

Estimating the causal effects of multiple intermittent treatments with application to COVID-19

Liangyuan Hu^{1,*}, Fan Li², Jiayi Ji¹, Himanshu Joshi³, Erick Scott⁴

¹*Department of Biostatistics and Epidemiology, Rutgers University, Piscataway, NJ 08854, U.S.A.*

²*Department of Biostatistics, Yale School of Public Health, New Haven, Connecticut 06510, U.S.A.*

³*Department of Population Health Science and Policy, Icahn School of Medicine at Mount Sinai, New York, New York 10029, U.S.A.*

⁴*Karius Inc., Redwood City, CA 94065, U.S.A.*

liangyuan.hu@rutgers.edu

SUMMARY

To draw real-world evidence about the comparative effectiveness of multiple time-varying treatment regimens on patient survival, we develop a joint marginal structural proportional hazards model and novel weighting schemes in continuous time to account for time-varying confounding and censoring. Our methods formulate complex longitudinal treatments with multiple “start/stop” switches as the recurrent events with discontinuous intervals of treatment eligibility. We derive the weights in continuous time to handle a complex longitudinal dataset on its own terms, without the need to discretize or artificially align the measurement times. We further propose using machine learning models designed for censored survival data with time-varying covariates and the kernel function estimator of the baseline intensity to efficiently estimate the

*To whom correspondence should be addressed.

continuous-time weights. Our simulations demonstrate that the proposed methods provide better bias reduction and nominal coverage probability when analyzing observational longitudinal survival data with irregularly spaced time intervals, compared to conventional methods that require aligned measurement time points. We apply the proposed methods to a large-scale COVID-19 dataset to estimate the causal effects of several COVID-19 treatment strategies on in-hospital mortality or ICU admission, and provide new insights relative to findings from randomized trials.

Key words: Causal inference; Observational study; Stabilized inverse probability weights; Time-varying treatments; Recurrent events; Machine learning

1. INTRODUCTION

The COVID-19 pandemic has been a rapidly evolving crisis challenging global health and economies. Public health experts believe that this pandemic has no true precedent in modern times. While multiple COVID-19 vaccines have been developed across the globe, no consensus has been reached on optimal clinical management of COVID-19 (Yousefi et al., 2020). The lack of evidence for effective treatment options warrants further investigation into the causal effects of multiple COVID-19 treatment strategies currently implemented in clinics. Although randomized controlled trials (RCTs) are considered as the gold standard for evaluating the efficacy of COVID-19 therapies, they are enormously expensive and time consuming, especially in a time of crisis. Stringent inclusion and exclusion criteria also limit the generalizability of RCTs to frailer populations at higher risk for severe morbidity and mortality. To overcome these challenges, we study the causal effects of COVID-19 treatment strategies on patient survival by leveraging the continuously growing observational data collected at the Mount Sinai Health System—New York City’s largest academic medical system. We focus on four commonly used medication classes that are of most clinical interest: (i) remdesivir; (ii) dexamethasone; (iii) anti-inflammatory medications other than cor-

ticosteroids; and (iv) corticosteroids other than dexamethasone.

The complex nature of COVID-19 treatments, owing to differential physician preferences and variability of treatment choices attributable to evolving clinical guidelines, poses three major challenges for statistical analysis of observational data that cannot be easily addressed by existing longitudinal causal inference methods. First, treatment is not randomly allocated and the treatment status over time may depend upon the evolving patient- and disease-specific covariates. Second, the measurement time points during the follow-up are irregularly spaced. Third, there is more than one treatment under consideration. Patients can be simultaneously prescribed to various treatment combinations, or can be switched to a different treatment. Figure 1 illustrates the observed treatment trajectories for nine randomly selected patients during their hospital stays.

While previous work has shown that a continuous-time marginal structural model is effective in addressing time-varying confounding and provides consistent causal effect estimators (Johnson and Tsiatis, 2005; Saarela and Liu, 2016; Hu et al., 2018; Hu and Hogan, 2019; Ryalen et al., 2020), the development has been restricted to a single longitudinal treatment and therefore may not be directly applicable. We consider a joint marginal structural model to accommodate multiple longitudinal treatments in continuous time. To estimate causal parameters in the joint marginal structural model, we derive a novel set of continuous-time stabilized inverse probability weights by casting each treatment process as a counting process for recurrent events, allowing for discontinuous intervals of eligibility. In addition, we propose to use machine learning and smoothing techniques designed for censored survival data to estimate such complex weights. Through simulations, we demonstrate that our approach provides valid causal effect estimates and can considerably alleviate the consequence of unstable inverse probability weights under parametric formulations. We further undertake a detailed analysis of a large longitudinal registry data of clinical management and patient outcomes to investigate the comparative effectiveness of multiple COVID-19 treatments on patient survival.

2. JOINT MARGINAL STRUCTURAL SURVIVAL MODEL

2.1 Notation and set up

We consider a longitudinal observational study with multiple treatments and a right-censored survival outcome. Denote t as the time elapsed from study entry (e.g., hospital admission), t^o the maximum follow-up time, and \mathcal{T} a collection of time points on the interval $[0, t^o]$. Suppose each individual has a p -dimensional covariate process $\{L(t) : t \in \mathcal{T}\}$, some elements of which may be time-varying; by definition, the time-fixed elements of $L(t)$ are constant over \mathcal{T} . Let T denote time to an outcome event of interest such as death, with $\{N^T(t) : t \in \mathcal{T}\}$ as its associated zero-one counting process. We consider W different medication classes (treatments), whose separate and joint causal effects on patient survival are of interest. We use $A_w(t)$ to denote the assignment of treatment $w \in \mathcal{W} = \{1, \dots, W\}$, which can be characterised as the counting process, where we let $A_w(t) = 1$ if an individual is treated with w at time t and $A_w(t) = 0$ otherwise (Johnson and Tsiatis, 2005). Let C denote the time to censoring due to, for example, discharge or loss to follow up. We use the overbar notation to represent the history of a random variable, for example, $\bar{A}_w(t) = \{A_w(s) : 0 \leq s \leq t\}$ corresponds to the history of treatment A_w from hospital admission up to time t and $\bar{L}(t) = \{L(s) : 0 \leq s \leq t\}$ corresponds to the covariate history up to time t . Following the convention in the longitudinal causal inference literature (Robins et al., 2008), we assume the treatment decision is made only after observing the most recent covariate information just prior to the treatment; that is, for a given t , $A_w(t)$ occurs after $L(t)$ for all w .

Let $T^{\bar{a}_1(t), \dots, \bar{a}_W(t)}$ represent the counterfactual failure time to event of interest had an individual been *assigned* treatment history $\{\bar{a}_1(t), \bar{a}_2(t), \dots, \bar{a}_W(t)\}$ rather than the *observed* treatment history $\{\bar{A}_1(t), \bar{A}_2(t), \dots, \bar{A}_W(t)\}$. Similarly, $T^{\bar{A}_1(t), \dots, \bar{A}_W(t)}$ represents the observed failure time to event for an individual given the observed treatment history. We similarly define $C^{\bar{a}_1(t), \dots, \bar{a}_W(t)}$ as the counterfactual censoring time under treatment $\{\bar{a}_1(t), \bar{a}_2(t), \dots, \bar{a}_W(t)\}$. The observed data available for drawing inferences about the distribution of potential outcomes are as follows: the

observed time to outcome event is $T^* = T \wedge C$, with the censoring indicator $\Delta^T = I(T \leq C)$. Note that both the treatment processes $\{A_w(t), w = 1, \dots, W\}$ and the covariate process $\bar{L}(t)$ are defined for all $t \in \mathcal{T}$ but are observed only at discrete and potentially irregularly spaced time points for each individual. For example, individual i may have covariates and treatment status observed at a set of discrete time points from study entry $t = 0$ to his or her last follow-up time $t_{iK_i} \leq t^o$. We denote the set of discrete time points with observed covariate and treatment information for individual i as $\mathcal{T}_i = \{0, t_{i1}, \dots, t_{iK_i}\}$, and therefore the observed covariate and treatment histories become $\bar{L}_i(\mathcal{T}_i) = \{L(t) : t \in \mathcal{T}_i\}$ and $\bar{A}_{w,i}(\mathcal{T}_i) = \{A_{w,i}(t) : t \in \mathcal{T}_i\}$.

2.2 Joint marginal structural model for survival outcomes

We consider a marginal structural model to estimate the joint causal effects of $\bar{A}_1(t), \dots, \bar{A}_W(t)$ on patient survival. The most popular model specification is a marginal structural Cox model, for its flexibility in handling baseline hazard and straightforward software implementation when used in conjunction with the stabilized inverse probability weights (Howe et al., 2012). When there is a strong concern that the proportional hazards assumption may not be satisfied across the marginal distribution of the counterfactual survival times, alternative strategies including the structural additive hazards model or accelerated failure time model can also be considered. For purposes of presenting our methodology, we focus on the marginal structural Cox model but extensions to alternative structural models are possible. For notational brevity but without loss of generality, we first consider $W = 2$ treatments. Expansion of the joint marginal structural model and weighting schema for $W \geq 3$ treatments is discussed in Section 3.4. Specifically, we assume $T^{\bar{a}_1(t), \bar{a}_2(t)}$ follows a marginal structural proportional hazards model of the form

$$\lambda^{T^{\bar{a}_1(t), \bar{a}_2(t)}}(t) = \lambda_0(t) \exp \{ \psi_1 a_1(t) + \psi_2 a_2(t) + \psi_3 a_1(t) a_2(t) \}, \quad (2.1)$$

where $\lambda^{T^{\bar{a}_1(t), \bar{a}_2(t)}}$ is the hazard function for $T^{\bar{a}_1(t), \bar{a}_2(t)}$ and $\lambda_0(t)$ is the unspecified baseline hazard function when treatment A_1 and A_2 are withheld during the study. The parameter ψ_1 encodes

the instantaneous effect of treatment A_1 on $T^{\bar{a}_1(t), \bar{a}_2(t)}$ in terms of log hazard ratio while A_2 is withheld during the study. Similarly, ψ_2 corresponds to the instantaneous treatment effect for A_2 in the absence of A_1 . The multiplicative interaction effect of A_1 and A_2 is captured by ψ_3 . In addition, the hazard function $\lambda^{T^{\bar{a}_1(t), \bar{a}_2(t)}}$ can depend on baseline covariates by elaborating model (2.1) or by using a stratified version of $\lambda_0(t)$. Model (2.1) implicitly assumes that the instantaneous treatment effect is constant in the course of follow-up. This model assumption is reasonable given that the COVID-related hospitalization is generally short and medications are prescribed for days in succession. Finally, model (2.1) is a continuous-time generalization of the discrete-time model considered by Howe et al. (2012) for estimating the joint survival effects of multiple time-varying treatments.

Model (2.1) offers two advantages for the estimation of treatment effects. First, the counterfactual survival function can be expressed as $S^{T^{\bar{a}_1(t), \bar{a}_2(t)}}(t) = \exp \left\{ - \int_0^t \lambda_{T^{\bar{a}_1(t), \bar{a}_2(t)}}(s) ds \right\}$. Therefore, causal contrasts can be performed based on any relevant summary measures of the counterfactual survival curves including median survival times and restricted mean survival times. Second, model (2.1) allows for the estimation of causal effects of interventions defined by varying treatment initiation timing and treatment duration. For example, an intervention may take the form of $\bar{a}_1(t^\dagger) = \{a_1(s) = 1, 0 \leq s \leq t^\dagger\}$, representing prescribing treatment A_1 until t^\dagger (e.g., $t^\dagger = \text{day 6}$). A more complex intervention strategy is $\{\bar{a}_1(t^\dagger), \bar{a}_2(t^\dagger)\} = \{a_1(s) = \mathbb{1}(0 \leq s \leq t^\dagger/2), a_2(s) = \mathbb{1}(t^\dagger/2 < s \leq t^\dagger)\}$, which refers to assigning treatment A_1 until $t^\dagger/2$ and then switching altogether to A_2 until t^\dagger .

3. ESTIMATING STRUCTURAL MODEL PARAMETERS IN CONTINUOUS TIME

To obtain a consistent estimator for $\boldsymbol{\psi} = \{\psi_1, \psi_2, \psi_3\}$ in model (2.1) using longitudinal observational data with two treatments, we introduce the following causal assumptions and maintain them throughout the rest of the article:

(A1) *Consistency.* The observed failure times, $T = \sum_{\mathcal{A}} T^{\bar{a}_1(t), \bar{a}_2(t)} \mathbb{1}(\bar{A}_1(t) = \bar{a}_1(t), \bar{A}_2(t) = \bar{a}_2(t))$, where $\mathcal{A} = \{\bar{a}_1(t), \bar{a}_2(t) : a_1(t) \in \{0, 1\}, a_2(t) \in \{0, 1\}, t \in \mathcal{T}\}$. Similarly for the observed censoring times, $C = \sum_{\mathcal{A}} C^{\bar{a}_1(t), \bar{a}_2(t)} \mathbb{1}(\bar{A}_1(t) = \bar{a}_1(t), \bar{A}_2(t) = \bar{a}_2(t))$. The consistency assumption implies that the observed outcome corresponds to the counterfactual outcome under a specific joint treatment trajectory $\{\bar{a}_1(t), \bar{a}_2(t)\}$ when an individual actually follows treatment $\{\bar{a}_1(t), \bar{a}_2(t)\}$. This is an extension of the consistency assumption developed with a single time-varying treatment (Robins, 1999) to two time-varying treatments.

(A2) *Conditional Exchangeability.* Alternatively referred to as *sequential randomization*, this assumption states that initiation of treatment at time t among those who are still alive and remain in the study is conditionally independent of the counterfactual survival time $T^{\bar{a}_1(t), \bar{a}_2(t)}$ conditional on observed treatment and covariate histories. Mathematically, let $\bar{\mathcal{O}}(t^-) = \{\bar{L}(t^-), \bar{A}_1(t^-), \bar{A}_2(t^-)\}$ denote the observed history up to t^- , then $\forall t \in \mathcal{T}$

$$\lambda^{A_1, A_2} \left(t \mid \bar{\mathcal{O}}(t^-), T > t^-, C > t^-, T^{\bar{a}_1(t), \bar{a}_2(t)} \right) = \lambda^{A_1, A_2} \left(t \mid \bar{\mathcal{O}}(t^-), T > t^-, C > t^- \right), \quad (3.2)$$

where $\lambda^{A_1, A_2}(t)$ is the joint intensity process of the joint counting process $A_1(t)$ and $A_2(t)$. Similarly, let $\bar{\mathcal{O}}(t) = \{\bar{L}(t), \bar{A}_1(t), \bar{A}_2(t)\}$ denote the observed history up to t , we assume conditional exchangeability for censoring such that $\forall t \in \mathcal{T}$,

$$\lambda^C \left(t \mid \bar{\mathcal{O}}(t), T > t, C > t, T^{\bar{a}_1(t), \bar{a}_2(t)} \right) = \lambda^C \left(t \mid \bar{\mathcal{O}}(t), T > t, C > t \right), \quad (3.3)$$

where $\lambda^C(t)$ is the intensity process corresponding to the counting process of censoring. Our conditional exchangeability assumption is a continuous-time generalization of the usual sequential randomization assumption for the discrete-time marginal structural models (Robins, 1999; Howe et al., 2012).

(A3) *Positivity.* We assume that at any given time t , there is a positive probability of initiating a treatment plan, among those who are subject to initiating at least one treatment, for all configurations $\bar{\mathcal{O}}(t^-)$: $P \left\{ \lambda^{A_1, A_2}(t \mid \bar{\mathcal{O}}(t^-), T > t^-, C > t^-) > 0 \right\} = 1$. For a pair of joint treat-

ments (A_1, A_2) , at a given time t , individuals with treatment status $(0, 0)$, $(0, 1)$ or $(1, 0)$ are subject to “initiating” at least one treatment. The treatment initiation patterns can be as follows: $(0, 0) \rightarrow \{(0, 1), (1, 0), (1, 1)\}$, $(0, 1) \rightarrow \{(1, 1), (1, 0)\}$, $(1, 0) \rightarrow \{(1, 1), (0, 1)\}$. Treatment discontinuation (from 1 to 0), however, is not considered a stochastic process in our study, as COVID medication is typically prescribed with a specific treatment duration, e.g., treat with dexamethasone at a dose of 6 mg once daily for 10 days (RECOVERY Collaborative Group, 2021). Furthermore, because an individual cannot be at risk for receiving the same treatment once he or she is on the treatment, we only need to assume the positivity when the individual is *off* that specific treatment, i.e., at risk for initiating that treatment.

3.1 Framing repeated treatment initiation as recurrent events

As Figure 1 suggests, the observed treatment pattern is complex due to considerable variability in COVID treatment protocols and clinician preferences over time. Individuals may discontinue a treatment and restart the same treatment at a later time; or they may be switched altogether to another treatment. Meanwhile, patients can take more than one treatment for a period of time. Each treatment can therefore be viewed as the counting process of recurrent events, with discontinuous intervals of treatment eligibility (Andersen and Gill, 1982). Specifically, casting treatment initiation as a recurrent event process captures two distinguishing features of our observational data: (i) having received a treatment would prevent an individual from receiving the same treatment again for the time period while the individual is *on* the treatment; and (ii) after the individual was *off* the treatment, he or she would be eligible or *at risk* for re-initiating the treatment.

To formalize the treatment initiation process, we first consider a univariate treatment process N^{A_w} . We assume that the jumps of $A_w(t)$, i.e., $dA_w(t)$, is observed on certain subintervals of \mathcal{T} only. Specifically for individual i , we observe the stochastic process $A_{w,i}(t)$ on a set of in-

tervals $\mathcal{E}_{w,i} = \bigcup_{j=1}^{J_i} (V_{w,ij}, U_{w,ij}]$, where $0 \leq V_{w,i1} \leq U_{w,i1} \leq \dots \leq V_{w,iJ_i} \leq U_{w,iJ_i} \leq t_{w,iK_i}$. This representation implies the following results. First, an individual can have at most $J_i \geq 1$ initiations of treatment w : if $U_{w,iJ_i} = t_{w,iK_i}$, then individual i has $J_i - 1$ treatment initiations; and if $U_{w,iJ_i} < t_{w,iK_i}$, then individual i has J_i treatment initiations. A special case where $J_i = 1$ and $U_{w,iJ_i} = t_{w,iK_i}$ corresponds to the situation where individual i is continuously eligible for treatment initiation and has not been treated with w during the follow-up. Second, once treatment is initiated, $A_{w,i}(t)$ is no longer stochastic until person i discontinues the treatment. This also suggests that the j th treatment initiation is observed at $U_{w,ij}$. Third, $A_{w,i}(t) = 1 \forall t \in (U_{w,ij}, V_{w,i(j+1)}], j = 1, \dots, J_i - 1$. In words, treatment status is equal to one deterministically on the discontinuous intervals of ineligibility (i.e., *on* treatment period). Define a censoring or filtering process by $D_i^{A_w}(t) = I(t \in \mathcal{E}_{w,i})$, and the filtered counting process by $N_{iq}^{A_w}(t) = \int_0^t D_i^{A_w}(u) dN_{w,iq}(u)$, where q indexes the q th treatment initiation.

Following Andersen et al. (1993), we assume conditional independence among occurrences of treatment initiation given all observed history, and that the set $\mathcal{E}_{w,i}$ is defined such that $D_i^{A_w}(t)$ is predictable. The observed data with occurrences on the set $\mathcal{E}_{w,i}$ can therefore be viewed as a *marked point process* generating the filtration $(\mathcal{F}_{w,t}^D)$. Similarly, we denote the filtration generated by the counting process $\{A_w(t) : t \in \mathcal{T}\}$ corresponding to $\mathcal{E}_{w,i} = \mathcal{T}$ by $(\mathcal{F}_{w,t})$. We assume $A_{w,iq}(t)$ follows Aalen's multiplicative intensity model (Aalen et al., 2008) $\lambda_{w,iq}(t, \theta) = \alpha_{iq}(\theta, t)Y_{w,iq}(t)$, with respect to (\mathcal{F}_t) , where $\lambda_{w,iq}(t, \theta)$ is the intensity process of $A_{w,iq}(t)$, $\alpha_{iq}(\theta, t)$ is the hazard rate function parameterized by θ , and $Y_{w,iq}(t)$ is the at-risk function with $Y_{w,iq}(t) = 1$ indicating person i is eligible just before time t for the q th initiation of treatment w in the interval $[t, t + dt)$, and $Y_{w,iq}(t) = 0$ indicating otherwise. It follows that the filtered counting process $N_{iq}^{A_w}(t)$ follows the multiplicative intensity model

$$\lambda_{iq}^{A_w}(t, \theta) = \alpha_{iq}(\theta, t)Y_{iq}^{A_w}(t) \quad (3.4)$$

with respect to $(\mathcal{F}_{w,t}^D)$ (Andersen et al., 1993). Here, $Y_{iq}^{A_w}(t) = Y_{w,iq}(t)D_i^{A_w}(t)$. With two treat-

ments, model (3.4) can be directly extended for the joint treatment initiation process as

$$\lambda_{iq}^{A_1, A_2}(t, \theta) = \alpha_{iq}(\theta, t) Y_{iq}^{A_1, A_2}(t), \quad (3.5)$$

where $Y_{iq}^{A_1, A_2}(t) = Y_{(1,2),iq}(t) D_i^{A_1, A_2}(t)$ is the at-risk process for the q th treatment initiation with the filtering process defined jointly by A_1 and A_2 .

3.2 Derivation of the continuous-time weights

Under assumptions (A1)-(A3), a consistent estimator of ψ can be obtained by solving the weighted partial score equations (Hu et al., 2018),

$$\sum_{i=1}^n \int_0^\infty \Omega^{A_1, A_2}(t_{K_i}) \{Z(A_{1i}, A_{2i}, t) - \bar{Z}^*(t; \psi)\} dN_i^T(t) = 0, \quad (3.6)$$

where $\Omega^{A_1, A_2}(t_{K_i})$ is the weight that corrects for potential time-varying confounding for time-varying treatments A_1 and A_2 , $Z(A_{1i}, A_{2i}, t)_{(3 \times 1)} = [A_{1i}(t), A_{2i}(t), A_{1i}(t)A_{2i}(t)]^\top$, and

$$\bar{Z}^* = \frac{\sum_{k \in \mathcal{R}_t^T} Z(A_{k1}, A_{k2}, t) Y_k^{*T}(t) r(A_{k1}, A_{k2}, t; \psi)}{\sum_{k \in \mathcal{R}_t^T} Y_k^{*T}(t) r(A_{k1}, A_{k2}, t; \psi)} \quad (3.7)$$

is a modified version of the weighted mean of Z over observations still at risk for the outcome event at time t . In equation (3.7), we define the weighted risk set indicator for outcome $Y_i^{*T}(t) = \Omega^{A_1, A_2}(t_{K_i}) Y_i^T(t)$, where $Y_i^T(t)$ is the at-risk function for the outcome event, and $r(a_1, a_2, t) = \exp\{\psi_1 a_1(t) + \psi_2(t) a_2(t) + \psi_3 a_1(t) a_2(t)\}$.

In the discrete-time setting with non-recurrent treatment initiation, the stabilized inverse probability weights (we suppress subscript i for brevity) are given by Howe et al. (2012)

$$\Omega^{A_1, A_2}(t) = \left\{ \prod_{\{k: t_k \leq t\}} \frac{P(A_1(t_k) = a_1(t_k) | \bar{A}_1(t_{k-1}), \bar{A}_2(t_{k-1}))}{P(A_1(t_k) = a_1(t_k) | \bar{A}_1(t_{k-1}), \bar{A}_2(t_{k-1}), \bar{L}(t_{k-1}), T \geq t, C \geq t)} \right\} \times \left\{ \prod_{\{k: t_k \leq t\}} \frac{P(A_2(t_k) = a_2(t_k) | \bar{A}_1(t_k), \bar{A}_2(t_{k-1}))}{P(A_2(t_k) = a_2(t_k) | \bar{A}_1(t_k), \bar{A}_2(t_{k-1}), \bar{L}(t_{k-1}), T \geq t, C \geq t)} \right\}, \quad (3.8)$$

where t_k 's are a set of ordered discrete time points common to all individuals satisfying $0 = t_0 < t_1 < t_2 < \dots \leq t$. While $\Omega^{A_1, A_2}(t)$ in (3.8) corrects for time-varying confounding by adjusting

for $\bar{L}(t)$ in the weights, it requires that the time points are well aligned across all individuals. In addition, it does not accommodate the recurrent nature of complex intervention strategies as in our observational study.

We now generalize the weights developed for the discrete-time setting to a continuous-time process, which do not require the time points to be well aligned. Partition the time interval $[0, t]$ into a number of small time intervals, and let $dA_w(s)$ be the increment of A_w over the small time interval $[s, s + ds)$, $\forall s \in [0, t]$. Recall that treatment initiation, or the jumps of $A_w(t)$, $dA_w(t)$, is observed on a number of subintervals of \mathcal{T} only. That is, conditional on history $\bar{L}(s)$, the occurrence of treatment initiation for an individual in $[s, s + ds)I(s \in \mathcal{E})$ is a Bernoulli trial with outcomes $dA_w(s) = 1$ and $dA_w(s) = 0$. Then the $P(A_w(t_k) = a_w(t_k) | \bullet)$ in equation (3.8) can be represented by

$$D^{A_w}(s) \{P(dA_w(s) = 1 | \bullet)\}^{dA_w(s)} \{P(dA_w(s) = 0 | \bullet)\}^{1-dA_w(s)},$$

which takes the form of the individual partial likelihood for the filtered counting process $\{D^{A_w}(s)A_w(s) : 0 \leq s \leq t\}$. When the number of time intervals in $[0, t]$ increases and ds approaches zero, the finite product over the number of time intervals of the individual partial likelihood will approach a product integral (Aalen et al., 2008), given by

$$\begin{aligned} & \prod_0^t \{D^{A_w}(s)\lambda^{A_w}(s | \bullet)ds\}^{dA_w(s)} \{D^{A_w}(s)(1 - \lambda^{A_w}(s | \bullet)ds)\}^{1-dA_w(s)} \\ &= \left[\prod_0^t \{D^{A_w}(s)\lambda^{A_w}(s | \bullet)\}^{\Delta A_w(s)} \right] \exp \left\{ - \int_0^t D^{A_w}(s)\lambda^{A_w}(s | \bullet)ds, \right\} \end{aligned} \quad (3.9)$$

where $\Delta A_w(t) = A_w(t) - A_w(t^-)$. For individual i , both factors in (3.9) need to be evaluated with respect to the individual's filtered counting process $\{N_{iq}^{A_w}(t) : 0 \leq t \leq t_{K_i}, q = 1, \dots, Q_{w,i}\}$, with the first quantity being equal to the finite product over the jump times and the second quantity being the survival function for treatment initiation. As described in Section 3.1, the number of treatment initiations for individual i , $Q_{w,i}$ can take three values: (i) $Q_{w,i} = 0$, (ii) $Q_{w,i} = J_i - 1$

or (iii) $Q_{w,i} = J_i$. Corresponding to the three cases, the quantity in (3.9) can be rewritten as

$$\text{Quantity (3.9)} = \begin{cases} S^{A_w}(t_{K_i} | \bullet) & \text{if } Q_{w,i} = 0 \\ f^{A_w}(U_{i,J_i-1} | \bullet) \{S^{A_w}(V_{iJ_i} | \bullet) - S^{A_w}(t_{iK_i} | \bullet)\} & \text{if } Q_{w,i} = J_i - 1 \\ f^{A_w}(U_{iJ_i} | \bullet) & \text{if } Q_{w,i} = J_i, \end{cases}$$

where S^{A_w} and f^{A_w} are the survival and density function of the filtered counting process for treatment A_w . Here we assume that initiations of different treatments are ordered. For example, whether to initiate A_2 at time t is decided upon observing the treatment status $A_1(t)$. This suggests that the hazard function λ^{A_2} is estimable by conditioning on $\bar{A}_1(t)$ and $\bar{A}_2(t^-)$; and the hazard function λ^{A_1} is estimable by conditioning on $\bar{A}_1(t^-)$ and $\bar{A}_2(t^-)$. For exposition brevity, let $\bar{\mathcal{O}}_1(t) = \{\bar{A}_1(t^-), \bar{A}_2(t^-), \bar{L}(t^-), T \geq t, C \geq t\}$, $\bar{\mathcal{O}}_2(t) = \{\bar{A}_1(t), \bar{A}_2(t^-), \bar{L}(t^-), T \geq t, C \geq t\}$, $\bar{\mathcal{O}}^{A_1}(t) = \{\bar{A}_1(t^-), \bar{A}_2(t^-), T \geq t, C \geq t\}$ and $\bar{\mathcal{O}}^{A_2}(t) = \{\bar{A}_1(t), \bar{A}_2(t^-), T \geq t, C \geq t\}$. Putting this all together, the individual continuous-time stabilized inverse probability weight that corrects for time-varying confounding by \bar{L} is given by $\Omega^{A_1, A_2}(t) = \Omega^{A_1}(t)\Omega^{A_2}(t)$ with Ω^{A_w} being:

$$\begin{aligned} & \Omega^{A_w}(t_{K_i}) \\ &= \begin{cases} \frac{S^{A_w}(t_{K_i} | \bar{\mathcal{O}}^{A_w}(t_{K_i}))}{S^{A_w}(t_{K_i} | \bar{\mathcal{O}}_w(t_{K_i}))} & \text{if } Q_{w,i} = 0 \\ \frac{f^{A_w}(U_{i,J_i-1} | \bar{\mathcal{O}}^{A_w}(U_{i,J_i-1})) \{S^{A_w}(V_{iJ_i} | \bar{\mathcal{O}}^{A_w}(V_{iJ_i})) - S^{A_w}(t_{iK_i} | \bar{\mathcal{O}}^{A_w}(t_{K_i}))\}}{f^{A_w}(U_{i,J_i-1} | \bar{\mathcal{O}}_w(U_{i,J_i-1})) \{S^{A_w}(V_{iJ_i} | \bar{\mathcal{O}}_w(V_{iJ_i})) - S^{A_w}(t_{iK_i} | \bar{\mathcal{O}}_w(t_{K_i}))\}} & \text{if } Q_{w,i} = J_i - 1 \\ \frac{f^{A_w}(U_{iJ_i} | \bar{\mathcal{O}}^{A_w}(U_{iJ_i}))}{f^{A_w}(U_{iJ_i} | \bar{\mathcal{O}}_w(U_{iJ_i}))} & \text{if } Q_{w,i} = J_i \end{cases} \end{aligned} \quad (3.10)$$

Turning to censoring, under the conditional exchangeability assumption (A2), the censoring process is covariate- and treatment-dependent. To correct for selection bias due to censoring, we additionally define a weight function associated with censoring,

$$\Omega^C(G_i) = \frac{S^C(G_i | C_i \geq G_i, T_i \geq G_i)}{S^C(G_i | \bar{A}_1(G_i), \bar{A}_2(G_i), \bar{L}(G_i), C_i \geq G_i, T_i \geq G_i)},$$

where S^C is the survival function associated with the censoring process, and

$$G_i = \mathbb{1}(\Delta_i^T = 1)T_i + \mathbb{1}(\Delta_i^T = 0, C_i > t_{K_i})t_{K_i} + \mathbb{1}(\Delta_i^T = 0, C_i \leq t_{K_i})C_i.$$

This leads to a final modification of the estimating equation for ψ ,

$$\sum_{i=1}^n \int_0^\infty \Omega^{A_1, A_2} \Omega^C(G_i) \{Z(A_{1i}, A_{2i}, t) - \bar{Z}^{**}(t; \psi)\} dN_i^T(t) = 0, \quad (3.11)$$

where $\bar{Z}^{**} = \frac{\sum_{k \in \mathcal{R}_t^T} Z(A_{k1}, A_{k2}, t) Y_k^{**T}(t) r(A_{k1}, A_{k2}, t; \psi)}{\sum_{k \in \mathcal{R}_t^T} Y_k^{**T}(t) r(A_{k1}, A_{k2}, t; \psi)}$ and $Y_i^{**T}(t) = \Omega^C(G_i) \Omega^{A_1, A_2}(t_{K_i}) Y_i^T(t)$.

3.3 Estimation of the causal survival effects

We consider four ways in which the continuous-time weights $\Omega^{A_1, A_2}(t)$ can be estimated: (i) fitting a usual Cox regression model for the intensity process of the counting process of treatment initiation $\{A_w(t) : t \in \mathcal{T}\}$, estimating the density function f^{A_w} and survival function S^{A_w} from the fitted model with the Nelson-Aalen estimator for the baseline intensity function, and finally calculating the weights following equation (3.10); (ii) smoothing the Nelson-Aalen estimator and in turn f^{A_w} and S^{A_w} estimated from the fitted Cox regression model by means of kernel functions (Ramlau-Hansen, 1983), and calculating the weights using the smoothed version of f^{A_w} and S^{A_w} ; (iii) fitting a multiplicative intensity tree-based model (Yao et al., 2020) in which the functional form of the intensity ratio for treatment initiation is flexibly captured to estimate the f^{A_w} and S^{A_w} and calculate the weights; (iv) smoothing the Nelson-Aalen estimator of the baseline intensity from the tree-based model and calculating the weights using the smoothed version of f^{A_w} and S^{A_w} . Among these approaches, (i) relies on the parametric assumptions about the intensity ratio relationships between the treatment initiation process and covariate process and may be subject to model misspecification and bias for estimating causal effects. Compared to the Nelson-Aalen estimator which includes discrete jumps at event occurrences, the kernel function estimator in (ii) may help alleviate the issue of extreme or spiky weights, and has also been shown to be a consistent and asymptotically normal baseline intensity estimator (Andersen et al., 1993). Approach (iii) leverages a recent random survival forests model (Yao et al., 2020) that can accommodate time-varying covariates to mitigate the parametric assumptions and attendant biases associated with

the usual Cox regression. With baseline time-fixed treatment, prior work has used similar machine learning techniques to improve propensity score weighting estimators (Lee et al., 2010) as well as to provide more accurate causal effect estimates with censored survival data (Hu et al., 2021). Finally, approach (iv) smooths the baseline intensity estimated from the survival forests for estimating the stabilized inverse probability weights, and serves as an additional step to smooth over the potentially spiky weights. In Section 4, we compare the performances of these four strategies to estimating the continuous-time weights to generate practical recommendations. In addition, the censoring weight function $\Omega^C(G_i)$ can be estimated in a similar fashion via any one of these four approaches. Additional details of kernel function smoothing in approach (ii) and random survival forests in approach (iii) are presented in Web Appendix S1.

To accommodate the time-varying covariate process and account for the recurrent nature of treatment initiation, we fit a survival model to the counting process style of data input. Each individual is represented by several rows of data corresponding to nonoverlapping time intervals of the form (start, stop]. To allow for discontinuous intervals of eligibility, when defining multiple time intervals $\mathcal{E}_{w,i} = \cup_{j=1}^{J_i} (V_{w,ij}, U_{w,ij}]$ on \mathcal{T} for individual i , the duration of a treatment is removed from \mathcal{T} when the individual is currently being treated and therefore no longer eligible for initiating the treatment. Finally, since our estimators for ψ is a solution to the weighted partial score equation (3.11), we can use the robust sandwich variance estimator to construct confidence intervals for the structural parameters. In practice, the robust sandwich variance estimator is at most conservative under the discrete-time setting (Shu et al., 2021), and we will empirically assess the accuracy of this variance estimator with continuous-time weights via simulations.

3.4 Extensions to more than two time-varying treatments

Although we introduce our methods with two longitudinal treatments, our approach can be extended to more than two time-varying treatments in a straightforward fashion. In theory, a

fully interacted version of model (2.1) can be formed to include all the main effects of $a_w(t)$ $\forall w \in \mathcal{W}$ and the interactions thereof. Clinical interests and data sparsity on combinations of treatments may also be used to guide the inclusion of interaction terms into the structural model. Suppose $\mathcal{B} = \{b_1(t), \dots, b_V(t)\}$ is a collection of causal interaction effects of interest, e.g., $b_1(t) = a_1(t)a_2(t)$, the general joint marginal structural proportional hazards model is

$$\lambda^{T^{\bar{a}_1(t), \dots, \bar{a}_W(t)}}(t) = \lambda_0(t) \exp \left\{ \sum_{w=1}^W \psi_{1w} a_w(t) + \sum_{v=1}^V \psi_{2v} b_v(t) \right\}, \quad (3.12)$$

where ψ_{1w} 's and ψ_{2v} 's capture the causal main and interaction effects on the counterfactual hazard function. A consistent estimator of $\boldsymbol{\psi} = \{\psi_{11}, \dots, \psi_{iW}, \psi_{21}, \dots, \psi_{2V}\}$ can be obtained by solving the general form of the estimating equation

$$\sum_{i=1}^n \int_0^\infty \Omega^{A_1, \dots, A_W} \Omega^C(G_i) \{Z(A_{1i}, \dots, A_{Wi}, t) - \bar{Z}^{**}(t; \boldsymbol{\psi})\} dN_i^T(t) = 0, \quad (3.13)$$

where $Z(A_{1i}, \dots, A_{Wi}, t)$ is a vector of length $W + J$ representing the time-varying treatment status $A_w(t)$ and multiplicative terms of the treatment status $A_v(t)A_{v'}(t)$, \bar{Z}^{**} is evaluated using weighted risk set indicators $Y_i^{**T}(t) = \Omega^C(G_i) \Omega^{A_1, \dots, A_W}(t_{K_i}) Y_i^T(t)$. The joint treatment weights $\Omega^{A_1, \dots, A_W}(t_{K_i})$ can be estimated by assuming a specific order in which treatments are initiated and calculating the weights using appropriate history information $\bar{O}_w(t)$ and $\bar{O}^{A_w}(t)$, similar as described in Section 3.2. The estimation of the censoring weights $\Omega^C(G_i)$ also follows the same strategy outlined in Section 3.2 with two longitudinal treatments.

4. SIMULATION STUDY

4.1 Simulation design

We carry out simulations to investigate the finite-sample properties of the proposed weight estimators for the marginal structure Cox model parameters. We simulate data compatible with the marginal structural Cox model by generating and relating data adhering to the structural nested accelerated failure time (SNAFT) model (Young et al., 2008). A general representation of

a SNAFT model for time-varying treatment a is (Hernán et al., 2005)

$$T^{\bar{0}} = \int_0^{T^{\bar{a}}} \exp[\psi_{\text{aft}} a(t)],$$

where $T^{\bar{0}}$ is the counterfactual failure time under no treatment. Robins (1992) developed a simulation algorithm to generate data adhering to the SNAFT model under the discrete-time version of the identifying assumptions (A1)-(A3) in Section 3. Young et al. (2008) showed that, under the same identifying assumptions, data adhering to a marginal structural Cox model of the form

$$\lambda^{T^{\bar{a}}}(t) = \lambda_0(t) \exp[\psi_{\text{msm}} a(t)]$$

can be simulated from a SNAFT model with $\psi_{\text{aft}} = \psi_{\text{msm}}$ by adding an additional quantity to the term $\exp[\psi_{\text{aft}} a(t)]$. In particular when $T^{\bar{0}}$ has an exponential distribution, the additional quantity is zero, hence the structural nested AFT model simulation algorithm (Robins, 1992) can be used to appropriately simulate data compatible with the marginal structural cox model under complex time-varying data structures. Building on these previous works, we extend the simulation algorithm described in Karim et al. (2017) to generate data from the joint marginal structural Cox model, while allowing for multiple time-varying treatments with discontinuous intervals of treatment eligibility and for both continuous and discrete time-varying confounders.

Throughout we simulate an observational study with $n = 1000$ patients and two time-varying treatments $A_1(t)$ and $A_2(t)$. We assume $\bar{L}(t)$ is appropriately summarized by a continuous time-varying confounding variable $L_1(t)$ and a binary time-varying confounding variable $L_2(t)$. The simulation algorithm includes two steps. In step (1), we consider nonlinear terms for the continuous variable $L_1(t_k)$ and the interaction term $A_1(t_{k-1}) \times L_1(t_k)$, $A_2(t_k) \times L_1(t_k)$, $A_1(t_{k-1}) \times L_2(t_k)$ and $A_2(t_k) \times L_2(t_k)$ in the true treatment decision model. In particular, past treatment status $\{A_1(t_{k-1}), A_2(t_k)\}$ is a predictor of $L(t_k)$, which then predicts future treatment exposure $\{A_1(t_k), A_2(t_{k+1})\}$ as well as future failure status $Y(t_{k+1})$ via $1/\log(T^{\bar{0}})$ (Karim et al., 2017). Therefore, $L(t_k)$ is a time-dependent confounder affecting both the future treatment choices and

counterfactual survival outcomes. The simulation of treatment initiation is placed in the recurrent events framework. Once treatment is initiated at time t_k , treatment duration following initiation is simulated from a zero-truncated Poisson distribution. In step (1), we generate a longitudinal data set with 100×1000 observations (100 aligned measurement time points for each of $n = 1000$ individuals). In step (2), we randomly discard a proportion of follow-up observations for a randomly selected subset of individuals (Lin et al., 2004); and in the resulting data set, the individuals will have varying number of follow-up measurement time points, which are also irregularly spaced. Web Appendix S2 provides the full pseudo-code for simulating data under the marginal structural Cox model with two time-varying treatments.

Our simulation parameters are chosen so that the simulated data possess similar characteristics to those observed in the COVID-19 data set. The treatments A_1 and A_2 are simulated to resemble dexamethasone and remdesivir such that: (i) about 20% patients did not take any of the anti-viral and anti-inflammatory medications aimed at treating COVID-19; (ii) among those who were treated, 62% took dexamethasone only, 25% took remdesivir only and 13% took both (either concurrently or with treatment switching); (iii) the number of initiations for both treatments ranges from 0 to 4 with the average medication duration about 5 days. The values of treatment effect parameters ψ_1 and ψ_2 were set to yield a 6.7% mortality proportion among those who received dexamethasone and a 4.9% mortality proportion in those treated with remdesivir.

4.2 Comparison of methods

We conduct two sets of simulations to investigate the finite-sample performance of our proposed joint marginal structural survival model in continuous time (JMSSM-CT). First, we compare how accurate the four weight estimators described in Section 3.3 estimate the structural parameter ψ . Second, we use the best weight estimator, suggested by the first set of simulation, for JMSSM-CT, and compare it with the joint marginal structural model that requires aligned discrete time

points (JMSM-DT). To ensure an objective comparison, we use the random forests (Yao et al., 2020) and adapt it into our proposed recurrent events framework to estimate the weights for JMSM-DT. In addition, we implement both JMSSM-CT and JMSM-DT on the “rectangular” simulation data with 100 aligned time points for each individual and on the “ragged” data with irregular observational time points. The performance on the rectangular data will be considered as the benchmark performance, based on which we will assess the relative accuracy of JMSSM-CT and JMSM-DT when estimating the structural parameters with the “ragged” data.

4.3 *Performance characteristics*

To assess the performance of each method, we simulate 250 observational data sets using the above approach, and evaluate the absolute bias, root mean squared error (RMSE) and coverage probability (CP) for estimating the ψ . The CP is evaluated on normality-based confidence intervals with the robust sandwich variance estimator. Figure 2 suggests that the weight estimator (iv) using both the flexible tree-based survival model and kernel function estimator of the treatment initiation intensity yielded the lowest biases in estimating both ψ_1 and ψ_2 . By contrast, the weight estimator (i) via the usual main-effects Cox regression model along with the Nelson-Aalen baseline intensity estimator produced the largest estimation bias. Applying the kernel function smoothing to the Nelson-Aalen estimator led to bias reduction for both the Cox (approach (ii)) and tree-based survival model (approach (iv)) for the treatment process. Flexible modeling of the intensity ratio function has a larger effect in reducing the bias in structural parameter estimates than smoothing the nonparametric baseline intensity estimator. For example, compared to approach (ii), approach (iii) further reduced the mean absolute bias (MAB) in estimating $\hat{\psi}_1$ by approximately 67%. Supplementary Table 1 summarizes the MAB, RMSE and CP for the four weight estimators and similarly suggests that approach (iv) led to the smallest MAB and RSME, and provided close to nominal CP with the robust sandwich variance estimator.

The second set of simulation benchmarks the performance of JMSSM-CT versus JMSM-DT on the data with fully aligned follow-up time points and compare how much each method can recover the benchmark performance in situations where the longitudinal measurements are irregularly spaced. Table 2 displays the MAB, RMSE and CP for each of the two methods under both data settings, and Supplementary Figure 1 visualizes the distributions of biases across 250 data replications. In the rectangular data setting with fully aligned time points, compared to JMSM-DT, JMSSM-CT had similar CP but smaller MAB and RMSE. As the sparsity of longitudinal measurements increased and the time intervals became unevenly spaced, the proposed JMSSM-CT could still recover the benchmark performance; whereas the JMSM-DT had a deteriorating performance (larger MAB and RMSE and lower CP), with larger performance decline under coarser discretization of the follow-up time. Supplementary Table 2 summarizes the distribution of estimated individual time-varying weights from one random replication of the ragged data for JMSSM-CT and JMSM-DT. Overall, JMSSM-CT with the weight estimator (iv) provided the smallest maximum/minimum weight ratio (2.36/0.68) and no extreme or spiky weights.

5. ESTIMATING CAUSAL EFFECTS OF MULTIPLE COVID-19 TREATMENTS

We apply the proposed JMSSM-CT with the best weight estimation approach (iv) (Section 3.3) to analyze a comprehensive COVID-19 data set drawn from the Epic electronic medical records system at the Mount Sinai Medical Center, with the goal of studying the comparative effectiveness of multiple COVID-19 treatment strategies. The data set includes 11,286 de-identified unique adult patients (≥ 18 years of age) who were diagnosed with COVID-19 and hospitalized within the Mount Sinai Health System between February 25, 2020 to February 26, 2021. A confirmed case of COVID-19 was defined as a positive test result from a real-time reverse-transcriptase PCR-based clinical test carried out on nasopharyngeal swab specimens collected from the patient (Wang et al., 2020). We focus on four treatment classes that are of most clinical interest: (i)

remdesivir; (ii) dexamethasone; (iii) anti-inflammatory medications other than corticosteroids; and (iv) corticosteroids other than dexamethasone. Detailed definitions of the four treatment classes are provided in Supplementary Table 3. The observed treatment patterns are visualized in Figure 1; patients could be simultaneously prescribed two or more treatment strategies, or they could switch from one treatment class to another during their hospital stays.

We considered a composite outcome, ICU admission or in-hospital death, whichever occurs first. The outcome may be right censored by hospital discharge or administratively censored on t^o = February 26, 2021, the date on which the database for the current analysis was locked. For modeling purposes, we assumed that $\bar{L}(t)$ was appropriately summarized by age, sex, race, ethnicity, D-dimer levels (the degradation product of crosslinked fibrin, reflecting ongoing activation of the hemostatic system), serum creatinine levels (a waste product that forms when creatine breaks down, indicating how well kidneys are working), whether the patient used tobacco at the time of admission, whether the patient was admitted to ICU, history of comorbidities represented by a set of binary variables: hypertension, coronary artery disease, cancer, diabetes, asthma and chronic obstructive pulmonary disease, hospital site, and patient oxygen levels (definition provided in Supplementary Table 4). The time-varying confounding variables were ICU admission, D-dimer levels, serum creatinine level and patient oxygen levels. When fitting the joint marginal structural proportional hazards model (3.12), pairwise treatment interactions were included if there were sufficient data points supporting the joint use of the pair of treatments.

Using the stabilized inverse probability weights to correct for time-varying confounding and censoring, the structural model parameter estimates $\hat{\psi}$ (log hazard ratio) and the associated 95% confidence intervals are provided in Table 2. Using the parameter estimates, we further computed the counterfactual survival curves under each treatment regimen. Figure 3 presents the counterfactual survival curves for ICU admission or death, whichever occurs first, among patients with COVID-19 infection, under five treatment regimens given upon admission to hospital. Among

the four separate treatment classes, remdesivir had significantly better treatment benefits followed by dexamethasone than two alternative treatment classes: anti-inflammatory medications other than corticosteroid and corticosteroids other than dexamethasone. Interestingly, remdesivir and corticosteroids other than dexamethasone had a significant treatment interaction effect suggesting additional survival benefit when they are used in combination. This is demonstrated by the highest counterfactual survival curve under the concomitant use of these two types of medications.

6. DISCUSSION

Motivated by inconclusive real-world evidence for the comparative effectiveness of multiple treatment strategies for COVID-19, we have developed a joint marginal structural survival model and novel weighting schemes to address time-varying confounding and censoring in continuous time. There are three main advantages of our proposed method. First, this approach casts the complex time-varying treatment with irregular “start/stop” switches into the process of recurrent events where treatment initiation can be considered under the recurrent event framework with discontinuous intervals of eligibility. This innovative formulation enables us to address complex time-varying confounding by modeling the intensity processes of the filtered counting processes for complex time-varying treatments. Second, the proposed method is able to handle a complex longitudinal dataset on its own terms, without discretizing and artificially aligning measurement times, which would lead to less accurate and efficient treatment effect estimates, as demonstrated by our simulations. Third, modern machine learning techniques designed for censored survival data and smoothing techniques of the baseline intensity function can be used easily for with our weighting method to further improve the treatment effect estimator under conventional parametric formulations. We have also introduced a simulation algorithm that is compatible with the complex data structures of our proposed modeling framework, and demonstrated the accuracy of the proposed method for estimating causal parameters.

Our approach can be extended in the following two directions. First, we considered a joint marginal structural proportional hazards model and a tailored simulation algorithm to generate datasets of complex time-varying structures that are compatible with the proportional hazards model. It may be worthwhile to develop alternative joint marginal structural survival models such as the structural additive hazards model, and assess the robustness of different structural models for estimating counterfactual survival functions under different data generating processes. Second, we have maintained the conditional exchangeability assumption in our work, and therefore developing sensitivity analyses to capture the effects of time-varying unmeasured confounding for our model would be a worthwhile and important contribution.

7. SOFTWARE

R codes to implement the proposed methods and replicate our simulation studies are provided in the GitHub page of the first author <https://github.com/liangyuanhu/JMSSM-CT>.

8. SUPPLEMENTARY MATERIAL

Web Appendices, Tables and Figures referenced in Section 3, 4, 5, are available online at <http://biostatistics.oxfordjournals.org>.

ACKNOWLEDGMENTS

This work was supported in part by award ME_2017C3_9041 from the Patient-Centered Outcomes Research Institute and by grants R21 CA245855-01 and P30CA196521-01 from the National Cancer Institute. *Conflict of Interest:* None declared.

REFERENCES

- Aalen, O., Borgan, O., and Gjessing, H. (2008). *Survival and event history analysis: a process point of view*. Springer Science & Business Media, New York.
- Andersen, P. K., Borgan, O., Gill, R. D., and Keiding, N. (1993). *Statistical models based on counting processes*. New York: Springer-Verlag.
- Andersen, P. K. and Gill, R. D. (1982). Cox’s regression model for counting processes: a large sample study. *The Annals of Statistics* **10**, 1100–1120.
- Hernán, M. A., Cole, S. R., Margolick, J., Cohen, M., and Robins, J. M. (2005). Structural accelerated failure time models for survival analysis in studies with time-varying treatments. *Pharmacoepidemiology and Drug Safety* **14**, 477–491.
- Howe, C. J., Cole, S. R., Mehta, S. H., and Kirk, G. D. (2012). Estimating the effects of multiple time-varying exposures using joint marginal structural models: Alcohol consumption, injection drug use, and HIV acquisition. *Epidemiology* **23**, 574.
- Hu, L. and Hogan, J. W. (2019). Causal comparative effectiveness analysis of dynamic continuous-time treatment initiation rules with sparsely measured outcomes and death. *Biometrics* **75**, 695–707.
- Hu, L., Hogan, J. W., Mwangi, A. W., and Siika, A. (2018). Modeling the causal effect of treatment initiation time on survival: Application to HIV/TB co-infection. *Biometrics* **74**, 703–713.
- Hu, L., Ji, J., and Li, F. (2021). Estimating heterogeneous survival treatment effect in observational data using machine learning. *Statistics in Medicine* **40**, 4691–4713.
- Johnson, B. A. and Tsiatis, A. A. (2005). Semiparametric inference in observational duration-response studies, with duration possibly right-censored. *Biometrika* **92**, 605–618.

- Karim, M. E., Platt, R. W., and Group, B. S. (2017). Estimating inverse probability weights using super learner when weight-model specification is unknown in a marginal structural cox model context. *Statistics in Medicine* **36**, 2032–2047.
- Lee, B. K., Lessler, J., and Stuart, E. A. (2010). Improving propensity score weighting using machine learning. *Statistics in Medicine* **29**, 337–346.
- Lin, H., Scharfstein, D. O., and Rosenheck, R. A. (2004). Analysis of longitudinal data with irregular, outcome-dependent follow-up. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **66**, 791–813.
- Ramlau-Hansen, H. (1983). Smoothing counting process intensities by means of kernel functions. *The Annals of Statistics* **11**, 453–466.
- RECOVERY Collaborative Group (2021). Dexamethasone in hospitalized patients with Covid-19. *New England Journal of Medicine* **384**, 693–704.
- Robins, J. (1992). Estimation of the time-dependent accelerated failure time model in the presence of confounding factors. *Biometrika* **79**, 321–334.
- Robins, J., Orellana, L., and Rotnitzky, A. (2008). Estimation and extrapolation of optimal treatment and testing strategies. *Statistics in Medicine* **27**, 4678–4721.
- Robins, J. M. (1999). Association, causation, and marginal structural models. *Synthese* **121**, 151–179.
- Ryalen, P. C., Stensrud, M. J., Fosså, S., and Røysland, K. (2020). Causal inference in continuous time: an example on prostate cancer therapy. *Biostatistics* **21**, 172–185.
- Saarela, O. and Liu, Z. (2016). A flexible parametric approach for estimating continuous-time inverse probability of treatment and censoring weights. *Statistics in Medicine* **35**, 4238–4251.

- Shu, D., Young, J. G., Toh, S., and Wang, R. (2021). Variance estimation in inverse probability weighted cox models. *Biometrics* **77**, 1101–1117.
- Wang, Z., Zheutlin, A., Kao, Y.-H., Ayers, K., Gross, S., Kovatch, P., Nirenberg, S., Charney, A., Nadkarni, G., De Freitas, J. K., et al. (2020). Hospitalised COVID-19 patients of the Mount Sinai Health System: a retrospective observational study using the electronic medical records. *BMJ open* **10**, e040441.
- Yao, W., Frydman, H., Larocque, D., and Simonoff, J. S. (2020). Ensemble methods for survival data with time-varying covariates. *arXiv preprint arXiv:2006.00567*.
- Young, J. G., Hernán, M. A., Picciotto, S., and Robins, J. M. (2008). Simulation from structural survival models under complex time-varying data structures. *JSM Proceedings, Section on Statistics in Epidemiology, Denver, CO: American Statistical Association* pages 1–6.
- Yousefi, B., Valizadeh, S., Ghaffari, H., Vahedi, A., Karbalaei, M., and Eslami, M. (2020). A global treatments for coronaviruses including COVID-19. *Journal of Cellular Physiology* **235**, 9133–9142.

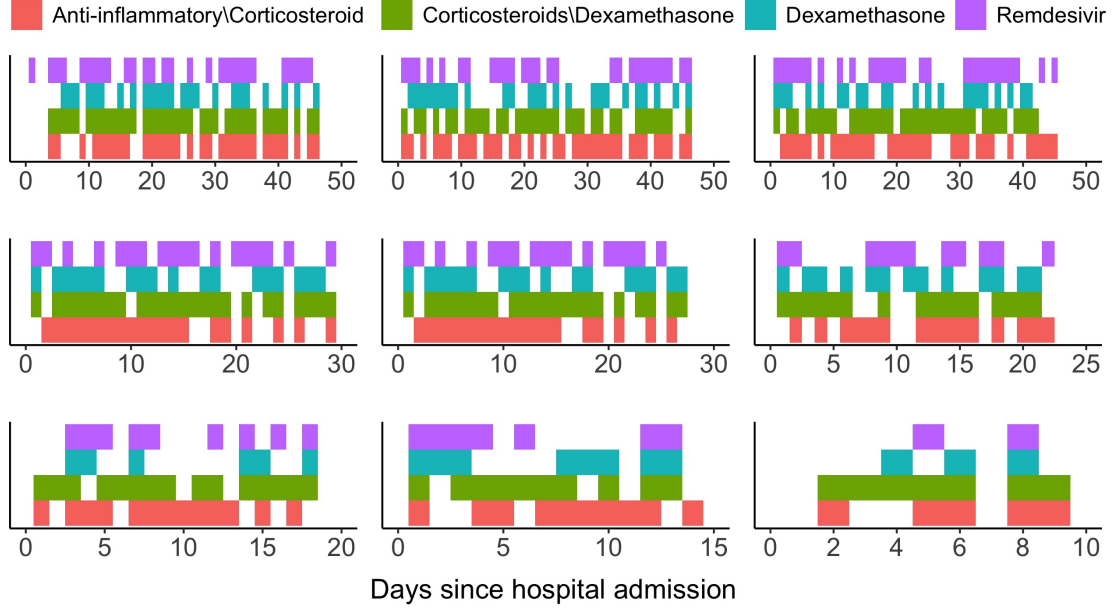


Fig. 1. Treatment processes for nine randomly selected patients visualized by heat maps. Colors indicate remaining on treatment. Lack of color corresponds to being switched off treatment.

Table 1. Comparing the proposed method JMSSM-CT in continuous time with JMSM-DT in discrete time in estimating the treatment effect ψ on the bases of mean absolute bias (MAB), root mean square error (RMSE) and coverage probability (CP) across 250 data replications. In the estimation of the weights, the weight estimator (iv) was used for JMSSM-CT and the random forests adapted into our recurrent events framework (Section 3.2) was used for JMSM-DT. Both methods were implemented on the “rectangular” simulation data with 100 aligned time points for each individual and on the “ragged” data with unaligned time points. With the ragged data, the follow-up time was discretized in the space of 0.5 , 1 and 2 days for JMSM-DT.

Data format	Methods	ψ_1			ψ_2		
		MAB	RMSE	CP	MAB	RMSE	CP
Rectangular	JMSM-DT	.021	.026	.944	.019	.023	.948
	JMSSM-CT	.015	.020	.948	.014	.018	.948
	JMSM-DT (2d)	.040	.047	.660	.035	.041	.668
Ragged	JMSM-DT (1d)	.033	.041	.732	.029	.035	.738
	JMSM-DT (0.5d)	.027	.034	.801	.024	.030	.804
	JMSSM-CT	.016	.022	.952	.015	.019	.952

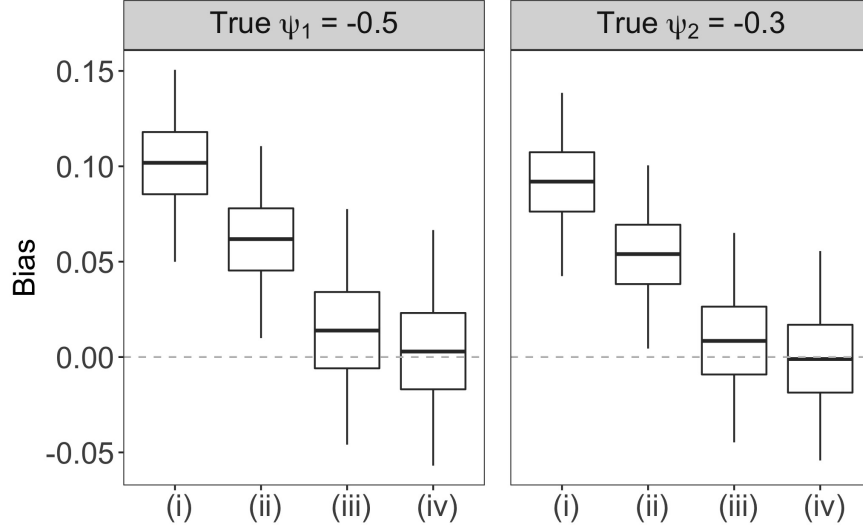


Fig. 2. Biases in the estimates of ψ_1 and ψ_2 among 250 data replications using four approaches to estimate the weights as described in Section 3.4. Approach (i) uses main-effects Cox regression model and Nelson-Aalen estimator for baseline intensity. Approach (ii) uses kernel function smoothing of the Nelson-Aalen estimator in approach (i). Approach (iii) uses a survival forests model that accommodates time-varying covariates and Nelson-Aalen estimator for baseline intensity. Approach(iv) uses kernel function smoothing of the Nelson-Aalen estimator in approach (iii).

Table 2. The joint and interactive effect estimates $\hat{\psi}$ (log hazard ratio) of COVID-19 treatments and associated 95% confidence intervals (CI), using the COVID-19 dataset drawn from the Epic electronic medical records system at the Mount Sinai Medical Center. The composite outcome of in-hospital death or admission to ICU was used for the general COVID-19 patients. Time to in-hospital death was used for subpopulations of those who had never been admitted into ICU (pre-ICU) and of post-ICU patients. Confidence intervals were estimated via the robust sandwich variance estimators. “ \times ” denotes treatment interaction and “—” indicates that the pairwise interaction was not included in the joint marginal structural survival model.

Treatment classes	$\hat{\psi}$ (95% Confidence Interval)		
	Overall	Pre-ICU	Post-ICU
Dexamethasone	-0.20(-0.35, -0.06)	-0.06(-0.71, 0.64)	-0.75(-1.42, -0.08)
Remdesivir	-0.53(-0.75, -0.31)	-0.62(-1.22, -0.02)	-0.54(-1.04, -0.04)
Corticosteroids other than dexamethasone	-0.08(-0.29, 0.19)	-0.13(-0.46, 0.21)	-0.19(-0.27, -0.03)
Anti-inflammatory medications other than corticosteroids	-0.05(-0.56, 0.47)	-0.28(-1.02, 0.45)	-0.08(-0.89, 0.72)
Remdesivir \times Corticosteroids other than dexamethasone	-0.74(-0.95, -0.52)	—	-0.71(-1.38, -0.04)
Dexamethasone \times Corticosteroids other than dexamethasone	—	—	-1.13(-1.78, -0.46)




Fig. 3. Counterfactual survival curves for each of five treatment strategies among the general COVID-19 patients. The composite outcome of ICU admission or death is used.

Web-based Supplementary Materials for “Estimating the causal effects of multiple intermittent treatments with application to COVID-19” by Hu et al.

email: liangyuan.hu@rutgers.edu

S1 Additional technical details

S1.1 The kernel function estimator

In Section 3.4, four approaches are described to estimate the stabilized time-varying inverse probability of treatment weights. The Nelson-Aalen estimator of the baseline intensity is smoothed by means of kernel functions in approach (ii) and approach (iv). For exposition brevity, consider a simple multiplicative intensity model for treatment initiation

$$\lambda(t) = \lambda_0(t) r(\bar{L}(t), \beta) Y(t),$$

where $\lambda_0(t)$ is the baseline intensity rate function, $r(\bar{L}(t), \beta) > 0$ is the (time-dependent) intensity ratio function parameterized by β , and $Y(t)$ is the at risk indicator. The cumulative baseline hazard is $\Lambda_0(t) = \int_0^t \lambda_0(s) ds$. The kernel function estimator for $\lambda_0(t)$ is derived by smoothing the increments of the Nelson-Aalen estimator of $\hat{\Lambda}_0$ as

$$\hat{\lambda}_0(t) = b^{-1} \int_{\mathcal{T}} K\left(\frac{t-s}{b}\right) d\hat{\Lambda}_0(s),$$

where K is the kernel function, which is a bounded function on $[-1, 1]$ and has integral 1, and the bandwidth b is a positive parameter governing the amount of smoothness. Commonly used kernel functions include the standard normal density function $K(x) = \phi(x)$ and the Epanechnikov kernel function $K(x) = 0.75(1 - x^2)$, $|x| \leq 1$. ? shows that the kernel function estimator $\hat{\lambda}_0(t)$ is consistent and asymptotically normal provided that there exists a sequence of positive constants $\{a_n\}$, increasing to infinity as $n \rightarrow \infty$, and that the bandwidth tends to zero more slowly than a_n^{-2} as $n \rightarrow \infty$.

Placed in the framework of recurrent events for a treatment A_w that can have multiple “start/stop” switches as described in Section 3.1, the intensity of the q th treatment initiation $\lambda_q^{A_w}(t)$ is smoothed by means of kernel function estimator of the baseline intensity function $\lambda_{0q}^{A_w}(t)$.

In our simulations (Section 5) and COVID-19 case study (Section 6), we used both the standard normal density $\phi(x)$ and Epanechnikov kernel functions and they yielded similar results. We presented results from the normal density kernel function. Following ?, we chose the optimal bandwidth b as the value that minimizes the mean integrated squared error (MISE)

$$\text{MISE}(\hat{\lambda}_0) = E \int_0^{t^o} \left[\hat{\lambda}_0(t) - \lambda(t) \right]^2 dt.$$

S1.2 Random survival forests accommodating time-varying covariates

Approach (iii) uses a recent random survival forests model (?) that accommodates time-varying covariates to reduce the parametric assumptions about the form of the intensity ratio function required by the usual proportional intensity regression model. ? proposed a forest method that estimates the survival function by three steps. In step (1), reformat the data in the counting process structure, that is, for individual i , the time-varying covariate $L^{(i)}(t)$ will be represented as $L^{(i)}(t) = l_j^{(i)}, t \in [t_j^{(i)}, t_{j+1}^{(i)})$, $j = 0, \dots, J^{(i)} - 1$. Then split the individual i observation into $J^{(i)}$ pseudo subject observations: $(t_j^{(i)}, t_{j+1}^{(i)}, \delta_j^{(i)}, l_j^{(i)})$ with left-truncated right-censored (LTRC) times $t_j^{(i)}, t_{j+1}^{(i)}$ and event indicator $\delta_j^{(i)}$ for the time interval $[t_j^{(i)}, t_{j+1}^{(i)})$. Pool the counting process styled records from N subjects to create a list of pseudo-subjects,

$$\left\{ t'_l, t'_{l+1}, \delta'_l, l'_l \right\}_{l=1}^n, \quad n = \sum_i J^{(i)}. \quad (1)$$

The set of pseudo-subjects is treated as if they were independent. Step (2) is to apply the forest algorithms on the reformatted dataset given in (1), to fit a model. In step (3), given a particular stream of covariate values at the corresponding time values, a survival function estimate is constructed based on the outputs of the forest algorithms.

We now briefly describe the forest algorithms. ? extended the relative risk forests, which combines the relative risk trees (?) with random forest methodology (?), for LTRC data by modifying the splitting criteria. The splitting criterion under the relative risk framework is to maximize the reduction in the one-step deviance between the log-likelihood of the saturated model and the maximized log-likelihood. Let \mathcal{R}_h denote the set of observations that fall into the node h . Let Λ_0 index the baseline cumulative hazard function, φ_h represent the nonnegative relative risk of the node h ,

and t_l and δ_l be the time and event indicator of the l th observation $\forall l \in \mathcal{R}_h$. Given the right censored observations (t_l, δ_l) , the full likelihood deviance residual for node h is defined as

$$d_h = \sum_{l \in \mathcal{R}_h} 2 \left\{ \delta_l \log \left(\frac{\delta_l}{\hat{\Lambda}_0(t_l) \hat{\varphi}_h} \right) - \left(\delta_l - \hat{\Lambda}_0(t_l) \hat{\varphi}_h \right) \right\}. \quad (2)$$

Two steps are needed to modify the splitting rule and obtain the deviance residual appropriate for LTRC data (?). First, compute the estimated cumulative hazard function $\hat{\Lambda}_0(\cdot)$ based on all pseudo-subject observations. Second, replace $\hat{\Lambda}_0(t_l)$ in (2) with $\hat{\Lambda}_0(t'_{l+1}) - \hat{\Lambda}_0(t'_l)$, and replace δ_l in (2) with δ'_l . We refer to ? for more detailed description of the random survival forests model.

S1.3 Variance estimation

Since our estimators for ψ (marginal structural proportional hazards model parameter) can be considered as solution to the weighted partial score equation, we consider the robust sandwich variance estimator as a convenience device to construct confidence intervals. The robust sandwich variance estimator has been considered, for example, in ?, and ??, and has been shown to be at most conservative under the discrete-time setting. We use this estimator for inference in conjunction with our continuous-time stabilized inverse probability weights, and formally evaluate its performance in our simulations. We briefly describe the robust sandwich variance estimator below. With the continuous-time weights, we focus on equation (11) in the main text with two treatments:

$$\sum_{i=1}^n \int_0^\infty \Omega^{A_1, A_2} \Omega^C(G_i) \left\{ Z(A_{1i}, A_{2i}, t) - \frac{S^{(1)}(t; \psi)}{S^{(0)}(t; \psi)} \right\} dN_i^T(t) = 0,$$

where $\Omega^{A_1, A_2} \Omega^C(G_i)$ is the weight for time-varying treatments A_1 and A_2 and censoring (in continuous time), $Z(A_{1i}, A_{2i}, t)_{(3 \times 1)} = [A_{1i}(t), A_{2i}(t), A_{1i}(t)A_{2i}(t)]^\top$, and

$$\begin{aligned} S^{(0)}(t; \psi) &= \sum_{k \in \mathcal{R}_t^T} Y_k^{**T}(t) r(A_{k1}, A_{k2}, t; \psi) \\ S^{(1)}(t; \psi) &= \sum_{k \in \mathcal{R}_t^T} Z(A_{k1}, A_{k2}, t) Y_k^{**T}(t) r(A_{k1}, A_{k2}, t; \psi). \end{aligned}$$

In the above definition, \mathcal{R}_t^T refers to the collection of subjects still at risk for the outcome event at time t , $Y_k^{**T}(t) = \Omega^C(G_i) \Omega^{A_1, A_2}(t_{K_i}) Y_k^T(t)$, where $Y_i^T(t)$ is the at-risk function for the outcome event, and $r(a_1, a_2, t) = \exp\{\psi_1 a_1(t) + \psi_2(t) a_2(t) + \psi_3 a_1(t) a_2(t)\}$. Now further define

$$S^{(2)}(t; \psi) = \sum_{k \in \mathcal{R}_t^T} Z(A_{k1}, A_{k2}, t)^{\otimes 2} Y_k^{**T}(t) r(A_{k1}, A_{k2}, t; \psi),$$

with $a^{\otimes 2} = aa^\top$ for any vector a . Then the robust sandwich variance estimator takes the form $\Sigma_0^{-1}\Sigma_1\Sigma_0^{-1}$ (?), where

$$\Sigma_0 = \sum_{i=1}^n \int_0^\infty \Omega^{A_1, A_2} \Omega^C(G_i) \left\{ \frac{S^{(2)}(t; \psi)}{S^{(0)}(t; \psi)} - \frac{S^{(1)}(t; \psi)^{\otimes 2}}{S^{(0)}(t; \psi)} \right\} dN_i^T(t),$$

and

$$\Sigma_1 = \sum_{i=1}^n \left[\int_0^\infty \Omega^{A_1, A_2} \Omega^C(G_i) \left\{ Z(A_{1i}, A_{2i}, t) - \frac{S^{(1)}(t; \psi)}{S^{(0)}(t; \psi)} \right\} dM_i^T(t) \right]^{\otimes 2},$$

and $M_i^T(t) = N_i^T(t) - \int_0^t Y_k^{**T}(u) \lambda_0(u) r(A_{k1}, A_{k2}, u; \psi) du$ is the martingale based on the counting process for outcome event. The sandwich variance estimator is obtained when both Σ_0 and Σ_1 are evaluated at the estimated weights and $\hat{\psi}$.

Because the above robust variance estimator considers the weights $\Omega^{A_1, A_2} \Omega^C(G_i)$ as fixed known values (?), it could result in conservative (but still valid) inference. With time-invariant weights estimated by logistic regression in the cross-sectional treatment setting, the corrected robust sandwich variance estimator has been derived to achieve improved variance estimation for hazard ratio parameters (?). However, an extension to continuous-time weights is not trivial and currently unavailable. Alternatively, resampling method such as the bootstrap method (??) could be used to make more robust inference for ψ that accounts for the uncertainty of the weights. Although theoretical valid, this resampling approach is not pursued in our current work due to the substantially more intensive computations associated with repeated estimation of complex weights under the recurrent event framework.

S2 Marginal Structural Cox Model Simulation Algorithm

Here we provide pseudocodes for the marginal structural cox model data simulation. We use the COVID-19 dataset as the foundation to set the values of the parameters in both the treatment assignment and marginal structural models. In our simulations, to reduce the impact of data sparsity, we set the maximum number of treatment initiations to be 4. To do this, the pseudocodes can be modified by setting treatment exposure status to one at all time points after the fourth treatment initiation throughout to the end of the follow-up period.

GET

$M \leftarrow 100$ (maximum follow-up); $\lambda_0 \leftarrow 0.005$; $n \leftarrow 1000$;
 $\beta \leftarrow (\beta_0, \beta_1, \beta_2, \beta_3) \leftarrow [\log(3/7), -0.5, -\log(1/2), \log(3/2)]$ (parameter vector for generating time-varying confounding variables L_2);

$\zeta \leftarrow (\zeta_0, \zeta_1, \zeta_2, \zeta_3, \zeta_4) \leftarrow [\log(2/7), -\log(1/2), -0.5, \log(3/2), \log(2/3)]$ (parameter vector for generating time-varying confounding variables L_1);

$\gamma \leftarrow (\gamma_0, \gamma_1, \gamma_2, \gamma_3, \gamma_4, \gamma_5, \gamma_6, \gamma_7, \gamma_8, \gamma_9, \gamma_{10}, \gamma_{11})$
 $\leftarrow [\log(2/7), 1/2, -1/2, -\log(3/5), 0.8, 0.5, 0.8, -0.5, 1/2, 1.2, -0.6, -0.3]$ (parameter vector for generating A_1);

$\eta \leftarrow (\eta_0, \eta_1, \eta_2, \eta_3, \eta_4, \eta_5, \eta_6, \eta_7, \eta_8, \eta_9, \eta_{10}, \eta_{11})$
 $\leftarrow [\log(3/7), 1/3, -1/3, -\log(2/5), 0.9, 0.6, 0.8, -0.5, 1/3, 0.9, -0.6, -0.4]$ (parameter vector for generating A_2);

$\psi_1 \leftarrow -0.5$ (true log-hazard value representing the effect of treatment A_1)

$\psi_2 \leftarrow -0.3$ (true log-hazard value representing the effect of treatment A_2)

COMPUTE

for ID $\leftarrow 1$ to n (for each individual) **do**

INIT: $L_1(-1) \leftarrow 0; L_2(-1) \leftarrow 0; A_1(-1) \leftarrow 0; A_2(-1) \leftarrow 0; Y(0) \leftarrow 0; H(m) \leftarrow 0;$

$n_{A_1} \leftarrow 0; n_{A_2} \leftarrow 0; \tau_{A_1} \leftarrow 0; \tau_{A_2} \leftarrow 0$

$T^{\bar{0}} \sim \text{Exponential}(\lambda_0)$

for $m \leftarrow 0$ to M **do**

$L_1(m) \leftarrow E(L_1(m) = l_1(m) \mid L_1(m-1), A_1(m-1), A_2(m-1), Y(m) = 0; \zeta)$

$\leftarrow \zeta_0 + \zeta_1(1/\log T^{\bar{0}}) + \zeta_2 A_1(m-1) + \zeta_3 L_1(m-1) + \zeta_4 A_2(m-1)$

$\text{logit}(p_{L_2}) \leftarrow \text{logit}P(L_2(m) = 1 \mid L_2(m-1), A_1(m-1), A_2(m-1), Y(m) = 0; \beta)$

$\leftarrow \beta_0 + \beta_1 A_1(m-1) + \beta_2 L_2(m-1) + \beta_3 A_2(m-1)$

$L_2(m) \sim \text{Bernoulli}(p_{L_2})$

$\text{logit}(p_{A_1}) = \text{logit}P(A_1(m) = 1 \mid L_1(m), L_1(m-1), L_2(m), L_2(m-1), A_1(m-1),$

$A_2(m-1), n_{A_1}, Y(m) = 0; \gamma)$

$= \gamma_0 + \gamma_1 A_1(m-1) + \gamma_2 (L_2(m-1))^2 + \gamma_3 (L_1(m-1))^2 + \gamma_4 (A_1(m-1)L_1(m))$

$+ \gamma_6 (L_1(m)L_2(m)) + \gamma_7 (A_1(m-1)L_2(m)) + \gamma_8 A_2(m-1)$

$+ \gamma_9 (A_2(m-1)L_1(m)) + \gamma_{10} (A_2(m-1)L_2(m)) + \gamma_{11} n_{A_1}$

if $A_1(m-1) = 0$ or $m-1 = \tau_{A_1}$ **then**

$A_1(m) \sim \text{Bernoulli}(p_{A_1})$

if $A_1(m) = 1$ **then**

$\delta_{A_1} \sim \text{zero-truncated Poisson}(10)$ (treatment duration after initiation)

$\tau_{A_1} \leftarrow m + \delta_{A_1}$

$A_1(m+1) : A_1(\max(\tau_{A_1}, M)) \leftarrow 1$

```

     $n_{A_1} \leftarrow n_{A_1} + 1$ 
  end if
end if
logit( $p_{A_2}$ ) = logit $P(A_2(m) = 1 \mid L_1(m), L_1(m-1), L_2(m), L_2(m-1), A_1(m),$ 
     $A_2(m-1), n_{A_2}, Y(m) = 0; \boldsymbol{\eta})$ 
     $= \eta_0 + \eta_1 A_1(m-1) + \eta_2 (L_2(m-1))^2 + \eta_3 (L_1(m-1))^2 + \eta_4 (A_1(m) L_1(m))$ 
     $+ \eta_6 (L_1(m) L_2(m)) + \eta_7 (A_1(m) L_2(m)) + \eta_8 A_2(m-1)$ 
     $+ \eta_9 (A_2(m-1) L_1(m)) + \eta_{10} (A_2(m-1) L_2(m)) + \eta_{11} n_{A_2}$ 
  if  $A_2(m-1) = 0$  or  $m-1 = \tau_{A_2}$  then
     $A_2(m) \sim \text{Bernoulli}(p_{A_2})$ 
    if  $A_2(m) = 1$  then
       $\delta_{A_2} \sim \text{zero-truncated Poisson}(9)$  (treatment duration after initiation)
       $\tau_{A_2} \leftarrow m + \delta_{A_2}$ 
       $A_2(m+1) : A_2(\max(\tau_{A_2}, M)) \leftarrow 1$ 
       $n_{A_2} \leftarrow n_{A_2} + 1$ 
    end if
  end if
end if
 $H_m \leftarrow H_m + \exp\{\psi_1 A_1(m) + \psi_2 A_2(m)\}$ 
  if  $T^{\bar{0}} \geq H_m$ 
     $Y_{m+1} \leftarrow 0$ 
  else
     $Y_{m+1} \leftarrow 1$ 
     $T \leftarrow m + (T^{\bar{0}} - H_m) \times \exp\{-\psi_1 A_1(m) - \psi_2 A_2(m)\}$ 
  end if
end for m
end for ID

```




Figure 1: The distributions of biases across 250 simulation replications of the ragged longitudinal data with unaligned time points, in estimating the parameters ψ_1 and ψ_2 using the proposed JMSSM-CT method and the comparison method JMSM-DT. When implementing JMSM-DT, the follow-up time was respectively discretized into time intervals of length 0.5, 1 and 2 days, as the method requires aligned measurement time points.

S3 Supplementary figures and tables

S3.1 Additional simulation results

Table 1: Mean absolute bias (MAB), root mean square error (RMSE) and coverage probability (CP) for the estimates of ψ across 250 data replications with unaligned follow-up time points, using four weight estimators (i)-(iv) described in Section 3.4.

Weight estimators	ψ_1			ψ_2		
	MAB	RMSE	CP	MAB	RMSE	CP
(i)	.102	.104	.048	.092	.094	.060
(ii)	.063	.067	.104	.055	.059	.112
(iii)	.023	.029	.936	.020	.025	.940
(iv)	.016	.022	.952	.015	.019	.952

Table 2: The distribution of the estimated individual time-varying weights from one random replication of the “ragged” longitudinal data with unaligned time points, for the proposed JMSSM-CT versus JMSM-DT. To estimate the weights, the random forests was used for JMSM-DT and four approaches (i)-(iv) (Section 3.4) were used for JMSSM-CT.

Methods	Weight estimators	Distribution of estimated weights				
		Minimum	First quartile	Mean	Third quartile	Maximum
JMSSM-CT	(i)	0.23	0.89	1.05	1.23	5.34
	(ii)	0.40	0.98	1.01	1.07	4.28
	(iii)	0.52	0.88	1.00	1.09	2.99
	(iv)	0.68	0.95	1.00	1.05	2.36
JMSM-DT	Random forests	0.43	0.90	1.03	1.11	3.65

S3.2 Treatment classes for COVID-19

In Table 3, we provide detailed definitions of the four treatment classes for COVID-19 whose comparative effectiveness on in-hospital death was investigated in Section 6.

Table 3: Definitions of four treatment classes for COVID-19. iv:intravenous; po: by mouth.

Treatment classes	Medication (route)
Dexamethasone	Dexamethasone (iv), Dexamethasone (po)
Remdesivir	Remdesivir (iv)
Corticosteroids other than dexamethasone	Hydrocortisone (po), Hydrocortisone (iv), Methylprednisolone (po), Methylprednisolone (iv), Prednisolone (iv), Prednisone (po), Prednisone (iv)
Anti-inflammatory medications other than corticosteroids	Alpha-1-Proteinase Inhibitor (iv), Anakinra (iv), Azathioprine (po), Belatacept (iv), Dexamethasone (iv), Dexamethasone (po), Eculizumab (iv), Envarsus (iv), Sarilumab (iv), Gengraf (po), Gengraf (iv), Hydrocortisone (po), Hydrocortisone (iv), Ibrutinib (po), Immune Globulin (iv), Infliximab (iv), Methylprednisolone (po), Methylprednisolone (iv), Montelukast (po), Prednisolone (iv), Prednisone (po), Prednisone (iv), Ruxolitinib (po), Tocilizumab (iv), Apremilast (po), Celecoxib (po), Dasatinib (po), Everolimus (iv), Gemtuzumab (iv), Ibuprofen (iv), Ibuprofen (po), Ifosfamide (iv), Leflunomide (po), Mesalamine (iv), Methotrexate (iv), Methylphenidate (iv), Mycophenolate (iv), Naproxen (po), Rituximab (iv), Sulfasalazine (po), Tacrolimus (iv), Pacritinib (po), Risankizumab (iv), Daratumumab (iv), Talquetamab (iv)

S3.3 Patient oxygen levels

We describe how patient oxygen levels are categorized based on the use of ventilator in Table 4.

Table 4: Definitions of patient oxygen levels based on the use of ventilator

Patient oxygen level	Ventilator status
0	Room air
1	Cannula
2	Mask, Blow-by, Face tent, Oxyhood, Non-rebreather, RAM cannula
3	Continuous positive airway pressure machine, High flow nasal cannula, Hudson prongs
4	Bilevel positive airway pressure machine, Tracheostomy mask
5	Tracheotomy, Transtracheal oxygen therapy, Ventilator, Endotracheal tube, T-shaped tubing connected to an endotracheal tube, Nasal synchronized intermittent mandatory ventilation