Bayesian Framework for Predictive and Causal Modeling with Application to the HIV Care Cascade

Yizhen Xu*
Department of Biostatistics, Johns Hopkins University
Joseph Hogan
Department of Biostatistics, Brown University
Michael Daniels
Department of Statistics, University of Florida
Ann Mwangi
College of Health Sciences, School of Medicine, Moi University
Tao Liu
Department of Biostatistics, Brown University
Rami Kantor
Division of Infectious Diseases, Brown University

July 31, 2022

Abstract

We develop a Bayesian method for semiparametric predictive and causal inference on longitudinal multi-state outcomes. We use a Bayesian causal formulation of the g computation algorithm (GCA) and incorporate Bayesian additive regression trees (BART) as the generative components to reduce the need for parametric model specifications and enable machine learning based prediction of events using time-evolving models. Our method provides a way to conduct predictive and causal inference based on posterior predictive distributions of the counterfactual outcomes over time. The work is motivated by the electronic health records (EHRs) from the Academic Model Providing Access to Healthcare (AMPATH) in Kenya. We use the data to investigate dynamic treatment regimes by comparing their causal effects on the progression of patients' engagement in care through the HIV care cascade, which is framed as a time-evolving multi-state outcome with dependent alternatives. Under settings involving dynamical systems that can be described by state transitions over time, the proposed framework can be applied broadly to understand complex interventions that

*The authors gratefully acknowledge that *funding of this work is provided by the US National Institutes of Health (NIH) under R01 AI136664, R01 AI108441, R01 AI167694, and P30 AI 42853.*

may depend on the progression of outcome and confounders when massive data are available.

Keywords: Bayesian G computation, BART, competing risks

1 Introduction

Prevention and treatment of HIV remains a high public health priority. The HIV care cascade (Perlman et al., 2017; WHO, 2012; Gardner et al., 2011) is a conceptual model that guides systematic evaluations of public health interventions and supports solution-seeking in pursuit of HIV care benchmarks (Mugglin et al., 2021; Vourli et al., 2020). The HIV care cascade outlines essential stages of HIV care as (a) HIV diagnosis through testing, (b) linkage to and engagement in care, (c) initiation of antiviral therapy, and (d) sustained suppression of viral load through continued engagement and retention in care. The Joint United Nations Program on HIV/AIDS (UNAIDS) set the 95-95-95 goal (UNAIDS, 2014) for HIV treatment in 2014, aiming to achieve three 95% targets in HIV care by 2030: 95% diagnosis rate among people with HIV infection, 95% on sustainable treatment among diagnosed individuals, and 95% viral suppression among those treated. To meet these goals, it is critical to understand patient-level dynamics of progression through the cascade to better manage HIV care programs (Kay et al., 2016).

Patient engagement in HIV care, reflected by regular clinic visits, implies a continuous and stable relationship with health care providers. It promotes adherence to medications and HIV viral load suppression, deters development of drug resistance, improves health outcomes and reduce health care costs, and encourages positive behavioral change that lowers HIV transmission rate and reduces HIV population burden (Horstmann et al., 2010). Measures of engagement in care are typically based on patient-level appointment history.

Antiretroviral treatment (ART) initiation strategies in HIV care can impact not only disease outcomes but behavioral ones, such as engagement and retention in care. The current World Health Organization guidelines (Organization et al., 2015) call for initiation of ART for all persons newly diagnosed with HIV regardless of disease status indicators

such as CD4 count. The Strategic Timing of AntiRetroviral Treatment (START) study (Insight Start Study Group, 2015), an international randomized trial, showed that immediate initiation of ART upon HIV diagnosis reduces the risk of serious illness and death after three years of follow-up. In real-world settings, continued adherence with treatment requires engagement and retention in care. However, studying the impact of treatment policies on engagement and retention is not feasible in a clinical trial because of the controlled environment. Our analysis is therefore designed to investigate the causal effect of early ART initiation on engagement and retention in care using observational data from a large HIV care program in western Kenya. To quantify this effect, our analysis compares two treatment strategies: one that initiates ART immediately at the first clinic encounter, and one based on prior recommendations to initiate treatment when CD4 cell count drops below 350 cells/mm^3.

Our analysis uses data from the electronic health records (EHRs) of the Academic Model Providing Access to Healthcare (AMPATH) program in Kenya (Tierney et al., 2013). The EHR cohort include adults living with HIV who were not previously in AMPATH HIV care and whose enrollment date was between June 1st 2008 and August 23rd 2016, and the date of data closure was August 24th 2016. Our goal is to draw inference about the effect of treatment regime on care status over time. At any given time, a patient is engaged, disengaged, transferred out, or died. We therefore formulate the evolution of patient status as a multi-state process following the framework introduced in Lee et al. (2017). We use Bayesian g computation algorithm (GCA) (Keil et al., 2017; Robins, 1986; Robins et al., 1999) to draw inference about the causal impact of treatment strategy on the evolution of patient outcomes.

The GCA is a method for estimating the effect of complex interventions that may depend

on the progression of outcome, treatment, and confounders. Like all model-based causal inference methods, it assumes correct specification of the distributions of key processes. GCA requires the users to specify the time-varying outcome and confounder distributions. To ensure flexibility in model specification, we use Bayesian additive regression trees (BART) (Chipman et al., 2010; Keil et al., 2017) . Given a treatment regime of interest, predictions at the individual level and inference for the entire population of interest can then be obtained from the posterior predictive distribution of the counterfactual trajectories.

GCA for longitudinal data has been used in multiple settings. Petersen and van der Laan (2011) and Young et al. (2011) used longitudinal observational HIV data to estimate the causal effect of dynamic treatment regimes on HIV/AIDS-related mortality, using linear model specifications for distributions of time-dependent variables conditional on covariates related to clinical histories. Zhou et al. (2019) combined the parametric GCA with multiple imputation of missing data using penalized splines. Keil et al. (2017) introduced the Bayesian approach to GCA for static interventions. Antonelli and Daniels (2019) applied Bayesian nonparametric models using the Dirichlet process for implementing Bayesian GCA. Bayesian semi-parametric GCA using BART is used by Josefsson and Daniels (2019) to infer the causal effect of deterministic longitudinal intervention assignment on a time-varying continuous outcome subject to missingness due to dropout.

In our application, we use multinomial probit BART and model retention in HIV care as a time-evolving multi-state variable. The states are not required to be independent, i.e. the ratio of likelihood between engagement versus disengagement may be affected by the chance of death. In the AMPATH data, the states are extremely imbalanced such that rate of death and transfer-out are extremely low compared to engagement and disengagement. Hence, allowing dependence between states enables us to borrow strength in handling sparse

categories from the more frequent ones. Moreover, Lemnaru and Potolea (2011) observed that pruning deteriorates the performance of the decision trees as imbalance increases in binary outcomes; this further encourages our usage of BART because it reduces the size of individual trees by a strongly influential prior distribution instead of pruning as in frequentist tree methods. The sampling strategy in Xu et al. (2021) implements a multinomial probit version of BART and is designed to address these complications.

We lay out the rest of this paper as follows. Section 2 describes the structure of data from AMPATH EHR. In Section 3, we introduce the causal model of interest and associated notation, describe our proposed Bayesian approach including the assumptions and formulation of Bayesian GCA, provide details of model specifications, and outline posterior inference as a three-step procedure. In Section 4, we present an analysis of the AMPATH data. Section 5 provides a discussion of the proposed Bayesian approach.

2 Notations and Data Structure for the Observed Data

Following Lee et al. (2017), we adopt a state transition framework that defines patient state at regular time intervals each of length 200 days, indexed by $k = 1, \ldots, K$ with $k = 0$ being the study baseline. We set $K = 10$ based on duration of follow up in cohort. The study baseline is defined as a 12-day interval following enrollment on day 0, during which the initial measurement of CD4 typically happens as part of the enrollment process; unlike other post-baseline intervals, the first interval is (12,200] instead of (0, 200]. Time-invariant covariates at baseline are denoted as V. We define three stochastic processes, $\{S_k : k \geq 0\}$, $\{X_k : k \geq 0\}$, and $\{A_k : k \geq 0\}$, representing outcome state, time-varying confounders, and treatment status over time, respectively.

Patient status S_k is a multinomial random variable coded as 0 for disengagement, 1 for

engagement in care, 2 for transfer-out, and 3 for reported death. Patients are classified as engaged if they had at least one care encounter in interval k, and disengaged otherwise if there is no indication of reported death or transfer. We set $S_0 \equiv 1$ for all patients at the initial care encounter. If a patient has a care encounter in interval k that indicates a transfer-out, the person is considered to be engaged at interval k and transferred out at $k+1$, i.e. $S_k = 1$ and $S_{k+1} = 2$. Reported death and transfer-out are absorbing states, i.e. $S_{k'} = S_k$ for $k' > k$ if $S_k \in \{2, 3\}$. Figure 1 presents possible transitions of S_k to S_{k+1} for $k \geq 0$. Figure 2 summarizes the distribution of observed outcomes in the data and shows that the total number of observations decreases over time due to staggered entry over time. Table 1 summarizes the observed aggregated transition probabilities over time.

For time interval k, the variable A_k is an indicator of being initiated on antiretroviral treatment in HIV care during the interval. Following Young et al. (2011), a jth time-varying confounder is modeled as (L_{kj}, R_{kj}) in interval k, where R_{kj} is the binary indicator of whether at least one update of the confounder occurred in the interval and L_{kj} is the carried-forward measurement. Suppose there are p time-varying confounders, then they are modeled as a vector-valued stochastic process $X_k = (R_k, L_k)$, where $L_k = (L_{k1}, \ldots, L_{kp})$ and $R_k = (R_{k1}, \ldots, R_{kp})$. For example, CD4 is a time-varying confounder and its most recent measurement stays unchanged ($L_{kj} = L_{k-1,j}$) when it is not updated in the current interval ($R_{kj} = 0$). We let X_{-1} and X_0 represent measurements during the pre-baseline period of days (-200, 0] and study baseline (0,12]. Let the overbar notation denote longitudinal paths, e.g. $\bar{A}_k = \{A_l : l = 0, \ldots, k\}$ is the longitudinal path of interventions up to time interval $k \in \{0, \ldots, K\}$.

3 Bayesian G-Computation for Multistate Model

3.1 Causal Quantities

The goal is to set up a causal framework to compare treatment regimes. As in Young et al. (2011), we define a treatment regime to be a function $q(\cdot)$ that dynamically maps data history to a current intervention status. The intervention sequence $\bar{a}_k^q = (a_1^q, \ldots, a_k^q)$ is sequentially determined by regime q, i.e. $a_k^q = q(\mathcal{F}_k)$ for $k = 0, \ldots, K$, where $\mathcal{F}_k = \{v, \bar{x}_0, s_0, a_{l-1}, x_{l-1}, s_l : l = 1, \ldots, k\} = \{v, \bar{a}_{k-1}, \bar{x}_{k-1}, \bar{s}_k\}$ for $k = 1, \ldots, K$ represents an accumulated history up to interval k. Here we use $\mathcal{F}_0$ to denote $(v, \bar{x}_0, s_0)$ at $t = 0$ for uniformity of notation. We use the superscript q to denote potential outcomes under any regime q, e.g., $\bar{S}_k^q$ and $\bar{X}_k^q$ represent the longitudinal paths of potential outcomes and confounders under a treatment regime q.

For the analysis here, we are interested in comparing the causal effectiveness of dynamic treatment regimes expressed in the form of $a_k^q = q(a_{k-1}, X_{k-1}, S_k)$. For example, one of the regimes being compared in Section 4 is to initiate treatment when CD4 cell count drops below 350 cells/mm^3, i.e. treatment status is a function of the time-varying confounder, CD4 cell count. The causal comparison of two different regimes q_1 and q_2 across time can be represented in terms of $P(S_k^{q_1})$ and $P(S_k^{q_2})$, the marginal distribution of the potential outcome under q_1 and q_2 at each interval $k = 1, \ldots, K$.

3.2 The G-Formula and Assumptions

In this section we describe calculation of target quantity $P(S_k^q)$, the marginal distribution of potential outcome at time interval k had all subjects followed regime q for $k = 1, \ldots, K$. We state the assumptions for identifying $P(S_k^q)$ and show how its posterior predictive

distribution is obtained using the G-formula.

Let $\mathcal{S}$ denote the set of non-absorbing state(s) of the outcome, e.g. $\mathcal{S} = \{0, 1\}$ under the data structure described in Section 2. The set of absorbing state(s) is then $\mathcal{S}^c = \{2, 3\}$. Similar to Robins et al. (1999) and Young et al. (2011), we make the following identifying assumptions:

1. Consistency: if $\bar{A}_k = \bar{a}_k^q$, then $\bar{S}_{k+1} = \bar{S}_{k+1}^q$ and $\bar{X}_{k+1} = \bar{X}_{k+1}^q$

2. Positivity : $P(A_k = a_k^q | \bar{X}_{k-1}, \bar{S}_k, \bar{A}_{k-1} = \bar{a}_{k-1}^q) > 0 \quad$ w.p.1 for $S_k \in \mathcal{S}$, $k = 0, \ldots, K$.

3. Sequential exchangeability:

$$(S_{k+1}^q, \ldots, S_K^q) \perp A_k \mid (\bar{X}_k, \bar{S}_k, \bar{A}_{k-1} = \bar{a}_{k-1}^q),$$

   for $S_k \in \mathcal{S}$, $k = 0, \ldots, K-1$.

Consistency states that counterfactuals follow observed variables' distribution when the regime of interest is observed. The positivity assumption is that no regime is systematically ruled out in practice. Under sequential exchangeability, the study resembles a sequentially randomized trial in which initiation of treatment at each time point is randomized with a probability that depends on the accumulated longitudinal history. We further assume that the generative component models for $[X_k | \bar{X}_{k-1}, \bar{S}_{k-1}, \bar{A}_{k-1}, V]$ and $[S_k | \bar{X}_k, \bar{S}_{k-1}, \bar{A}_{k-1}, V]$ can be specified via flexible parametric models indexed by finite-dimensional parameters, such as BART. For now we denote these models as

$$P(X_k | \bar{X}_{k-1}, \bar{S}_{k-1}, \bar{A}_{k-1}, V, \gamma_k) \text{ and} \tag{1}$$

$$P(S_k | \bar{X}_k, \bar{S}_{k-1}, \bar{A}_{k-1}, V, \theta_k), \tag{2}$$

for $k = 1, \ldots, K$. Let $\bar{\theta}_k = (\theta_1, \ldots, \theta_k)$ and $\bar{\gamma}_k = (\gamma_1, \ldots, \gamma_k)$ denote the parameters for models up to time interval k, and assume that $(\bar{\theta}_k, \bar{\gamma}_k)$ are a priori independent.

The marginal outcome distribution $P(S_k^q = s_k|\bar{\theta}_k, \bar{\gamma}_k)$ can be calculated using the G-formula as follows (see Appendix A.1):

$$P(S_k^q|\bar{\theta}_k, \bar{\gamma}_k) = \int \cdots \int_{(V,X_0,S_0,\ldots,X_{k-1},S_{k-1},X_k)} \left\{ \prod_{j=1}^{k} P(S_j|\bar{X}_j, \bar{S}_{j-1}, \bar{a}_{j-1}^q, V, \theta_j) \right.$$
$$\left. \times P(X_j|\bar{X}_{j-1}, \bar{S}_{j-1}, \bar{a}_{j-1}^q, V, \gamma_j) \right\} \times P(V, X_0, S_0) dX_k dS_{k-1} dX_{k-1} \ldots dS_0 dX_0 dV. \tag{3}$$

Given a posterior distribution of $(\bar{\theta}_k, \bar{\gamma}_k)$, we can simulate outcomes from the posterior predictive distribution of S_k^q as follows. Write the observed data as $\mathcal{D} = \{V_i, \bar{A}_{iK_i}, \bar{X}_{iK_i}, \bar{S}_{iK_i}; i = 1, \ldots, N\}$, where K_i is the length of follow-up for the ith subject. Define $\overline{\mathcal{D}}_k$ as all available observations up to interval k, and $\mathcal{D}_k$ is the subset of $\overline{\mathcal{D}}_k$ where every subject had at least k intervals of follow-up, i.e. $\mathcal{D}_k = \{\overline{\mathcal{D}}_k \mid K_i \geq k\}$. The posterior predictive distribution of S_k^q given the observed data $\mathcal{D}$ follows by integrating over the posterior of parameters:

$$P(S_k^q|\overline{\mathcal{D}}_k) = \int_{(\bar{\theta}_k, \bar{\gamma}_k)} P(S_k^q|\bar{\theta}_k, \bar{\gamma}_k) P(\bar{\theta}_k, \bar{\gamma}_k|\overline{\mathcal{D}}_k) d(\bar{\theta}_k, \bar{\gamma}_k), \tag{4}$$

where $P(S_k^q|\bar{\theta}_k, \bar{\gamma}_k)$ is calculated using the G-formula given in (3) and $P(\bar{\theta}_k, \bar{\gamma}_k|\overline{\mathcal{D}}_k)$ is the posterior density of the model parameters. Under the assumption of a priori independence between $\bar{\theta}$ and $\bar{\gamma}$, the posteriors are independent:

$$P(\bar{\theta}_k, \bar{\gamma}_k|\overline{\mathcal{D}}_k) = \prod_{j=1}^{k} P(\theta_j|\mathcal{D}_j) P(\gamma_j|\mathcal{D}_j).$$

Please see Appendix A.2 for the detailed derivation.

To draw from $P(S_k^q|\overline{\mathcal{D}}_k)$, we first draw a sequence $(\bar{\theta}_k^*, \bar{\gamma}_k^*)$ from the posterior $P(\bar{\theta}_k, \bar{\gamma}_k|\overline{\mathcal{D}}_k)$. Using (4), we obtain $P(S_k^q|\bar{\theta}_k^*, \bar{\gamma}_k^*)$ as a posterior sample of $P(S_k^q|\overline{\mathcal{D}}_k)$. In this way, we can obtain the posterior predictive distribution for the target quantity $P(S_k^q|\overline{\mathcal{D}}_k)$ at $k = 1, \ldots, K$, and use it to compare the causal effectiveness of different regimes defined by q.

3.3 Specifying Component Models using BART

Model construction involves specifying generative component models and priors for R_k, L_k, and S_k, which are binary, continuous, and multinomial variables, respectively. For a predictive model at interval k, denote the corresponding model covariates for subject i as H_{ik}, the components of which may differ between models for R_k, L_k, and S_k based on the specifications for (1) and (2). Note that the model for L_k is conditional on $R_k = 1$. Please see Appendix A.3 for details of the prior specifications in the component models.

For binary time-varying confounders R_{ik}, which indicate whether the covariate L_{ik} has been measured, we assume a probit model

$$P(R_{ik} = 1|H_{ik}, \gamma_k^R) = \Phi[G(H_{ik}; T_k^R, M_k^R)], \tag{5}$$

where $\Phi(\cdot)$ is the standard normal cumulative density function and γ_k^R is the vector of model parameters. The regression function $G(\cdot; T_k^R, M_k^R)$ is parameterized as a sum of binary trees,

$$G(H_{ik}; T_k^R, M_k^R) = \sum_{j=1}^{n_T} g(H_{ik}; T_{kj}^R, M_{kj}^R),$$

where n_T is a user-specified total number of trees, $T_k^R = (T_{k1}^R, \ldots, T_{k,n_T}^R)$, $M_k^R = (M_{k1}^R, \ldots, M_{k,n_T}^R)$, T_{kj}^R is the jth tree consisting of interior nodes as decision rules and B_{kj}^R terminal nodes with parameters $M_j^R = (m_{1,j}^R, \ldots, m_{B_{kj}^R,j}^R)$. The jth component of G, $g(H_{ik}; T_{kj}^R, M_{kj}^R)$, contributes $m_{b,j}^R$ to the sum of trees if H_{ik} belongs to the b_{th} terminal node of the tree structure T_{kj}^R by falling through its decision rules from root to bottom.

For the continuous time-varying confounder L_{ik}, the component model is parameterized as

$$[L_{ik}|H_{ik}, \gamma_k^L, R_{ik} = 1] \sim N[G(H_{ik}; T_k^L, M_k^L), \sigma^2], \tag{6}$$

where $N(\cdot,\cdot)$ is the normal density function, $\gamma_k^L = (T_k^L, M_k^L, \sigma)$, and $G(H_{ik}; T_k^L, M_k^L)$ is defined as the sum of trees.

The multistate outcome S_{ik}, is a multinomial variable taking value in $\{0,\ldots,C\}$; we model S_{ik} as a function of covariates H_{ik}. Because the categories of S_{ik} might not be independent, we use a MNP having correlated latent utilities (van Dyk, 2010). Without loss of generality, we assume the reference category is $S_{ik} = 0$. The outcome is related to a latent vector variable W_{ik} via

$$S_{ik}(H_{ik}) = \begin{cases} l & \text{if } \max(W_{ik}) = W_{ikl} > 0 \\ 0 & \text{if } \max(W_{ik}) < 0, \end{cases} \tag{7}$$

where $W_{ik} = (W_{ik1},\ldots,W_{ikC})$ is a vector of latent utilities that follows multivariate Gaussian distribution

$$W_{ik}|H_{ik},\gamma_k^S \sim MVN(G(H_{ik}; T_k^S, M_k^S), \Sigma_k),$$

and $G(H_{ik}; T_k^S, M_k^S) = (G_1(H_{ik}; T_{k1}^S, M_{k1}^S),\ldots,G_C(H_{ik}; T_{kC}^S, M_{kC}^S))$. Each mean component is parameterized as sum of trees:

$$G_l(H_{ik}; T_{kl}^S, M_{kl}^S) = \sum_{j=1}^{n_T} g(H_{ik}; T_{klj}^S, M_{klj}^S),$$

where $l = 1,\ldots,C$. The corresponding prior distributions for the trees follows similar standard priors as for continuous response in BART. The covariance matrix Σ of the latent variables is a $C \times C$ positive definite matrix under the constraint that $\text{trace}(\Sigma) = C$ (Burgette and Nordheim, 2012). Please refer to Xu et al. (2021) for more details to our implementation of multinomial probit Bayesian additive regression trees (MPBART).

3.4 Posterior Inference

Posterior inference about causal effects can be carried out in three steps: (a) generate posterior draws of the parameters in generative component models (1) and (2), i.e.

$(\bar{\theta}_K^*, \bar{\gamma}_K^*) \sim P(\bar{\theta}_K, \bar{\gamma}_K|\mathcal{D})$, (b) validate the component models, and (c) under a given regime q of interest, generate the posterior predictive distribution of potential outcomes $(S_1^q, \ldots, S_K^q)$ using (3). The difference between posterior predictive outcome distributions under two different regimes then indicates the causal effect of implementing one regime versus another. Details for calculation are as follows.

Step I: Bayesian posterior inference for the predictive models. Based on the specifications (5), (6), and (7), we fit the generative models (1) and (2) at each time interval. There are several considerations in actual implementation. First, irregularly measured time-varying confounders L_{ik} (such as CD4 count) are not recorded at every interval k; the modeling of L_{ik} is conditional on $R_{ik} = 1$, i.e. a confounder is being updated. Second, model specifications need to be aligned with assumptions backed by the scientific problem of interest; for example, state outcome S_k may involve multiple models when the support of S_k varries conditional on the value of S_{k-1}. We generate N_{post} posterior draws for each specified generative component model, producing N_{post} posterior samples of the parameter sequence $(\bar{\gamma}_K, \bar{\theta}_K) = (\gamma_1, \theta_1, \ldots, \gamma_K, \theta_K)$. The posterior model parameters generated on the training data are then applied to the validation data for model evaluation.

Step II: Model validation. Assume N' is the sample size of the validation data. Stratified by outcome level l, we evaluate each categorical component model by comparing the observed rate $\frac{1}{N'}\sum_{i=1}^{N'} \mathbb{1}\{S_{ik} = l\}$ to the mean and 95% credible intervals of the posterior agreement accuracy, defined as $\{\frac{1}{N'}\sum_{i=1}^{N'} \mathbb{1}\{\hat{S}_{ijk} = l\}; j = 1, \ldots, N_{post}\}$ (Xu et al., 2021), where $\hat{S}_{ijk}$ is the jth posterior prediction for the outcome of the ith subject at the kth time interval. For binary outcomes, we set $l = 1$. For continuous outcomes, we compare the observed mean $\frac{1}{N'}\sum_{i=1}^{N'} S_{ik}$ to its posterior samples $\{\frac{1}{N'}\sum_{i=1}^{N'} \hat{S}_{ijk}; j = 1, \ldots, N_{post}\}$.

Step III: Generate posterior predictive counterfactual marginal outcome dis-

tribution under a dynamic treatment regime. Two treatment regimes, q_1 and q_2, can be compared by estimating the populating average treatment effect at each interval k, $P(S_k^{q_1}) - P(S_k^{q_2})$. For any regime q and a target population of interest, the calculation of the counterfactual marginal outcome distribution $P(S_k^q)$ starts from the baseline distribution $F(X_0, V)$. Given a sampled data from the target population as in our application, we initiate the inference from a bootstrapped empirical distribution of (X_0, V) for each posterior iteration; this Bayesian bootstrap procedure accounts for sampling variability in the target population baseline distribution $F(X_0, V)$. For a specific realization of baseline variables and generative model parameters, $(v^*, \bar{x}_0^*, s_0^*; \gamma_k^*, \theta_k^*)$, the procedure fixes the treatment sequence under q and makes posterior predictive draws of longitudinal counterfactual variables $(x_k^*(q), s_k^*(q))$ as follows

$$
\begin{aligned}
a_{k-1}^*(q) &= q(\mathcal{F}_{k-1}^*) \\
x_k^*(q) &\sim P(X_k | \bar{X}_{k-1} = \bar{x}_{k-1}^*(q), \bar{S}_{k-1} = \bar{s}_{k-1}^*(q), \bar{A}_{k-1} = \bar{a}_{k-1}^*(q), V = v^*, \gamma_k^*) \qquad (8) \\
s_k^*(q) &\sim P(S_k | \bar{X}_k = \bar{x}_k^*(q), \bar{S}_{k-1} = \bar{s}_{k-1}^*(q), \bar{A}_{k-1} = \bar{a}_{k-1}^*(q), V = v^*, \theta_k^*), \qquad (9)
\end{aligned}
$$

where $\mathcal{F}_{k-1}^*(q) = \{\bar{a}_{k-2}^*(q), \bar{x}_{k-2}^*(q), \bar{s}_{k-1}^*(q), v^*\}$ for $k = 1, \ldots, K$, with the definition that $\bar{a}_{-1}^* = 0$.

For the jth bootstrap subsample, we draw a counterfactual posterior longitudinal path for the ith individual, $\{a_{ij,k-1}^*(q), x_{ijk}^*(q), s_{ijk}^*(q); k = 0, \ldots, K\}$, where $j = 1, \ldots, N_{post}$ and k indexes time intervals. Write $\tilde{N}$ as the sample size of the bootstrap subsample. The target quantity is then estimated by averaging over the counterfactual outcomes at interval k as follows,

$$
p_{ljk}^{(q)} = \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} \mathbb{1}\{s_{ijk}^*(q) = l\}, \qquad (10)
$$

and $\{p_{ljk}^{(q)}; j = 1, \ldots, N_{post}\}$ is a sample from the posterior predictive distribution of the

target quantity $P(S_k^q = l|\overline{\mathcal{D}}_k)$. As a result, we can summarize the causal effectiveness of different dynamic treatment regimes at each interval k and outcome level l by the posterior mean

$$p_{lk}^{(q)} = \frac{1}{N_{post}} \sum_{j=1}^{N_{post}} p_{ljk}^{(q)}, \tag{11}$$

and its 95% posterior credible intervals.

4 Analysis of AMPATH Data

We apply the Bayesian approach in Section 3 to the AMPATH data described in Section 2. The goal of the application is to compare the causal effectiveness of different antiretroviral treatment initiation regimes on the longitudinal progression of the multistate engagement outcome over time. To accomplish this, we compare the posterior predictive distribution of the marginal probability of engagement status at each time interval under different treatment regimes. We compare three regimes: treat immediately, never treat, and treat when CD4 drops below 350 cells/mm^3. The regime "treat when CD4 drops below 350 cells/mm^3" assumes that: (a) treatment initiation at the baseline is based on CD4 measured during days (-200,0], (b) post-baseline initiation requires engagement in care and a previous or current measured CD4 less than 350 cells/mm^3 during that specific time interval, and (c) once a patient is treated, the person is always on treatment. In addition to existing literature (Young et al., 2011; Hontelez et al., 2011; The INSIGHT START Study Group, 2015) advocating benefits of early treatment initiation on individual patient outcomes, early treatment initiation can also improve patients' retention in care by generating an incentive to engage with the care system (e.g., filling prescription).

For model training and validation, we randomly separate individual records into two parts: $N = 50,000$ individuals are used to fit component models and $N' = 26,740$ indi-

viduals are used for model validation. We bootstrap a subsample of $\tilde{N} = 10,000$ from a total of $76,740$ individuals for each posterior iteration and generate posterior predictive comparisons across the three regimes of interest using the bootstrapped empirical baseline distributions. We assume a first-order Markov structure for the generative components (1) and (2) such that the time-varying confounders X_k and outcomes S_k only depend on $(X_{k-1}, S_{k-1}, A_{k-1}, V)$ and $(X_k, S_{k-1}, A_{k-1}, V)$, respectively.

4.1 Model Estimation and Validation

We include the following as time-invariant baseline variables V: gender, marriage status (married or otherwise), calendar year of enrollment (2008, 2009,..., 2016), baseline age in years, travel time to clinics ($<$ 30min, 30min to 1h, 1h to 2h, $>$ 2h), baseline WHO stage (primary infection, clinically asymptomatic stage, symptomatic HIV infection, and AIDS), height in cm, baseline weight in kg, and baseline viral load in log copies per mL. Because viral load has a lower detection limit of 50 copies per mL, we also include an indicator of undetectable baseline viral load. We do not model viral load as a time-varying confounder due to its infrequent measurement. Table 2 summarizes $(V, \bar{X}_0, A_0)$ at the baseline. Missingness in baseline covariates is handled by including missing indicators as the model covariates. For example, we code a variable V as $(U, (1-U)V)$, where $U = 1$ if the covariate is missing. Time-varying CD4 is modeled as (R_k, L_k) where R_k is an indicator of having a new CD4 measurement, and $L_k = \log(CD4+1)$ for the most recently measured CD4 count. The assumption of no unmeasured confounders in this application states that given patients' baseline information and their longitudinal history of CD4 progression, their future counterfactual state outcomes are independent of the current treatment status. Model fitting for the multistate outcome S_k is done using two separate conditional models,

one for S_{k-1} = 'engaged' and the other for S_{k-1} = 'disengaged', because by definition the probability of transition from 'disengaged' to 'transfer out' is zero.

Using the training data of 50,000 individuals, we fit the generative component models (1) and (2) under BART specifications (5), (6), and (7). For each model, we set the number of trees m as 100, an MCMC burn in of 10,000, and the number of posterior draws as 5000; then we thin the MCMC procedure by one draw per five samples, obtaining 1000 posterior samples for the GCA simulation. For implementation, we developed the R package *GcompBART* . Using the package, the model estimation step can be accomplished within 3.5 hours when 1 CPU is used for each paralleled computation of a chain under RedHat Linux system. We set the reference level in MPBART to the outcome level with the largest observed frequency for each multistate outcome model.

Next, we use diagnostic plots to assess convergence of the MCMC chains and validate model accuracy by comparing observed outcome distributions to the posterior predictive distributions of the validation data. We track the MCMC convergence of a BART model using the average depth among all posterior trees, which should be stable across MCMC iterations under convergence (Kapelner and Bleich, 2013). For the convergence of MPBART, we check the average tree depth for every latent utility W_{ikl} in (7). Average tree depths for all the three latent utilities are stable within a small range of fluctuation around 2.1, showing adequate stability and sufficient regularization from the tree priors that encourage shallow trees.

Figure 3 shows model validation results for time-varying confounders and outcomes. Average values of observed responses are represented by blue dots. The grey area represents 95% posterior credible intervals. To supplement the visual check, Table 3 summarizes

posterior outcome agreement,

$$\frac{1}{N'N_{post}}\sum_{i=1}^{N'}\sum_{j=1}^{N_{post}}\mathbb{1}\{\hat{S}_{ijk}=S_{ik}\},$$

and the posterior mode accuracy $\frac{1}{N'}\sum_{i=1}^{N'}\mathbb{1}\{\tilde{S}_{ik}=S_{ik}\}$, where

$$\tilde{S}_{ik}=\text{argmax}_{l\in\{0,1,2,3\}}\sum_{j=1}^{N_{post}}\mathbb{1}\{\hat{S}_{ijk}=l\}$$

of the outcome predictions on the validation data across time intervals. From the out-of-sample accuracies in Table 3, we conclude that all estimated predictive models have reasonably well out-of-sample posterior predictive accuracy across time.

4.2 Posterior Predictive Inference about Causal Effects

Using the validated models, we simulate regime-specific potential longitudinal paths of time-varying states and confounders, under the three treatment initiation policies: treat immediately, treat when CD4 drops below 350 cells/mm^3, and never treat. As described in Section 3.1, for a regime q, (x_k^*, s_k^*) are counterfactuals at interval k under an intervention path $\bar{a}_{k-1}^* = (a_0^*, \ldots, a_{k-1}^*)$ sequentially determined by $a_j^* = q(\mathcal{F}_j^*)$, $j = 0, \ldots, k-1$, where $\mathcal{F}_j^*$ is the counterfactual accumulated history before any intervention decision at time interval j. The three regimes of interest then correspond to $a_j^* \equiv 1$ (treat immediately), $a_j^* \equiv 0$ (never treat), and

$$a_j^* = q(a_{j-1}^*, x_{j-1}^*, s_j^*) = \begin{cases} \mathbb{1}\{r_{-1}^*=1, l_{-1}^* \le \log(351)\} & j=0 \\ \mathbb{1}\{l_{j-1}^* \le \log(351), s_j^*=1\} & j>1, a_{j-1}^*=0 \\ 1 & j>1, a_{j-1}^*=1, \end{cases}$$

(treat when CD4 falls below 350), where (r_j^*, l_j^*) are simulated values of (R_j^q, L_j^q).

We base the posterior predictive inference on the bootstrapped empirical distribution of the baseline variables (V, X_{-1}, X_0, S_0). In the sequential sampling scheme, letting i, j, and

k index person, posterior draw, and time interval, (a) if a s^*_{ijk} is drawn to be transfer-out or reported death, outcomes in the corresponding longitudinal path after interval k stays in the absorbed state; (b) for a path indicating that a patient is still active in care at interval k, the most recent CD4 stays unchanged unless the posterior simulation indicates there is a measurement update; and (c) simulation of s^*_{ijk} is done conditioning on the simulated value of $s^*_{ij,k-1}$. Figure 4 summarizes the progression of counterfactual engagement outcomes over time by showing the posterior mean of marginal outcome distributions $p_k^{(l)}$ as defined in equation (11), for $k = 1, \ldots, 10$ and $l = 0, 1, 2, 3$. From Figure 4 we can see that the immediate initiation of treatment yields superior patient engagement over time compared to the other two policies.

Figure 5 illustrates the progression of counterfactual engagement status through the care cascade in terms of time-specific between-state transition probabilities. Once disengaged, probability of continued disengagement tends to be high, regardless of treatment regime. However, among those currently engaged, treatment regime can greatly affect retention for the next time interval; early treatment initiation improves engagement at approximately one year in HIV care and reduces disengagement throughout the simulation time window. Note that in the observed data, the highest risk for disengagement is within one year of enrollment, as shown in Figure 6.

4.3 Causal Comparison of Treatment Initiation Policies

For the comparison of any two regimes q_1 and q_2, at each time interval $k \in \{1, \ldots, K\}$ and outcome level $l \in \{0, 1, 2, 3\}$, we obtain from Bayesian GCA the posterior predictive samples of $P(S_k^{q_1} = l|\overline{\mathcal{D}}_k)$ and $P(S_k^{q_2} = l|\overline{\mathcal{D}}_k)$, denoted by $\{p^{(1)}_{ljk}; j = 1, \ldots, N_{post}\}$ and $\{p^{(2)}_{ljk}; j = 1, \ldots, N_{post}\}$, respectively.

Causal comparativeness effectiveness between the two policies is quantified as the difference $P(S_k^{q_1} = l|\overline{\mathcal{D}}_k) - P(S_k^{q_2} = l|\overline{\mathcal{D}}_k)$, whose posterior predictive distribution can then be sampled as $\{d_{ljk}; j = 1, \ldots, N_{post}\}$, where $d_{ljk} = p_{ljk}^{(q_1)} - p_{ljk}^{(q_2)}$. Figure 7 shows the causal comparative effectiveness of two treatment initiation policies: treat immediately and treat when CD4 drops below 350 cells/mm^3; it displays the posterior mean and 95% posterior confidence interval of $\{d_{ljk}; j = 1, \ldots, N_{post}\}$ over $k = 1, \ldots, K$, stratified by outcome levels $l = 0, 1, 2, 3$.

In Figures 4 and 7, we observe that starting from the second interval, treat immediately significantly enhances engagement in care and reduces disengagement in the long run, which potentially decreases the under-reporting death or transfer-out due to loss to follow-up. Figure 7 also shows that the posterior predictive uncertainty increases over time because there is less data further away from enrollment. Specifically, Figure 5 illustrates that compared to the never treat and treat based on CD4, the treat immediately regime is the best at retaining engaged patients and re-engaging disengaged patients over time. In Table 4, the observed data in the first time interval has a lower engagement rate for people with treatment initiated at baseline; this is also reflected in the treatment comparison results from Figures 4 and 7. A possible explanation is that, patients who did not have treatment initiated during the baseline visits may be more concerned about starting therapy, thus having a higher tendency to re-visit during the first 200 days. The trend of a mild increasing transfer-outs in Figure 7 may have resulted from the higher engagement rate under immediate treatment, because patients have to be engaged in care in order to file a transfer-out to the care program, otherwise classified as disengagement.

5 Discussion

The main contribution of this work is to develop a Bayesian framework for the causal inference of multistate outcomes under dynamic treatment regimes, by using BART in the Bayesian GCA for the time-evolving generative components. Compared to parametric GCA, our proposed method uses additive regression trees as generative components, reducing concerns for model mis-specifications and allowing for flexible model specification. The proposed framework combines multiple machine-learning-based component models across time and enables causal comparisons from simulated potential longitudinal trajectories. There are two important assumptions: no unmeasured confounding and the first order Markov assumption. The former is untestable (Robins et al., 1999); the latter is testable in principle but is difficult to assess in practice. In addition, we programmed an R package "GcompBART" to support the implementation of the proposed method.

Our work demonstrates a connection between mathematical modeling and statistical modeling in causal inference. Mathematical modeling (Granich et al., 2009; Hontelez et al., 2011), such as microsimulation models and agent-based models, simulates individual outcomes for population-level inference. It derives a series of distributional assumptions from literature across multiple populations before calibrating model parameters to match the target population. The model parameters usually have causal interpretations that represent processes such as demography, natural history, intervention provision, and outcome progression. The advantage of mathematical modeling is its flexibility to describe a highly complex dynamic system; however, the strong assumptions carry the risk of losing coherence and liability for representing the specific population of interest, which may eventually induce biases in risk and causal estimations (Murray et al., 2017). Whereas in our method, inferences from statistical models make a relatively minimal set of assumptions and are more

data-dependent; the data used to fit the generative component models and the samples for which we generate counterfactual distributions should represent the target population.

The proposed approach has successfully identified substantial long-term positive effect of immediate ART at the initial clinic visits, such as enhancing patients' retention in HIV care. It is thought that better retention in HIV care will result in a decrease in HIV transmission (Skarbinski et al., 2015), higher rates of sustained suppression of viral load, and lower risk of drug resistance and mortality(Tripathi et al., 2011). It is worth noting that death and transfer-out are reported outcome states in this implementation, which means that disengaged individuals may have died or transferred out without informing the HIV care system (Bakoyannis and Yiannoutsos, 2015). Hence, reduced disengagement might lead to fewer unreported cases of death and transfer-out. Our proposal also provides a framework for future research on addressing such incomplete outcome ascertainment from under-reported or unobserved states. We can partially recover information about under-reported states by importing prior knowledge on the mechanisms underlying those categories, for example, using data from double-sampling designs on unobserved deaths (Bakoyannis et al., 2020).

Tables

Table 1: Observed state transition rates and marginal state probabilities over the follow-up period of 76,740 people (678,815 observations).

	S_k				
S_{k-1}	Disengaged	Engaged	Transferred Out	Died	n
Disengaged	0.936	0.061	0	0.002	214301
Engaged	0.129	0.851	0.007	0.013	387774
Transferred Out	0	0	1	0	2804
Died	0	0	0	1	5450

Table 2: Summary table of baseline variables V , time-varying confounders (X_{-1},X_0), and treatment indicator A_0 during the care initiation period. V are gender, age, marriage status, year of enrollment, travel time to clinic, WHO stage, weight, height, and viral load. X_{-1} is the latest CD4 before and up to the baseline, i.e. during days (-200, 0]. X_0 is the latest CD4 measurement during the HIV care initiation period of days (0,12].

		n	percent	25%	50%	75%
Gender	Female	50941	66.4			
	Male	25799	33.6			
Age		76740	100	28.5	35.2	43.2
Married	0	29668	38.7			
	1	40403	52.6			
	Missing	6669	8.7			

		n	percent	25%	50%	75%
Year of Enrollment	2008	7189	9.4			
	2009	13616	17.7			
	2010	12747	16.6			
	2011	11762	15.3			
	2012	9302	12.1			
	2013	7512	9.8			
	2014	7547	9.8			
	2015	6488	8.5			
	2016	577	0.8			
Travel Time	$>$ 2 h	5764	7.5			
	1-2 h	10835	14.1			
	30-60 min	20202	26.3			
	$<$ 30 min	17975	23.4			
	Missing	21964	28.6			
WHO stage	1	4111	5.4			
	2	1502	2.0			
	3	1751	2.3			
	4	287	0.4			
	Missing	69089	90.0			
Weight		73412	95.7	50.0	56.0	63.0
Height		63690	83.0	159.0	165.0	171.0
Viral Load		639	0.8	0	519	33129
CD4 in (-200,0]		48923	63.8	107.0	253.0	442.0
CD4 in (0,12]		6424	8.4	97.0	249.0	441.8
on ARV in [0,12]	0	64247	83.7			
	1	12493	16.3			

Table 3: Posterior mean accuracy of outcome models across time intervals.

k	1	2	3	4	5	6	7	8	9	10
Posterior Mean Accuracy	0.81	0.71	0.79	0.83	0.84	0.85	0.86	0.86	0.87	0.87
Posterior Mode Accuracy	0.89	0.79	0.86	0.89	0.91	0.91	0.92	0.92	0.92	0.93

Table 4: Observed engagement rate and the total number of patients at day 200, stratified by A_0 and CD4 measurement status.

	CD4$\leq$350	CD4>350 or missing	All
$A_0 = 0$	0.899 (33042)	0.907 (31138)	0.903 (64180)
$A_0 = 1$	0.818 (4317)	0.732 (8159)	0.762 (12476)

Figures

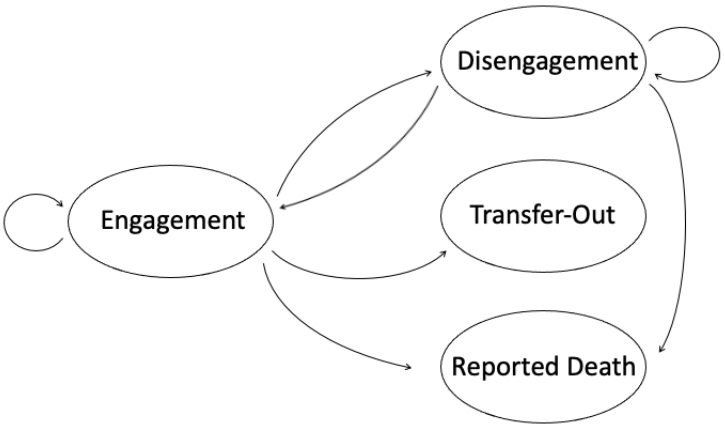

Figure 1: Multinomial state transitions in the progression of patients' engagement in care through the HIV care cascade.

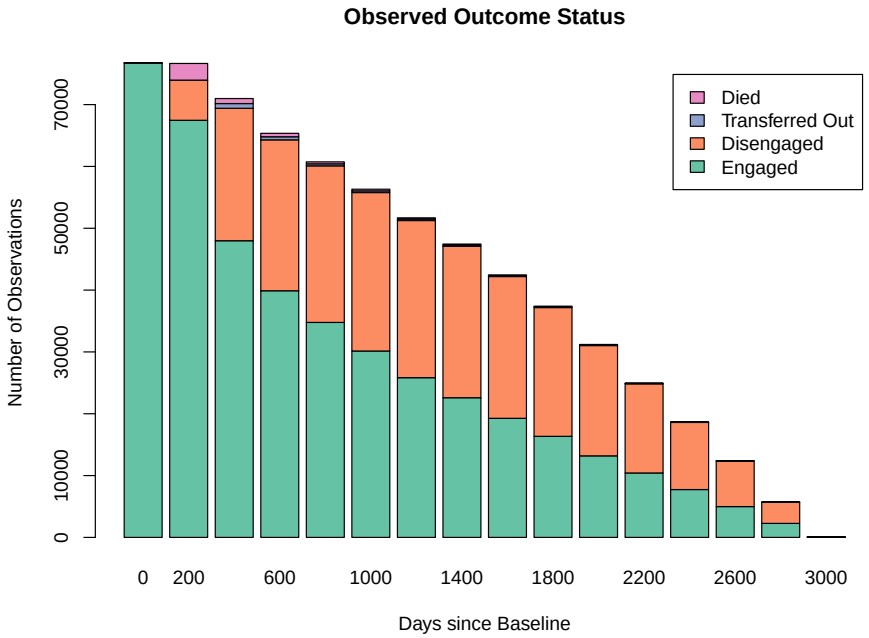

Figure 2: Number of subjects in each state over time. At baseline, 92,050 unique individual records are available from AMPATH after data processing. By definition, all individuals are engaged in care at baseline.

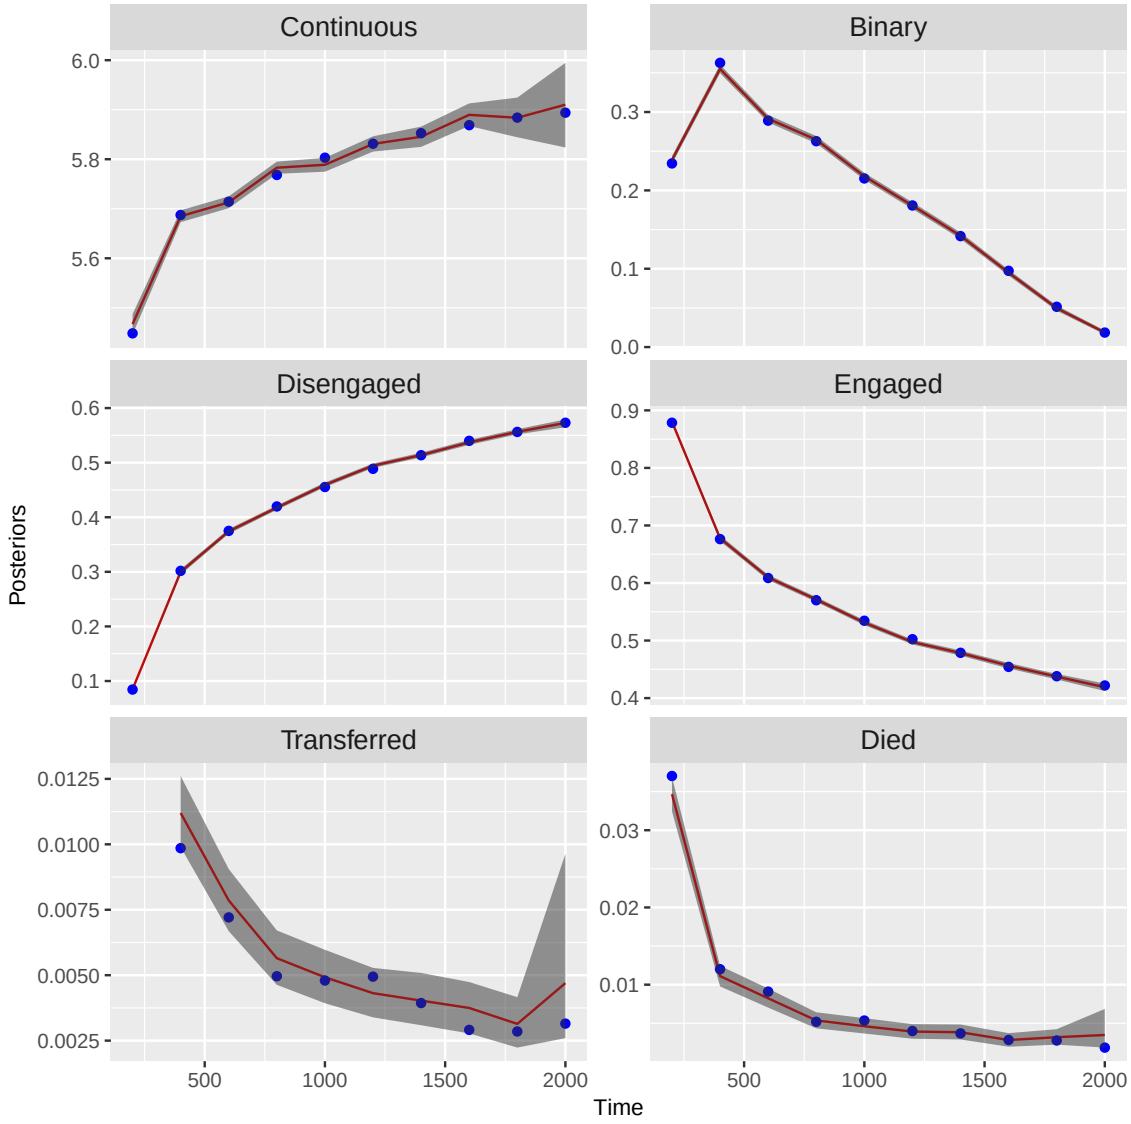

Figure 3: Accuracy plots for validating models of the time-varying confounders and multinomial outcome on validation data of 26,740 patients, across time intervals $t = 1, \ldots, 10$. For the first row, the left and right plots correspond to L_t (continuous) and R_t (binary), respectively. For L_t at the tth interval, blue dot represents observed response mean $\frac{1}{N_t'} \sum_{i \in I_t} L_{it}$, red line is the mean and grey band is the 95% confidence interval of posterior predictive samples $\{\frac{1}{N_t'} \sum_{i \in I_t} \hat{L}_{ijt}; j = 1, \ldots, N_{post}\}$, where I_t is the subgroup of test set subjects who were at risk for an event at time t, N_t' is the sample size of I_t, and $\hat{L}_{ijt}$ is the jth posterior prediction of the continuous confounder L_t for the ith subject; similarly for R_t. The second and third rows correspond to the outcome being disengaged, engaged, transferred out, and died. For disengagement, the blue dot represents the observed marginal probability $\frac{1}{N_t'} \sum_{i \in I_t} \mathbb{1}\{S_{it} = \text{Diseng}\}$, red line is the mean and grey band as the 95% confidence interval of posterior predictive samples $\{\frac{1}{N_t'} \sum_{i \in I_t} \mathbb{1}\{\hat{S}_{ijt} = \text{Diseng}\}; j = 1, \ldots, N_{post}\}$; likely for engagement, transfer-out, and death.

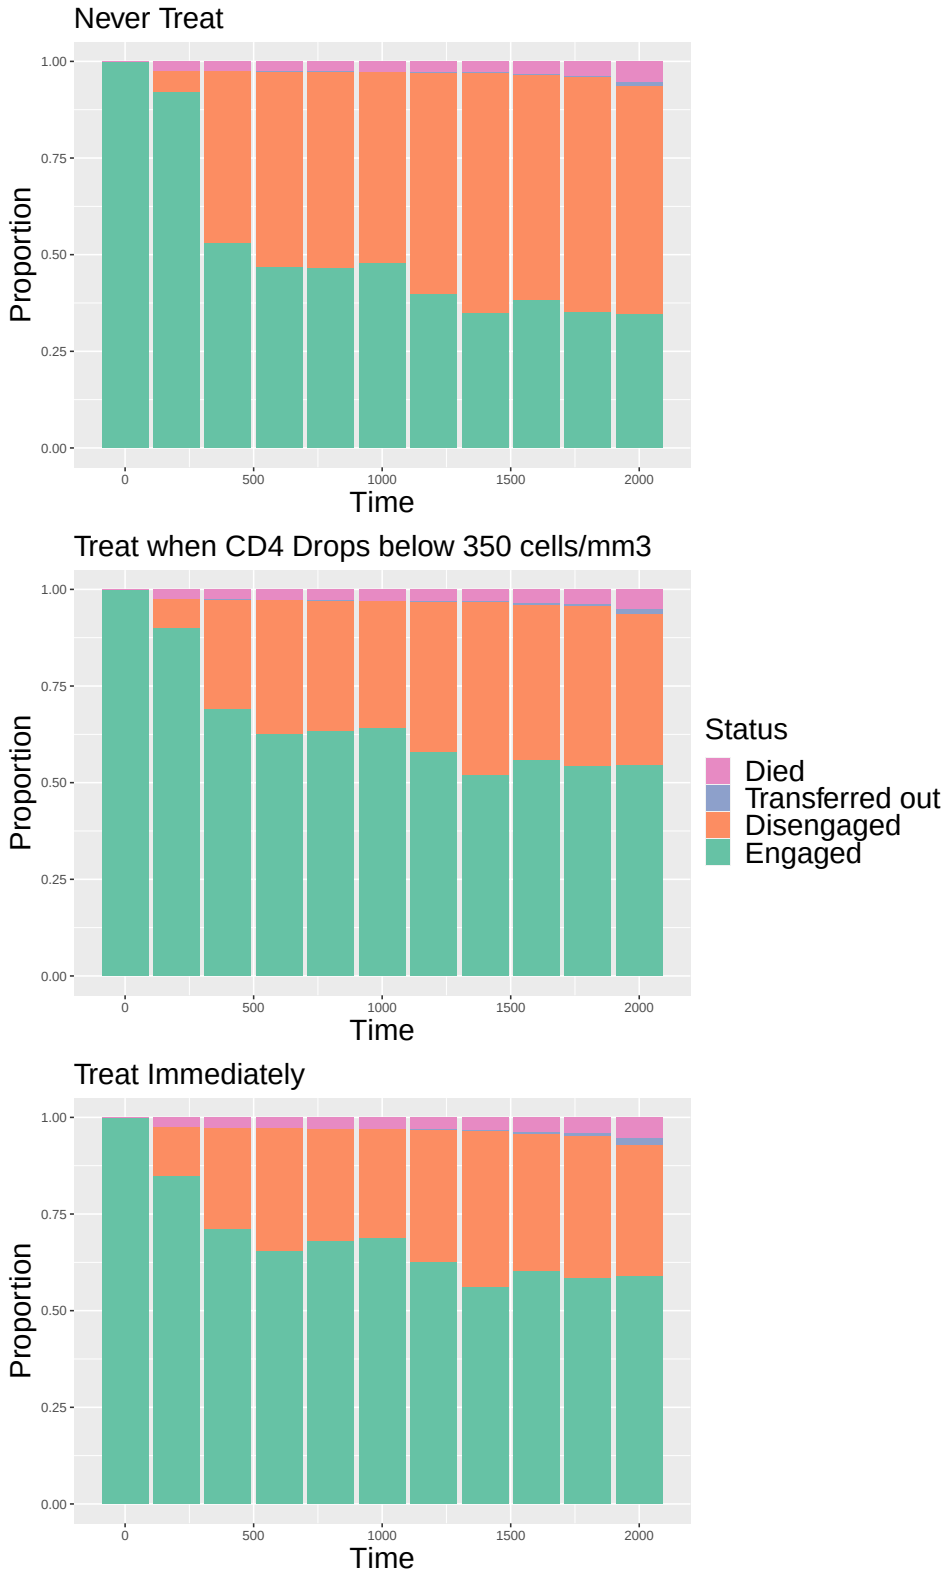

Figure 4: Predicted marginal state probabilities for an out-of-sample 30,000 individuals engaged in AMPATH-supported HIV care at baseline, under the never treat, treat when CD4 drops below 350 cells/mm^3, and treat immediately policies (in the order of display, left to right, up to bottom).

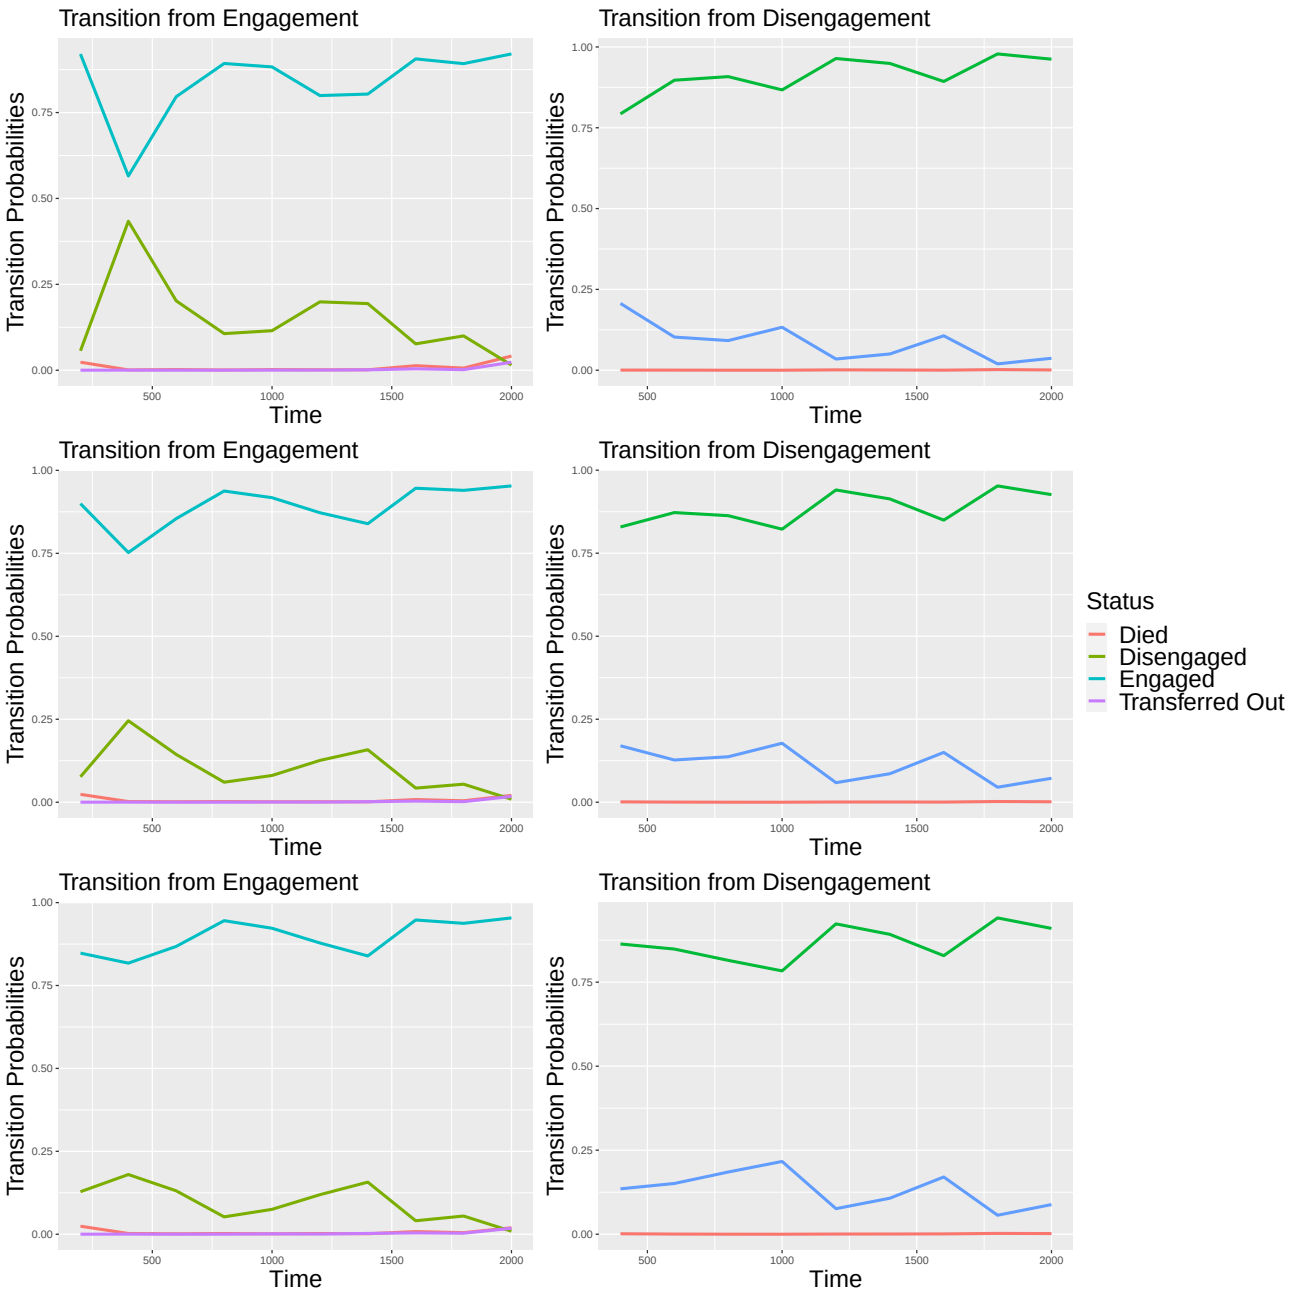

Figure 5: Counterfactual transition probabilities for an 10,000 individuals engaged in AMPATH-supported HIV care at baseline, under the never treat, treat when CD4 drops below 350 cells/mm^3, and treat immediately regimes (in the order of display, up to bottom).

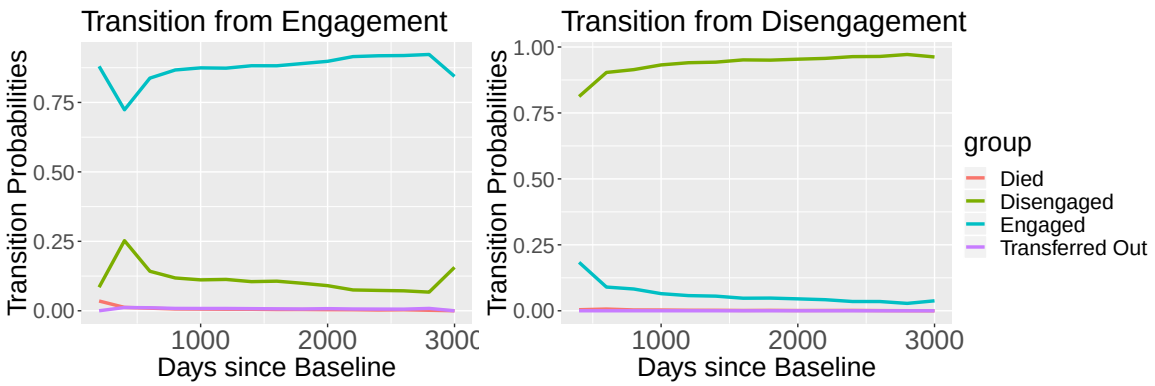

Figure 6: Plot of unadjusted transition probabilities. Plot on the left depicts observed probability of state transition from engagement calculated every 200 days since the first primary encounter and plot on the right shows that from disengagement.

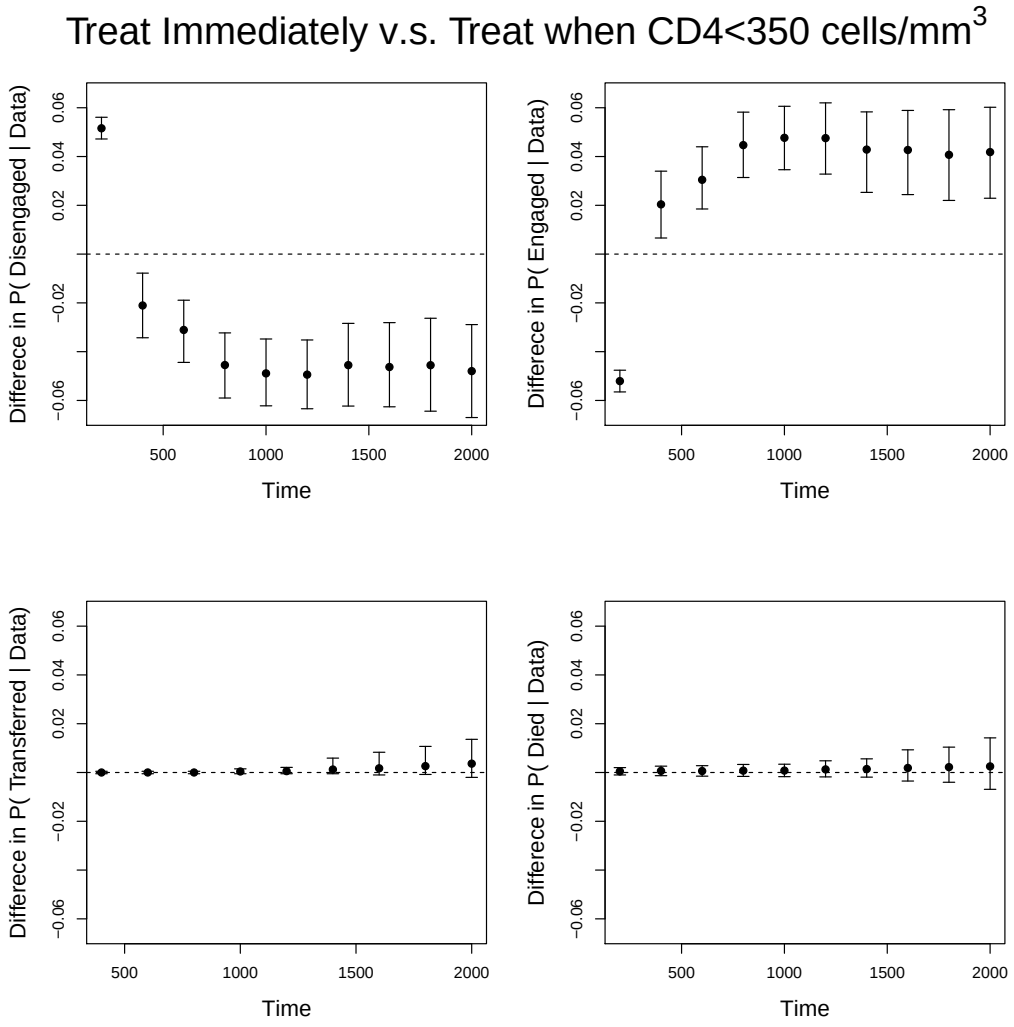

Figure 7: Difference in percentage of engagement statuses between treat immediately and the treat when CD4 drops below 350 cells/mm^3, calculated based on counterfactual trajectories for 10,000 individuals engaged in AMPATH-supported HIV care at baseline.

References

Antonelli, J. and Daniels, M. J. (2019). Discussion of PENCOMP. *Journal of the American Statistical Association*, 114:24–27.

Bakoyannis, G., Diero, L., Mwangi, A., Wools-Kaloustian, K. K., and Yiannoutsos, C. T. (2020). A semiparametric method for the analysis of outcomes during a gap in hiv

care under incomplete outcome ascertainment. *Statistical communications in infectious diseases*, 12(s1).

Bakoyannis, G. and Yiannoutsos, C. T. (2015). Impact of and Correction for Outcome Misclassification in Cumulative Incidence Estimation. *PloS one*, 10(9):e0137454.

Burgette, L. F. and Nordheim, E. V. (2012). The Trace Restriction: An Alternative Identification Strategy for the Bayesian Multinomial Probit Model. *Journal of Business & Economic Statistics*, 30(3):404–410.

Chipman, H. A., George, E. I., and McCulloch, R. E. (2010). BART: Bayesian additive regression trees. *The Annals of Applied Statistics*, 4(1):266–298.

Gardner, E. M., McLees, M. P., Steiner, J. F., Del Rio, C., and Burman, W. J. (2011). The spectrum of engagement in HIV care and its relevance to test-and-treat strategies for prevention of HIV infection. *Clinical Infectious Diseases: An Official Publication of the Infectious Diseases Society of America*, 52(6):793–800.

Granich, R. M., Gilks, C. F., Dye, C., De Cock, K. M., and Williams, B. G. (2009). Universal voluntary HIV testing with immediate antiretroviral therapy as a strategy for elimination of HIV transmission: a mathematical model. *The Lancet*, 373(9657):48–57.

Hontelez, J. A. C., Vlas, S. J. d., Tanser, F., Bakker, R., Bärnighausen, T., Newell, M.-L., Baltussen, R., and Lurie, M. N. (2011). The Impact of the New WHO Antiretroviral Treatment Guidelines on HIV Epidemic Dynamics and Cost in South Africa. *PLOS ONE*, 6(7):e21919.

Horstmann, E., Brown, J., Islam, F., Buck, J., and Agins, B. D. (2010). Retaining HIV-

Infected Patients in Care: Where Are We? Where Do We Go from Here? *Clinical Infectious Diseases*, 50(5):752–761.

Insight Start Study Group (2015). Initiation of antiretroviral therapy in early asymptomatic hiv infection. *New England Journal of Medicine*, 373(9):795–807.

Josefsson, M. and Daniels, M. J. (2019). Bayesian semi-parametric G-computation for causal inference in a cohort study with non-ignorable dropout and death. *arXiv:1902.10787 [stat]*. arXiv: 1902.10787.

Kapelner, A. and Bleich, J. (2013). bartMachine: Machine Learning with Bayesian Additive Regression Trees. *arXiv:1312.2171 [cs, stat]*. arXiv: 1312.2171.

Kay, E. S., Batey, D. S., and Mugavero, M. J. (2016). The HIV treatment cascade and care continuum: updates, goals, and recommendations for the future. *AIDS Research and Therapy*, 13.

Keil, A. P., Daza, E. J., Engel, S. M., Buckley, J. P., and Edwards, J. K. (2017). A Bayesian approach to the g-formula. *Statistical Methods in Medical Research*. arXiv: 1512.04809.

Lee, H., Genberg, B. L., Nyambura, M., Hogan, J., Braitstein, P., and Sang, E. (2017). State-Space Models for Engagement, Retention, and Reentry in the HIV Care Cascade. *CROI Conference*.

Lemnaru, C. and Potolea, R. (2011). Imbalanced Classification Problems: Systematic Study, Issues and Best Practices. In *Enterprise Information Systems*, pages 35–50. Springer, Berlin, Heidelberg.

Mugglin, C., Kläger, D., Gueler, A., Vanobberghen, F., Rice, B., Egger, M., and epidemiology Databases to Evaluate AIDS in Southern Africa, I. (2021). The hiv care cascade

in sub-saharan africa: systematic review of published criteria and definitions. *Journal of the International AIDS Society*, 24(7):e25761.

Murray, E. J., Robins, J. M., Seage, G. R., Freedberg, K. A., and Hernán, M. A. (2017). A Comparison of Agent-Based Models and the Parametric G-Formula for Causal Inference. *American Journal of Epidemiology*, 186(2):131–142.

Organization, W. H. et al. (2015). *Guideline on when to start antiretroviral therapy and on pre-exposure prophylaxis for HIV*. World Health Organization.

Perlman, D. C., Jordan, A. E., and Nash, D. (2017). Conceptualizing Care Continua: Lessons from HIV, Hepatitis C Virus, Tuberculosis and Implications for the Development of Improved Care and Prevention Continua. *Frontiers in Public Health*, 4.

Petersen, M. L. and van der Laan, M. J. (2011). Case Study: Longitudinal HIV Cohort Data. In *Targeted Learning*, Springer Series in Statistics, pages 397–417. Springer, New York, NY.

Robins, J. (1986). A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect. *Mathematical Modelling*, 7(9):1393–1512.

Robins, J. M., Greenland, S., and Hu, F.-C. (1999). Estimation of the Causal Effect of a Time-Varying Exposure on the Marginal Mean of a Repeated Binary Outcome. *Journal of the American Statistical Association*, 94(447):687–700.

Skarbinski, J., Rosenberg, E., Paz-Bailey, G., Hall, H. I., Rose, C. E., Viall, A. H., Fagan, J. L., Lansky, A., and Mermin, J. H. (2015). Human immunodeficiency virus transmission

at each step of the care continuum in the United States. *JAMA internal medicine*, 175(4):588–596.

The INSIGHT START Study Group (2015). Initiation of Antiretroviral Therapy in Early Asymptomatic HIV Infection. *New England Journal of Medicine*, 373(9):795–807.

Tierney, W. M., Nyandiko, W. N., Siika, A. M., Wools-Kaloustian, K., Sidle, J. E., Kiplagat, J., Bell, A., and Inui, T. S. (2013). "these are good problems to have...": establishing a collaborative research partnership in east africa. *Journal of general internal medicine*, 28(3):625–638.

Tripathi, A., Youmans, E., Gibson, J. J., and Duffus, W. A. (2011). The impact of retention in early HIV medical care on viro-immunological parameters and survival: a statewide study. *AIDS research and human retroviruses*, 27(7):751–758.

UNAIDS (2014). Fast-track – ending the AIDS epidemic by 2030.

van Dyk, D. A. (2010). MARGINAL MARKOV CHAIN MONTE CARLO METHODS. *Statistica Sinica*.

Vourli, G., Noori, T., Pharris, A., Porter, K., Axelsson, M., Begovac, J., Cazein, F., Costagliola, D., Cowan, S., Croxford, S., et al. (2020). Human immunodeficiency virus continuum of care in 11 european union countries at the end of 2016 overall and by key population: Have we made progress? *Clinical Infectious Diseases*, 71(11):2905–2916.

WHO (2012). Meeting report on Framework for metrics to support effective treatment as prevention.

Xu, Y., Hogan, J. W., Daniels, M. J., Kantor, R., and Mwangi, A. (2021). Inference for BART with Multinomial Outcomes. *arXiv:2101.06823 [cs, stat]*. arXiv: 2101.06823.

Young, J. G., Cain, L. E., Robins, J. M., O'Reilly, E. J., and Hernán, M. A. (2011). Comparative Effectiveness of Dynamic Treatment Regimes: An Application of the Parametric G-Formula. *Statistics in Biosciences*, 3(1):119.

Zhou, T., Elliott, M. R., and Little, R. J. A. (2019). Penalized Spline of Propensity Methods for Treatment Comparison. *Journal of the American Statistical Association*, 114(525):1–38.