.
- J. M. Robins and A. A. Tsiatis (1991). Correcting for non-compliance in randomized trials using rank preserving structural failure time models. *Communications in Statistics-Theory and Methods* 20(8), pp. 2609–2631.

- D. Schoenfeld (1981). The asymptotic properties of nonparametric tests for comparing survival distributions. *Biometrika* 68(1), pp. 316–319.
- S. Seaman, O. Dukes, R. Keogh, and S. Vansteelandt (2020). Adjusting for time-varying confounders in survival analysis using structural nested cumulative survival time models. *Biometrics* 76(2), pp. 472–483.
- T. R. Sullivan, N. R. Latimer, J. Gray, M. J. Sorich, A. B. Salter, and J. Karnon (2020). Adjusting for Treatment Switching in Oncology Trials: A Systematic Review and Recommendations for Reporting. *Value in Health* 23(3), pp. 388–396.
- K. T. Tashima, K. R. Mollan, L. Na, R. T. Gandhi, K. L. Klingman, C. J. Fichtenbaum, A. Andrade, V. A. Johnson, J. J. Eron, L. Smeaton, et al. (2015). Regimen selection in the OPTIONS trial of HIV salvage therapy: drug resistance, prior therapy, and race–ethnicity determine the degree of regimen complexity. *HIV clinical trials* 16(4), pp. 147–156.
- K. T. Tashima, L. M. Smeaton, C. J. Fichtenbaum, A. Andrade, J. J. Eron, R. T. Gandhi, V. A. Johnson, K. L. Klingman, J. Ritz, S. Hodder, et al. (2015). HIV salvage therapy does not require nucleoside reverse transcriptase inhibitors: a randomized, controlled trial. *Annals of Internal Medicine* 163(12), pp. 908–917.
- E. J. Tchetgen Tchetgen, B. Sun, and S. Walter (2017). The GENIUS approach to robust Mendelian randomization inference. *arXiv preprint arXiv:1709.07779*.
- E. J. Tchetgen Tchetgen, S. Walter, S. Vansteelandt, T. Martinussen, and M. Glymour (2015). Instrumental variable estimation in a survival context. *Epidemiology (Cambridge, Mass.)* 26(3), pp. 402–410.
- A. W. Van der Vaart (2000). *Asymptotic Statistics*. Vol. 3. Cambridge university press.
- A. W. Van Der Vaart and J. A. Wellner (1996). Weak convergence. *Weak Convergence and Empirical Processes*. Springer, pp. 16–28.
- S. Vansteelandt, M. Joffe, et al. (2014). Structural nested models and G-estimation: the partially realized promise. *Statistical Science* 29(4), pp. 707–731.

- J. M. Wooldridge (2010). *Econometric Analysis of Cross Section and Panel Data*. MIT press.
- T. Yamaguchi and Y. Ohashi (2004). Adjusting for differential proportions of second-line treatment in cancer clinical trials. Part I: structural nested models and marginal structural models to test and estimate treatment arm effects. *Statistics in Medicine* 23(13), pp. 1991–2003.
- A. Ying (2021). *ivsacim: Instrumental Variables Estimation Using Structural Additive Cumulative Intensity Models*. R package version 1.1.
- A. Ying, R. Xu, and J. Murphy (2019). Two-stage residual inclusion for survival data and competing risks—An instrumental variable approach with application to SEER-Medicare linked data. *Statistics in Medicine* 38(10), pp. 1775–1801.

Table 1: Time-constant exposure effect. Bias of $\hat{B}_D(t)$, empirical standard error, $see(\hat{B}_D(t))$, average estimated standard error, $sd(\hat{B}_D(t))$, and coverage probability of 95% confidence intervals $CP(\hat{B}_D(t))$, Bias of $\hat{\beta}$, empirical standard error, $see(\hat{\beta})$, average estimated standard error, $sd(\hat{\beta})$, and coverage probability of 95% pointwise confidence intervals $CP(\hat{\beta})$, in function of sample size n .

Sample Sizes		t = 1	t = 2	t = 3		
n = 800	Bias($\hat{B}_D(t)$)	0.0005	-0.0052	-0.0342	Bias($\hat{\beta}$)	-0.0044
	$see(\hat{B}_D(t))$	0.0670	0.1407	0.3064	$see(\hat{\beta})$	0.1214
	$sd(\hat{B}_D(t))$	0.0721	0.1492	0.3098	$sd(\hat{\beta})$	0.1303
	95% CP($\hat{B}_D(t)$)	95.5	96.8	97.8	95% CP($\hat{\beta}$)	97.6
n = 1600	Bias($\hat{B}_D(t)$)	-0.0021	-0.0113	-0.0244	Bias($\hat{\beta}$)	-0.0073
	$see(\hat{B}_D(t))$	0.0553	0.1102	0.2256	$see(\hat{\beta})$	0.0545
	$sd(\hat{B}_D(t))$	0.0546	0.1105	0.2228	$sd(\hat{\beta})$	0.0532
	95% CP($\hat{B}_D(t)$)	94.7	95.2	95.7	95% CP($\hat{\beta}$)	95.7

Table 2: Primary outcome of virologic failure and NRTI assignment change.

	Patients, n (%)		Difference (95% CI), percentage points
	Add NRTIs (n = 180)	Omit NRTIs (n = 177)	
First severe or worse sign or symptom	51 (28.3)	35 (19.8)	8.6 (-0.8 to 17.9)
Change in NRTI assignment	10 (5.3)	17 (9.5)	-4.0 (-10.1 to 2.0)
Change in NRTI assignment before failure	9 (3.4)	8 (2.8)	0.5 (-4.4 to 5.4)

Figure 1: Time to first severe or worse sign or symptom between groups.

Figure 2: Time to change NRTI assignment.

Figure 3: Estimated causal effect of NRTIs on time to first severe or worse sign or symptom, along with 95% pointwise confidence bands. The straight line corresponds to the constant effects estimator.

Figure 4: Test process $\sqrt{n}(\hat{B}_D(t) - B_D(t))$ along with 20 resampled processes from its limit distribution under the null.

A Proof of Proposition 1

It is straightforward to show that (1) implies that

$$\begin{aligned}
 & \mathbb{E}(d\tilde{N}_{\bar{D}(t_{m-1}),0}(t) | \bar{D}(t_m), Z, L, T_{\bar{D}(t_{m-1}),0} \geq t) \\
 &= \mathbb{E}(d\tilde{N}_{\bar{D}(t_m),0}(t) | \bar{D}(t_m), Z, L, T_{\bar{D}(t_m),0} \geq t) \\
 &\quad - D(t_m) \mathbb{1}(t_m \leq t < t_{m+1}) dB_D(t).
 \end{aligned} \tag{6}$$

For any time t satisfying $t_m \leq t < t_{m+1}$, we have that under model (6)) and independent censoring

$$\begin{aligned}
 & \mathbb{E} \left\{ Z^c e^{\int_0^{t-} D(s) dB_D(s)} Y(t) \left[dN(t) - D(t_m) dB_D(t) \right] \right\} \\
 = & \mathbb{E} \left\{ Z^c e^{\int_0^{t-} D(s) dB_D(s)} Y(t) \left[d\tilde{N}_{\bar{D}(t_m),0}(t) - D(t_m) dB_D(t) \right] \right\} \\
 = & \mathbb{E} \left\{ Z^c e^{\int_0^{t-} D(s) dB_D(s)} Y(t) \left[\mathbb{E}(d\tilde{N}_{\bar{D}(t_m),0}(t) | \bar{D}(t_m), Z, L, \tilde{T}_{\bar{D}(t_m),0} \geq t) - D(t_m) dB_D(t) \right] \right\} \\
 = & \mathbb{E} \left\{ Z^c e^{\int_0^{t_m-} D(s) dB_D(s)} \frac{\mathbb{1}(\tilde{T}(\bar{D}(t_m), 0) \geq t)}{e^{-\int_{t_m}^{t-} D(s) dB_D(s)}} \mathbb{1}(C \geq t) \right. \\
 & \cdot \left. \left[\mathbb{E}(d\tilde{N}_{\bar{D}(t_m),0}(t) | \bar{D}(t_m), Z, L, \tilde{T}_{\bar{D}(t_m),0} \geq t) - D(t_m) dB_D(t) \right] \right\} \\
 = & \mathbb{E} \left\{ Z^c e^{\int_0^{t_m-} D(s) dB_D(s)} \mathbb{P}(\tilde{T}(\bar{D}(t_{m-1}), 0) \geq t | \bar{D}(t_m), Z, L, \tilde{T}_{\bar{D}(t_m),0} \geq t) \mathbb{1}(C \geq t) \right. \\
 & \cdot \left. \left[\mathbb{E}(d\tilde{N}_{\bar{D}(t_m),0}(t) | \bar{D}(t_m), Z, L, \tilde{T}_{\bar{D}(t_m),0} \geq t) - D(t_m) dB_D(t) \right] \right\} \\
 = & \mathbb{E} \left\{ Z^c e^{\int_0^{t_m-} D(s) dB_D(s)} Y_{\bar{D}(t_{m-1}),0}(t) \mathbb{1}(C \geq t) \right. \\
 & \cdot \left. \left[\mathbb{E}(d\tilde{N}_{\bar{D}(t_m),0}(t) | \bar{D}(t_m), Z, L, \tilde{T}_{\bar{D}(t_m),0} \geq t) - D(t_m) dB_D(t) \right] \right\} \\
 = & \mathbb{E} \left\{ Z^c e^{\int_0^{t_m-} D(s) dB_D(s)} Y_{\bar{D}(t_{m-1}),0}(t) \mathbb{1}(C \geq t) \right. \\
 & \cdot \left. \left[\mathbb{E}(d\tilde{N}_{\bar{D}(t_{m-1}),0}(t) | \bar{D}(t_m), Z, L, \tilde{T}_{\bar{D}(t_{m-1}),0} \geq t) \right] \right\} \\
 = & \mathbb{E} \left\{ Z^c e^{\int_0^{t_m-} D(s) dB_D(s)} Y_{\bar{D}(t_{m-1}),0}(t) \mathbb{1}(C \geq t) d\tilde{N}_{\bar{D}(t_{m-1}),0}(t) \right\}.
 \end{aligned}$$

Now we can repeat these steps in order to blip down the effect of treatment $D(t_{m-1})$ at t_{m-1} . However, note that for $t_{m-1} < t_m \leq t$, this effect is null by assumption, and $\int_0^{t_m-} D(s) dB_D(s)$

is a function of $\bar{D}(t_{m-1})$. It follows that

$$\begin{aligned}
 & \mathbb{E} \left\{ Z^c e^{\int_0^{t_m-} D(s) dB_D(s)} Y_{\bar{D}(t_{m-1}),0}(t) \mathbb{1}(C \geq t) d\tilde{N}_{\bar{D}(t_{m-1}),0}(t) \right\} \\
 & \mathbb{E} \left\{ Z^c e^{\int_0^{t_m-} D(s) dB_D(s)} Y_{\bar{D}(t_{m-1}),0}(t) \mathbb{1}(C \geq t) \mathbb{E}(d\tilde{N}_{\bar{D}(t_{m-1}),0}(t) | \bar{D}(t_{m-1}), Z, L, \tilde{T}_{\bar{D}(t_{m-1}),0} \geq t) \right\} \\
 & \mathbb{E} \left\{ Z^c e^{\int_0^{t_{m-1}-} D(s) dB_D(s)} \frac{\mathbb{P}(T_{\bar{D}(t_{m-1}),0} \geq t | \bar{D}(t_{m-1}), Z, L, \tilde{T}_{\bar{D}(t_{m-1}),0} \geq t)}{e^{-\int_{t_{m-1}}^{t \wedge (t_{m-1})} D(s) dB_D(s)}} \mathbb{1}(C \geq t) \right. \\
 & \quad \cdot \mathbb{E}(d\tilde{N}_{\bar{D}(t_{m-1}),0}(t) | \bar{D}(t_{m-1}), Z, L, \tilde{T}_{\bar{D}(t_{m-1}),0} \geq t) \left. \right\} \\
 & \mathbb{E} \left\{ Z^c e^{\int_0^{t_{m-1}-} D(s) dB_D(s)} \mathbb{P}(T_{\bar{D}(t_{m-2}),0} \geq t | \bar{D}(t_{m-1}), Z, L, \tilde{T}_{\bar{D}(t_{m-1}),0} \geq t) \mathbb{1}(C \geq t) \right. \\
 & \quad \cdot \mathbb{E}(d\tilde{N}_{\bar{D}(t_{m-1}),0}(t) | \bar{D}(t_{m-1}), Z, L, \tilde{T}_{\bar{D}(t_{m-1}),0} \geq t) \left. \right\} \\
 & = \mathbb{E} \left\{ Z^c e^{\int_0^{t_{m-1}-} D(s) dB_D(s)} Y_{\bar{D}(t_{m-2}),0}(t) \mathbb{1}(C \geq t) \mathbb{E}(d\tilde{N}_{\bar{D}(t_{m-1}),0}(t) | \bar{D}(t_{m-1}), Z, L, \tilde{T}_{\bar{D}(t_{m-1}),0} \geq t) \right\} \\
 & = \mathbb{E} \left\{ Z^c e^{\int_0^{t_{m-1}-} D(s) dB_D(s)} Y_{\bar{D}(t_{m-2}),0}(t) \mathbb{1}(C \geq t) \mathbb{E}(d\tilde{N}_{\bar{D}(t_{m-2}),0}(t) | \bar{D}(t_{m-1}), Z, L, \tilde{T}_{\bar{D}(t_{m-2}),0} \geq t) \right\} \\
 & = \mathbb{E} \left\{ Z^c e^{\int_0^{t_{m-1}-} D(s) dB_D(s)} Y_{\bar{D}(t_{m-2}),0}(t) \mathbb{1}(C \geq t) d\tilde{N}_{\bar{D}(t_{m-2}),0}(t) \right\}.
 \end{aligned}$$

A recursive application of the above argument yields

$$\begin{aligned}
 & \mathbb{E} \left\{ Z^c \exp \left(\int_0^{t-} D(s) dB_D(s) \right) \left[dN_{\bar{D}(t_m),0}(t) - Y(t) D(t_m) dB_D(t) \right] \right\} \\
 & = \mathbb{E} \left\{ Z^c Y_0(t) \mathbb{1}(C \geq t) dN_0(t) \right\} = 0,
 \end{aligned}$$

by the IV independence assumption.

Consequently, (4) holds provided that $\mathbb{E}\{Z^c Y(t) \exp[\int_0^{t-} D(s) dB_D(s)] D(t)\} \neq 0$ for all $t > 0$. □

B Asymptotic Results

Throughout the paper, $\int_s^t$ is evaluated on $(s, t]$. The estimator can be written

$$\hat{B}_D(t, \hat{\theta}) = \int_0^t \frac{\sum_{i=1}^n Z_i^c(\hat{\theta}) \exp \left[\int_0^{s-} D_i(u) d\hat{B}_D(u, \hat{\theta}) \right] dN_i(s)}{\sum_{j=1}^n Z_j^c(\hat{\theta}) Y_j(s) \exp \left[\int_0^{s-} D_j(u) d\hat{B}_D(u, \hat{\theta}) \right] D_j(s)}.$$

We introduce additional notation, some technical assumptions and a necessary theorem.

Let $\mu(L; \theta) = \mathbb{E}(Z|L; \theta)$ be the conditional mean of the instrument given observed covariates L , which is a function of an unknown finite-dimensional parameter θ for which we have available a consistent estimator $\hat{\theta}$. We assume that this estimator is also asymptotically normal with influence function ϵ_i^θ such that $\sqrt{n}(\hat{\theta} - \theta) = \frac{1}{\sqrt{n}} \sum_{i=1}^n \epsilon_i^\theta + o_p(1)$.

We write $\|g\|_\infty = \sup_{t \in [0, \tau]} |g(t)|$ and use $\mathcal{V}(g)$ to denote the total variation of g over the interval $[0, \tau]$. Let $B^\circ(t)$ denote the true value of $B(t)$, and let $M^\circ = \|B^\circ\|_\infty < \infty$.

Assumption 5. *The instrument Z is bounded by Z_{max} .*

Define $a(s, H(\cdot)) = \mathbb{E}[Z^c Y(s) \exp(\int_0^{s-} D(u) dH(u)) D(s)]$ for any $H \in \mathbb{H}$, where $\mathbb{H}$ is the set of functions on $[0, \tau]$ that have total variations bounded by some $M > M^\circ$. We make the following assumption.

Assumption 6. *There exists $\nu > 0$ such that $\inf_{s \in [0, \tau], H \in \mathbb{H}} a(s, H) > \nu$.*

Define

$$\begin{aligned} A(s, H(\cdot)) &= \frac{1}{n} \sum_{i=1}^n Z_i^c Y_i(s) \exp \left(\int_0^{s-} D_i(u) dH(u) \right) D_i(s), \\ \Upsilon_n(H(\cdot), t) &= \int_0^t \frac{\sum_{i=1}^n Z_i^c \exp \left[\int_0^{s-} D_i(u) dH(u) \right] dN_i(s)}{nA(s, H)}, \\ \bar{\Upsilon}_n(H(\cdot), t) &= \int_0^t \frac{\sum_{i=1}^n Z_i^c \exp \left[\int_0^{s-} D_i(u) dH(u) \right] dN_i(s)}{na(s, H)}, \end{aligned}$$

and

$$\Upsilon(H(\cdot), t) = \int_0^t \frac{\mathbb{E} \left\{ Z^c \exp \left[\int_0^{s-} D(u) dH(u) \right] dN(s) \right\}}{a(s, H)}.$$

Note that $B_n(t, \theta)$ and $B^\circ(t, \theta)$ are solutions to $B(t) = \Upsilon_n(B, t)$ and $B(t) = \Upsilon(B, t)$, respectively. We note for later reference that for any two functions H_1 and H_2 ,

$$\|\Upsilon(\xi(B_1), t) - \Upsilon(\xi(B_2), t)\|_\infty \leq 4\tau Z_{\max} \exp(M)/\nu \|B_1 - B_2\|_\infty. \quad (30)$$

We further assume

Assumption 7. $B^\circ(t)$ is the unique continuous solution to the equation $B(t) = \Upsilon(B, t)$.

We shall use Helly's selection theorem to establish Theorem 1.

Theorem (Helly's Selection Theorem). *Let $\{f_n\}$ be a sequence of functions on $[0, \tau]$ such that $\|f_n\|_\infty \leq A_1$ and $\sup_n \mathcal{V}(f_n) \leq A_2$, where A_1 and A_2 are finite constants. Then*

1. *There exists a subsequence $\{f_{n_j}\}$ of $\{f_n\}$ which converges pointwise to some function f .*
2. *If f is continuous, the convergence is uniform.*

B.1 Proof of Theorem 1

We give a roadmap for proof of consistency.

1. We construct a modified version $\tilde{B}_n(t, \theta)$ of the estimator that is uniformly of bounded variation over n .
2. We show that

$$\sup_{s \in [0, \tau], H \in \mathbb{H}} |\Upsilon_n(H, s) - \Upsilon(H, s)| \rightarrow 0 \quad \text{a.s.} \quad (31)$$

3. These, together with an application of the Helly's Selection Theorem imply $\|\tilde{B}_n(t, \theta) - B^\circ(t, \theta)\|_\infty \rightarrow 0$ a.s..
4. Finally we show that $\tilde{B}_n(t)$ is equal to $\hat{B}_n(t)$ in large samples, and therefore, $\sup_{t \in [0, \tau]} |\hat{B}_n(t, \theta) - B_D^\circ(t, \theta)| \rightarrow 0$, a.s..

STEP 1: Let $\xi(y) = \text{sgn}(y) \min(|y|, M)$. We define the modified estimator $\tilde{B}_n$ to be the solution to the equation $B(t) = \Upsilon_n(\xi(B), t)$. Note that the treatment process can only take values in 0 and 1. Then the total variation

$$\mathcal{V}(\tilde{B}_n) \leq 4\tau Z_{\max} \exp(M)/\nu$$

Therefore $\tilde{B}_n$ is uniformly of bounded variation over n . Also, since $\tilde{B}_n(0) = 0$, $\tilde{B}_n$ is uniformly bounded, therefore Helly's selection theorem applies. Note that $\Upsilon(\xi(B^\circ), t) = \Gamma(B^\circ, t) = B^\circ(t)$.

STEP 2: We want to show

$$\sup_{s \in [0, \tau], H \in \mathbb{H}} |\tilde{\Upsilon}_n(H, s) - \Upsilon(H, s)| \rightarrow 0, \text{ a.s..} \quad (32)$$

To this end, we first define

$$\psi_{H,t}(T, \Delta, Z, L, D(\cdot)) := \frac{Z^c \exp(\int_0^{T-} D(s) dH(s)) N(t)}{a(T, H(\cdot))},$$

and therefore $\tilde{\Upsilon}_n(H, t) = \mathbb{P}_n \psi_{H,t}$ and $\Upsilon(H, t) = \mathbb{P} \psi_{H,t}$. Then it suffices to show that $\{\psi_{H,t} : H \in \mathbb{H}, t \in [0, \tau]\}$ is Gilvenko-Cantelli. This result is an immediate consequence of the following facts:

1. All functions that are of variation bounded form a Donsker class Van der Vaart 2000, Example 19.11. Therefore the function class $\{Z^c \exp(\int_0^{T-} D(s) dH(s)) N(t) : H \in \mathbb{H}, t \in [0, \tau]\}$ is Donsker.

2. The function $1/a(t, H)$ is a uniformly bounded, measurable function by Assumption 6.
3. A Donsker class multiplied by a uniformly bounded, measurable function remains Donsker Van Der Vaart and Wellner 1996, Example 2.10.10.
4. Donsker classes are Gilvenko-Cantelli.

Similary we can show that $\{\phi_{H,s} : H \in \mathbb{H}, s \in [0, \tau]\}$ is also Gilvenko-Cantelli, where

$$\phi_{H,s}(T, \Delta, Z, L, D(\cdot)) := Z^c Y(s) \exp \left(\int_0^{s-} D(u) dH(u) \right) D(s),$$

and hence

$$\sup_{s \in [0, \tau], H \in \mathbb{H}} |A(H, s) - a(H, s)| \rightarrow 0 \text{ a.s..}$$

This, together with Assumption 6 we get

$$\inf_{s \in [0, \tau], H \in \mathbb{H}} A(s, H) \geq \nu,$$

for sufficiently large n . We arrive at

$$\sup_{s \in [0, \tau], H \in \mathbb{H}} |\tilde{\Upsilon}_n(H, s) - \Upsilon_n(H, s)| \rightarrow 0, \text{ a.s..}$$

This together with (32), implies (31).

STEP 3: We prove $\|\tilde{B}_n(t, \theta) - B(t, \theta)\|_\infty \rightarrow 0$ a.s. by contradiction. Without loss of generality, we assume

$$\liminf_{n \rightarrow \infty} \|\tilde{B}_n(t, \theta) - B^\circ(t, \theta)\|_\infty > 0.$$

By Helly's Selection Theorem, there exists a subsequence $\{n_j\}$ such that $\tilde{B}_{n_j}(t, \theta)$ converges to some limit $B(t)$. We further claim this limit is continuous, in fact, Lipschitz continuous.

To see this, for any $t_1 < t_2$ in $[0, \tau]$, when n_j is large enough

$$\begin{aligned} |B(t_2) - B(t_1)| &\leq |B(t_2) - \tilde{B}_{n_j}(t_2)| + |\tilde{B}_{n_j}(t_2) - \tilde{B}_{n_j}(t_1)| + |\tilde{B}_{n_j}(t_1) - B(t_1)| \\ &\leq \left| \int_{t_1}^{t_2} \frac{\sum_{i=1}^{n_j} Z_i^c \exp(\int_0^{s-} D_i(u) d\xi(B_{n_j})(u)) dN_i(s)}{n_j A(s, \xi(B_{n_j}))} \right| + 2\varepsilon \\ &\leq 4Z_{\max} \exp(M)/\nu(t_2 - t_1) + 2\varepsilon, \end{aligned}$$

for any $\varepsilon > 0$. Therefore, by the second part of Helly's theorem, the convergence of the sub-subsequence is uniform. Going further, the limit B satisfies $B(t) = \Upsilon(B, t)$ since

$$\begin{aligned} & \|B(t) - \Upsilon(\xi(B), t)\|_\infty \\ \leq & \|B(t) - \tilde{B}_{n_j}(t)\|_\infty + \|\tilde{B}_{n_j}(t) - \Upsilon_{n_j}(\xi(\tilde{B}_{n_j}(t)), t)\|_\infty \\ & + \|\Upsilon_{n_j}(\xi(\tilde{B}_{n_j}(t)), t) - \Upsilon(B, t)\|_\infty + \|\Upsilon(\xi(\tilde{B}_{n_j}(t)), t) - \Upsilon(B, t)\|_\infty \rightarrow 0, \end{aligned}$$

where the last statement results from the uniform convergence of $\tilde{B}_{n_j}(t)$, the definition of $\tilde{B}_{n_j}(t)$, (31) and (30). Now since the solution to $B(t) = \Upsilon(B, t)$ is unique by Assumption 7 and is given by B° , this leads to a contradiction. It hence follows that $\|\tilde{B}_n(t, \theta) - B(t, \theta)\|_\infty \rightarrow 0$ a.s..

STEP 4: Since $\|B^\circ\|_\infty \leq M^\circ$ and we just showed $\|\tilde{B}_D(t, \theta) - B_D(t, \theta)\|_\infty \rightarrow 0$ a.s., for sufficiently large n we have $\|\tilde{B}_n\|_\infty \leq M^\circ + \frac{1}{2}(M - M^\circ) < M$, and therefore $\xi(\tilde{B}_n) = \tilde{B}_n$. Therefore, for sufficiently large n , $\tilde{B}_n(t)$ satisfies $\tilde{B}_n = \Upsilon(\xi(\tilde{B}_n), t) = \Upsilon(\tilde{B}_n, t)$, or in other words, $\hat{B}_n = \tilde{B}_n$. We thus have $\sup_{t \in [0, \tau]} |\hat{B}_n(t, \theta) - B^\circ(t, \theta)| \rightarrow 0$, a.s.

The consistency of $\hat{B}_D(t, \hat{\theta})$ then follows immediately by a Taylor series expansion since $\hat{\theta}$ is consistent for θ .

B.2 Proof of Theorem 2

We aim to provide a sum of i.i.d. representation of $\sqrt{n}(\hat{B}(t, \hat{\theta}) - B(t, \theta))$ heuristically. To that end, we rewrite the normalized residual $V_n(t, \theta) := \sqrt{n}(\hat{B}_D(t, \theta) - B_D(t, \theta))$ at a fixed θ as a solution to a Volterra equation, which shall yield a sum of i.i.d. representation at each time point t in this case. To establish this, we integrate by part the Riemann–Stieltjes integral (the integral is interpreted pathwise for the Càdlàg stochastic process $D_i(s)$, B_D is assumed to be continuous),

$$\int_0^{s-} D_i(u) dB_D(u) = D_i(s-)B_D(s) - \int_0^{s-} B_D(u) dD_i(u) =: G_{1,i}(s-) + G_{2,i}(s-). \quad (39)$$

By Theorem 1 and an application of Slutsky's Theorem, we can write

$$V_n(t, \theta) = \frac{1}{\sqrt{n}} \int_0^t \sum_{i=1}^n H_i(s, \hat{B}_D)(dN_i(s) - D_i(s)dB_D(s, \theta)) + o_P(1),$$

where

$$H_i(s, B_D) := \frac{Z_i^c \exp(\int_0^{s-} D_i(u) d\hat{B}_D(u, \theta))}{\mathbb{E}(Z^c Y(s) \exp(\int_0^{s-} D(u) dB_D(u, \theta)) D(s))}.$$

Now we are ready to write down the Volterra equation. By consistency and a Taylor expansion, it is easy to see that

$$\begin{aligned} V_n(t, \theta) &= \frac{1}{\sqrt{n}} \int_0^t \sum_{i=1}^n H_i(s, B_D)[dN_i(s) - D_i(s)dB_D(s, \theta)] \\ &\quad + \int_0^t V_n(s-, \theta)(1 + o_P(1)) \sum_{i=1}^n \frac{\partial H_i(s, B_D)}{\partial B_D(s-, \theta)} dN_i(s) + o_P(1), \end{aligned}$$

which by (39) yields

$$\begin{aligned} V_n(t, \theta) &= \frac{1}{\sqrt{n}} \int_0^t \sum_{i=1}^n H_i(s, B_D)[dN_i(s) - D_i(s)dB_D(s)] \\ &\quad + \int_0^t V_n(s-, \theta)(1 + o_P(1)) \sum_{i=1}^n \frac{\partial H_i(s, B_D)}{\partial G_{1,i}(s-)} D_i(s-)dN_i(s) \\ &\quad + \int_0^t \sum_{i=1}^n \frac{\partial H_i(s, B_D)}{\partial G_{2,i}(s-)} \int_0^{s-} V_n(u, \theta)(1 + o_P(1))dD_i(u)dN_i(s) + o_P(1) \\ &= \frac{1}{\sqrt{n}} \int_0^t \sum_{i=1}^n H_i(s, B_D)[dN_i(s) - D_i(s)dB_D(s)] \\ &\quad + \int_0^t V_n(s-, \theta)(1 + o_P(1)) \sum_{i=1}^n \frac{\partial H_i(s, B_D)}{\partial G_{1,i}(s-)} D_i(s-)dN_i(s) \\ &\quad + \int_0^t V_n(s, \theta) \sum_{i=1}^n \int_{s+}^t \frac{\partial H_i(u, B_D)}{\partial G_{2,i}(u-)} dN_i(u)dD_i(s), \end{aligned}$$

where the last equation follows by an application of Fubini's theorem. Together we have the

Volterra-equation,

$$\begin{aligned} V_n(t, \theta) &= \frac{1}{\sqrt{n}} \int_0^t \sum_{i=1}^n H_i(s, B_D) [dN_i(s) - D_i(s)dB_D(s)] \\ &\quad + \int_0^t V_n(s-, \theta) \sum_{i=1}^n \left\{ \frac{\partial H_i(s, B_D)}{\partial G_{1,i}(s-)} D_i(s-) dN_i(s) + \left[\int_{s+}^t \frac{\partial H_i(u, B_D)}{\partial G_{2,i}(u-)} dN_i(u) \right] dD_i(s) \right\}. \end{aligned}$$

which admits a solution with explicit form given by

$$V_n(t, \theta) = \frac{1}{\sqrt{n}} \int_0^t \mathcal{F}(s, t) \sum_{i=1}^n H_i(s, B_D) [dN_i(s) - D_i(s)dB_D(s)] + o_P(1),$$

where

$$\mathcal{F}(s, t) := \prod_{(s, t]} \left[1 + \sum_{i=1}^n \left\{ \frac{\partial H_i(u, B_D)}{\partial G_{1,i}(u-)} dN_i(u) + \left[\int_u^t \frac{\partial H_i(\cdot, B_D)}{\partial G_{2,i}(\cdot-)} dN_i(\cdot) \right] dD_i(u) \right\} \right].$$

This leads to an i.i.d. representation

$$V_n(t, \theta) = \frac{1}{\sqrt{n}} \sum_{i=1}^n \epsilon_i^B(t) + o_P(1),$$

with the $\epsilon_i^B(t)$'s being zero-mean i.i.d. terms. Specifically

$$\epsilon_i^B(t) := \int_0^t \mathcal{F}(s, t) H_i(s, B_D) [dN_i(s) - D_i(s)dB_D(s)].$$

This, together with a Taylor expansion, gives

$$\begin{aligned} \sqrt{n}(\hat{B}(t, \hat{\theta}) - B(t, \theta)) &= \sqrt{n}(\hat{B}(t, \theta) - B(t, \theta)) + \sqrt{n}(\hat{B}(t, \hat{\theta}) - \hat{B}(t, \theta)) \\ &= \sqrt{n}(\hat{B}(t, \theta) - B(t, \theta)) + \frac{\partial \hat{B}(t, \theta)}{\partial \theta} \Big|_{\hat{\theta}} \sqrt{n}(\hat{\theta} - \theta) + o_P(1). \end{aligned}$$

Finally we have

$$\sqrt{n}(\hat{B}(t, \hat{\theta}) - B(t, \theta)) = \frac{1}{\sqrt{n}} \sum_{i=1}^n \epsilon_i(t, \theta) + o_P(1),$$

where

$$\epsilon_i(t, \theta) := \epsilon_i^B(t) + \frac{\partial B^\circ(t, \theta)}{\partial \theta} \Big|_{\theta} \epsilon_i^\theta.$$

With this i.i.d. representation, similar arguments as in Lin et al. (2000) and Martinussen et al. (2017) can be adopted to establish the convergence of $V_n(t)$ in distribution to a Gaussian process.

C Supplementary Materials for Simulation

We show that our data generating process is compatible with our model (1), where $B_D(t) = 0.1 \cdot t$. To that end, it suffices to show that,

$$\frac{\mathbb{P}\left(\tilde{T}_i(\bar{D}_i(t_m), 0) > t | \bar{D}_i(t_m), Z_i, U_i, \tilde{T}_i \geq t_m\right)}{\mathbb{P}\left(\tilde{T}_i(\bar{D}_i(t_{m-1}), 0) > t | \bar{D}_i(t_m), Z_i, L_i, U_i, \tilde{T}_i \geq t_m\right)} = \exp\left\{-0.1 \int_{t_m}^{t \wedge t_{m+1}} D(t_m) dt\right\}, \quad (52)$$

which after integrating out U_i , yields (1). We work on the numerator of the LHS of (52) first.

$$\begin{aligned} & \mathbb{P}\left(\tilde{T}_i(\bar{D}_i(t_m), 0) > t | \bar{D}_i(t_m), Z_i, U_i, \tilde{T}_i \geq t_m\right) \\ &= \frac{\mathbb{P}\left(\tilde{T}_i(\bar{D}_i(t_m), 0) > t | \bar{D}_i(t_m), Z_i, U_i\right)}{\mathbb{P}\left(\tilde{T}_i > t_m | \bar{D}_i(t_m), Z_i, U_i\right)} = \frac{\mathbb{P}\left(\tilde{T}_i(\bar{D}_i(t_m), 0) > t | Z_i, U_i\right)}{\mathbb{P}\left(\tilde{T}_i(\bar{D}_i(t_m), 0) > t_m | Z_i, U_i\right)} \\ &= \exp\left(-0.25 \cdot (t - t_m) - 0.1 \cdot \int_{t_m}^t d(s) ds - 0 \cdot Z_i \cdot t - 0.15 \cdot U_{2,i} \cdot (t - t_m)\right), \end{aligned}$$

where the second equation follows by consistency and unconfoundedness (conditional on U) guaranteed by our data generating process. Now we switch to the denominator of the LHS of (52),

$$\begin{aligned} & \mathbb{P}\left(\tilde{T}_i(\bar{D}_i(t_{m-1}), 0) > t | \bar{D}_i(t_m), Z_i, U_i, \tilde{T}_i \geq t_m\right) \\ &= \frac{\mathbb{P}\left(\tilde{T}_i(\bar{D}_i(t_{m-1}), 0) > t | \bar{D}_i(t_m), Z_i, U_i\right)}{\mathbb{P}\left(\tilde{T}_i \geq t_m | \bar{D}_i(t_m), Z_i, U_i\right)} = \frac{\mathbb{P}\left(\tilde{T}_i(\bar{D}_i(t_{m-1}), 0) > t | Z_i, U_i\right)}{\mathbb{P}\left(\tilde{T}_i(\bar{D}_i(t_m), 0) \geq t_m | Z_i, U_i\right)} \\ &= \exp\left(-0.25 \cdot (t - t_m) - 0 \cdot Z_i \cdot t - 0.15 \cdot U_{2,i} \cdot (t - t_m)\right), \end{aligned}$$

now (52) is straightforward.

C.1 Variance Estimation

A variance estimator can be obtained by plugging in empirical counterparts to unknown quantities into

$$\hat{\epsilon}_i(t, \hat{\theta}) := \hat{\epsilon}_i^B(t) + \frac{\partial \hat{B}(t, \hat{\theta})}{\partial \theta} \Big|_{\hat{\theta}} \hat{\epsilon}_i^\theta. \quad (53)$$

Note that Volterra's equation (49) is actually a matrix equation since $\hat{B}_D(t)$ only changes over distinct event times $\{t_1, \dots, t_K\}$, by writing $\hat{\epsilon}_{1,i}(t) = dN_i(t) - D_i(t)d\hat{B}_D(t, \hat{\theta})$,

$$\begin{pmatrix} dV_n(t_1) \\ \dots \\ dV_n(t_K) \end{pmatrix} = \frac{1}{\sqrt{n}} \begin{pmatrix} \sum_{i=1}^n H_i(t_1, B_D) \epsilon_{1,i}(t_1) \\ \dots \\ \sum_{i=1}^n H_i(t_K, B_D) \epsilon_{1,i}(t_K) \end{pmatrix} + \begin{pmatrix} 0 & \dots & 0 & 0 \\ \sum_{i=1}^n \frac{\partial H_i(t_1, B_D)}{\partial G_{1,i}(t_1-)} D_i(t_1-) dN_i(t_1) & \dots & 0 & 0 \\ \dots & \dots & \dots & 0 \\ \sum_{i=1}^n \frac{\partial H_i(t_K, B_D)}{\partial G_{1,i}(t_K-)} D_i(t_K-) dN_i(t_K) & \dots & \sum_{i=1}^n \frac{\partial H_i(t_K, B_D)}{\partial G_{1,i}(t_K-)} D_i(t_{K-1}-) dN_i(t_K) & 0 \end{pmatrix} \begin{pmatrix} dV_n(t_1) \\ \dots \\ dV_n(t_K) \end{pmatrix}.$$

That is, we have

$$dV_n(t) = \varepsilon_1 + H_n dV_n(t).$$

The solution is therefore

$$dV_n(t) = (I - H_n)^{-1} \varepsilon_1.$$

Note that $I - H_n$ is a triangular matrix whose inverse can be computed recursively and efficiently.