

RESEARCH ARTICLE OPEN ACCESS

Causal Inference for Continuous Multiple Time Point Interventions

Michael Schomaker^{1,2,3} | Helen McIlleron^{4,5} | Paolo Denti⁴ | Iván Díaz⁶

¹Department of Statistics, Ludwig-Maximilians University, Munich, Germany | ²Centre for Infectious Disease Epidemiology and Research, University of Cape Town, Cape Town, South Africa | ³Institute of Public Health, Medical Decision Making and Health Technology Assessment, UMIT—University for Health Sciences, Medical Informatics and Technology, Hall in Tirol, Austria | ⁴Division of Clinical Pharmacology, Department of Medicine, University of Cape Town, Cape Town, South Africa | ⁵Wellcome Centre for Infectious Diseases Research in Africa (CIDRI-Africa), Institute of Infectious Disease and Molecular Medicine, University of Cape Town, Cape Town, South Africa | ⁶Division of Biostatistics, Department of Population Health, New York University Grossman School of Medicine, New York, New York, USA

Correspondence: Michael Schomaker (michael.schomaker@stat.uni-muenchen.de)

Received: 15 October 2023 | **Revised:** 20 September 2024 | **Accepted:** 29 September 2024

Funding: Michael Schomaker is supported by the German Research Foundations (DFG) Heisenberg Programm (grants 465412241 and 465412441).

ABSTRACT

There are limited options to estimate the treatment effects of variables which are continuous and measured at multiple time points, particularly if the true dose–response curve should be estimated as closely as possible. However, these situations may be of relevance: in pharmacology, one may be interested in how outcomes of people living with—and treated for—HIV, such as viral failure, would vary for time-varying interventions such as different drug concentration trajectories. A challenge for doing causal inference with continuous interventions is that the positivity assumption is typically violated. To address positivity violations, we develop projection functions, which reweigh and redefine the estimand of interest based on functions of the conditional support for the respective interventions. With these functions, we obtain the desired dose–response curve in areas of enough support, and otherwise a meaningful estimand that does not require the positivity assumption. We develop *g*-computation type plug-in estimators for this case. Those are contrasted with *g*-computation estimators which are applied to continuous interventions without specifically addressing positivity violations, which we propose to be presented with diagnostics. The ideas are illustrated with longitudinal data from HIV positive children treated with an efavirenz-based regimen as part of the CHAPAS-3 trial, which enrolled children < 13 years in Zambia/Uganda. Simulations show in which situations a standard *g*-computation approach is appropriate, and in which it leads to bias and how the proposed weighted estimation approach then recovers the alternative estimand of interest.

1 | Introduction

Causal inference for multiple time-point interventions has received considerable attention in the literature over the past few years: if the intervention of interest is binary, popular estimation approaches include inverse probability of treatment weighting approaches (IPTW) [1], *g*-computation estimators [2, 3]

and longitudinal targeted maximum likelihood estimators (LTMLE) [4], among others. For IPTW, the treatment and censoring mechanisms need to be estimated at each time point, parametric *g*-computation requires fitting of both the outcome and confounder mechanisms, sequential *g*-computation is based on the iterated outcome regressions only, and LTMLE needs models for the outcome, censoring, and treatment mechanisms,

This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial](https://creativecommons.org/licenses/by-nc/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited and is not used for commercial purposes.

© 2024 The Author(s). *Statistics in Medicine* published by John Wiley & Sons Ltd.

iteratively at each point in time. All of those methods are relatively well-understood and have been successfully applied in different fields (e.g., in [5–10]).

However, suggestions on how to estimate treatment effects of variables that are continuous and measured at multiple time points are limited. This may, however, be of interest to construct causal dose–response curves (CDRC). For example, in pharmacoepidemiology, one may be interested in how counterfactual outcomes vary for different dosing strategies for a particular drug, see Section 2.

Early work on continuous interventions includes the seminal paper of Robins, Hernán, and Brumback [1] on inverse probability weighting of marginal structural models (MSMs). For a single time point, this requires the estimation of stabilized weights based on the conditional density of the treatment, given the confounders. This density may be estimated with parametric regression models, such as linear regression models. The MSM, describing the dose–response relationship (i.e., the CDRC), can then be obtained using weighted regression. The suggested approach can also be used for the longitudinal case, where stabilized weights can be constructed easily, and one may work with a working model where, for example, the effect of cumulative dose over all time points on the response is estimated.

There are several other suggestions for the point treatment case (i.e., a single time point): for example, the use of the generalized propensity score (GPS) is often advocated in the literature [11]. Similar to the MSM approach described above, both the conditional density of the treatment, given the confounders, and the dose–response relationship have to be specified and estimated. Instead of using stabilized weights, the GPS is included in the dose–response model as a covariate. To reduce the risk of bias due to model mis-specification, it has been suggested to incorporate machine learning (ML) in the estimation process [12]. However, when combining GPS and ML approaches, there is no guarantee for valid inference, that is, related confidence intervals may not achieve nominal coverage [4]. Doubly robust (DR) approaches allow the integration of ML while retaining valid inference. As such, the DR-estimator proposed by Kennedy et al. [13] is a viable alternative to both MSM and GPS approaches. It does not rely on the correct specification of parametric models and can incorporate general machine learning approaches. As with other, similar, approaches [14] both the treatment and outcome mechanisms need to be modeled. Then a pseudo-outcome based on those two estimated nuisance functions is constructed, and put into relationship with the continuous intervention using kernel smoothing. Further angles in the point treatment case are given in the literature [15–18].

There are fewer suggestions as to how to estimate CDRCs for multiple time point interventions. As indicated above, one could work with inverse probability weighting of marginal structural models and specify a parametric dose–response curve. Alternatively, one may favor the definition of the causal parameter as a projection of the true CDRC onto a specified working model [19]. Both approaches have the disadvantage that, with mis-specification of the dose–response relationship or an inappropriate working model, the postulated curve may be far

away from the true CDRC. Moreover, the threat of practical positivity violations is even more severe in the longitudinal setup, and working with inverse densities—which can be volatile—remains a serious concern [20]. It has thus been suggested in the literature to avoid these issues by changing the scientific question of interest, if meaningful, and to work with alternative definitions of causal effects.

For example, Young, Hernan, and Robins [21] consider so-called modified treatment policies, where treatment effects are allowed to be stochastic and depend on the natural value of treatment, defined as the treatment value that would have been observed at time t , had the intervention been discontinued right before t . A similar approach relates to using representative interventions, which are stochastic interventions that maintain a continuous treatment within a pre-specified range [22]. Diaz et al. [23] present four estimators for longitudinal modified treatment policies (LMTs), based on IPTW, g-computation and doubly robust considerations. These estimators are implemented in an R-package (lmt). An advantage of those approaches is that the positivity assumption can sometimes be relaxed, depending on how interventions are being designed. Moreover, the proposed framework is very general, applicable to longitudinal and survival settings, and one is not forced to make arbitrary parametric assumptions if the doubly robust estimators are employed. It also avoids estimation of conditional densities as it recasts the problem at hand as a classification procedure for the density ratio of the post-intervention and natural treatment densities. A disadvantage of the approach is that it does not aim at estimating the CDRC, which may, however, relate to the research question of interest, see Section 2.

In this paper, we are interested in counterfactual outcomes after intervening on a continuous exposure (such as drug concentration) at multiple time points. For example, we may be interested in the probability of viral suppression after 1 year of follow-up, had HIV-positive children had a fixed concentration level of efavirenz during this year. In this example, it would be desirable to estimate the true CDRC as closely as possible, to understand and visualize the underlying biological mechanism and decide what preferred target concentrations should be. This comes with several challenges: most importantly, violations of the positivity assumption are to be expected with continuous multiple time-point interventions. If this is the case, estimands that are related to modified treatment policies or stochastic interventions, as discussed above, can be tailored to tackle the positivity violation problem, but redefine the question of interest. In pharmacoepidemiology and other fields this may be not ideal, as interpretations of the true CDRC are of considerable clinical interest, for example to determine appropriate drug target concentrations.

We propose g-computation based approaches to estimate and visualize causal dose–response curves. This is an obvious suggestion in our setting because developing a standard doubly robust estimator, for example, a targeted maximum likelihood estimator, is *not* possible as the CDRC is not a pathwise-differentiable parameter; and developing non-standard doubly robust estimators is not straightforward in a multiple time-point setting. Our suggested approach has two angles:

- i. As a first step, we simply consider computing counterfactual outcomes for multiple values of the continuous intervention (at each time point) using standard parametric and sequential g-computation. We evaluate with simulation studies where this “naive” estimation strategy can be successful and useful, and where not; and which diagnostics may be helpful in judging its reliability. To our knowledge, this standard approach has not been evaluated in the literature yet.
- ii. We then define regions of low support through low values of the conditional treatment density, evaluated at the intervention trajectories of interest. Our proposal is to redefine the estimand of interest (i.e., the CDRC) only in those regions, based on suitable weight functions. Such an approach entails a compromise between identifiability and interpretability: it is a tradeoff between estimating the CDRC as closely as possible [as in (i)], at the risk of bias due to positivity violations and minimizing the risk of bias due to positivity violations, at the cost of redefining the estimand in regions of low support. We develop a g-computation based plug-in estimator for this weighted approach.

We introduce the motivating question in Section 2, followed by the theoretical framework in Section 3. After presenting our extensive Monte-Carlo simulations (Section 4), we analyze the illustrative data example in Section 4.3. We conclude in Section 5.

2 | Motivating Example

Our data comes from *CHAPAS-3*, an open-label, parallel-group, randomized trial (*CHAPAS-3*) [24]. Children with HIV (aged 1 month to 13 years), from four different treatment centers (one in Zambia, three in Uganda) were randomized to receive one out of 3 different antiretroviral therapy (ART) regimens, given as fixed-dose-combination tablets. Each regimen consisted of 3 drugs. Every child received lamivudine (first drug), and either nevirapine or efavirenz (second drug), which was chosen at the discretion of the treating physician (and based on age). The third drug was randomly assigned (1:1:1) and either stavudine, zidovudine, or abacavir. The primary endpoint in the trial were adverse events, both clinical (grade 2/3/4) and laboratory (confirmed grade 3, or any grade 4).

Several substudies of the trial explored pharmacokinetic aspects of the nonnucleoside reverse transcriptase inhibitors (NNRTI) efavirenz and nevirapine (the second, non-randomized drug), in particular the relationship between the respective drug concentrations and elevated viral load, that is, viral failure [25, 26]. Our analyses are motivated by these substudies and are based on the subsets of children, who received efavirenz (i.e., 125 out of 478 patients). We use the same data as Bienczak et al. [25].

In the trial, efavirenz dose was recommended to be based on weight using 200 mg for those weighing 10–13.9 kg, 300 mg for 14–19.9 kg, 400 mg for 20–34.9 kg, and 600 mg (the adult dose) above 34.9 kg. While children may receive the same dose, EFV concentrations vary individually and depend on the child’s

metabolism (i.e., specifically the single nucleotide polymorphisms in the CYP2B6 gene encoding the key metabolizing enzyme). Too low concentrations reduce antiviral activity of the drug and thus lead to viral failure (a negative outcome, which typically leads to a change in drug regimen). This is the reason why one is interested in the minimum concentration that is still effective against viral replication; or, more generally in the range of concentrations that should be targeted.

Our analysis evaluates the relationship between EFV concentrations (not doses) and viral failure, over a follow-up period of 84 weeks. We are specifically interested in the counterfactual probability of viral load (VL) > 100 copies/ml at 84 weeks if children had concentrations (12/24 h after dose) of x mg/L at each follow-up visit, where x ranges from 0 to 6 mg/L. This question translates into the longitudinal causal dose–response curve (CDRC), which –in our example–, is actually a concentration-response curve. That is, we want to know how the probability of failure varies for different concentration trajectories.

However, drawing a particular CDRC is challenging for the following reasons: (i) the form of the curve should be flexible and as close as possible to the truth; (ii) there exist time-dependent confounders (e.g., adherence, weight) that are themselves affected by prior concentration levels, making regression an invalid method for causal effect estimation [27]; (iii) with long follow-up and moderate sample size, and given the continuous nature of the concentration variables, positivity violations are an issue of concern that have to be addressed.

More details on the data analysis are given in Section 4.3.

3 | Framework

3.1 | Notation

We consider a longitudinal data setup where at each time point t , $t = 0, 1, \dots, T$, we measure the outcome Y_t , a continuous intervention A_t and covariates L_t^j , $j = 1, \dots, q$ for $i = 1, \dots, n$ individuals. We denote $\mathbf{L}_0 = \{L_0^1, \dots, L_0^q\}$ as “baseline variables” and $\mathbf{L}_t = \{L_t^1, \dots, L_t^q\}$ as follow-up variables, with $q, q_0 \in \mathbb{N}$. The intervention and covariate histories of a unit i (up to and including time t) are $\bar{A}_{t,i} = (A_{0,i}, \dots, A_{t,i})$ and $\bar{L}_{t,i}^s = (L_{0,i}^s, \dots, L_{t,i}^s)$, $s = 1, \dots, q$, $i = 1, \dots, n$, respectively. The observed data structure is

$$O = (\mathbf{L}_0, A_0, Y_0, \mathbf{L}_1, A_1, Y_1, \dots, \mathbf{L}_T, A_T, Y_T)$$

That is, we consider n units, followed-up over T time points where at each time point we work with the order $L_t \rightarrow A_t \rightarrow Y_t$.

We are interested in the counterfactual outcome $Y_{t,i}^{\bar{a}_t}$ that would have been observed at time t if unit i had received, possibly contrary to the fact, the intervention history $\bar{A}_{t,i} = \bar{a}_t$, $a_{t,i} \in \mathbb{R}$. For a given intervention $\bar{A}_{t,i} = \bar{a}_t$, the counterfactual covariates are denoted as $\bar{L}_{t,i}^{\bar{a}_t}$. We use \mathbf{H}_t to denote the history of all data up to before A_t .

3.2 | Estimands

3.2.1 | Estimands for One Time Point

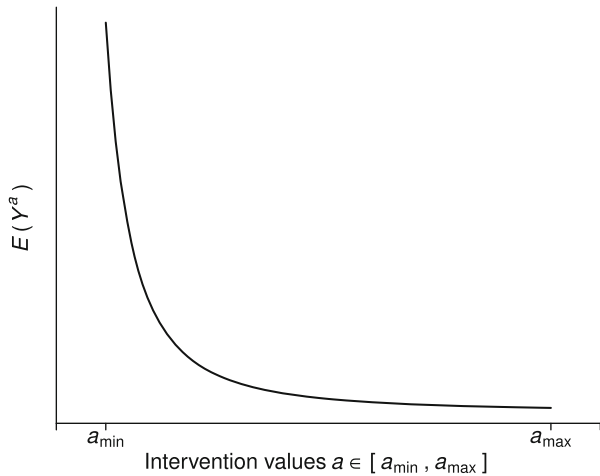
In order to illustrate the ideas, we first consider a simple example with only one time point where $O = (\mathbf{L}, A, Y)$. Suppose we observe an independent and identically distributed sample $O = (O_1, \dots, O_n)$, where O has support $\text{supp}(O) = \{\mathcal{L} \times \mathcal{A} \times \mathcal{Y}\}$. Our estimand of interest is the causal dose–response curve:

Estimand 1:

$$m : a \mapsto E(Y^a) \quad \forall a \in \mathcal{A} \quad (1)$$

where $\mathcal{A} = \text{supp}(A) := \{a \in \mathcal{A} : f(a) \neq 0\}$. Assume \mathbf{L} has density $p_0(\mathbf{l})$ with respect to some dominating measure $\nu(\mathbf{l})$. The dose–response curve (1) can be identified [13] as

$$\begin{aligned} m(a) &= E_{\mathbf{L}}[E(Y^a | A = a, \mathbf{L})] = E_{\mathbf{L}}[E(Y | A = a, \mathbf{L})] \\ &= \int_{\mathbf{l} \in \mathcal{L}} E(Y | A = a, \mathbf{L} = \mathbf{l}) p_0(\mathbf{l}) d\nu(\mathbf{l}) \end{aligned} \quad (2)$$

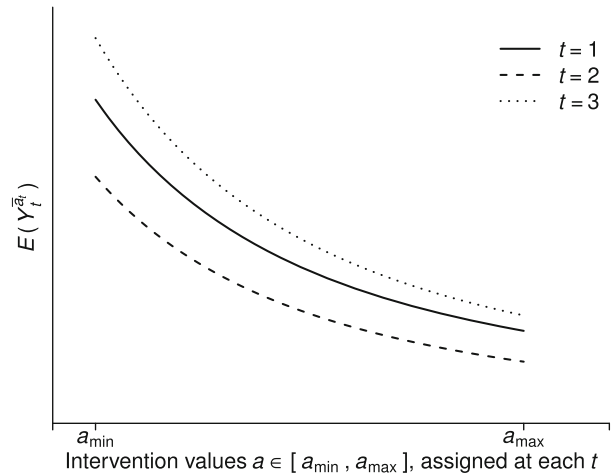


(a) CDRC for one time point

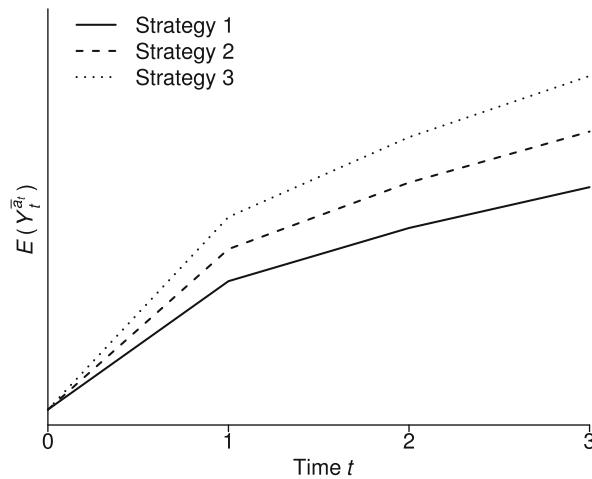
where the first equality follows by the law of iterated expectation, the assumption that Y^a is independent of A conditional on \mathbf{L} (conditional exchangeability) and positivity (see below). The second equality follows because $Y^a = Y$ in the event of $A = a$ (consistency). An example of such a curve for a particular interval $[a_{\min}, a_{\max}]$ is given in Figure 1a.

This quantity is undefined if there is an intervention level $a \in \mathcal{A}$ such that the conditional density function $g(a | \mathbf{l})$ is zero for some \mathbf{l} with $p_0(\mathbf{l}) > 0$. The *strong positivity assumption* reflects this consideration by requiring that $\inf_{a \in \mathcal{A}} g(a | \mathbf{l}) > 0$ almost everywhere [28]. This assumption may be weakened by making additional parametric model assumptions or restricting the range of a . In this paper, we however address violations of the positivity assumption through a redefinition of the estimand of interest, where we replace the marginal distribution function of \mathbf{L} in (2) by a user-given distribution function. That is, we instead target

$$m_w(a) = \int E(Y | A = a, \mathbf{L} = \mathbf{l}) w(a, \mathbf{l}) p_0(\mathbf{l}) d\nu(\mathbf{l}) \quad (3)$$



(b) Possible CDRC visualization
– if intervention values are constant over time



(c) Possible CDRC visualization
– if intervention values vary over time (or for survival data)

FIGURE 1 | Considerations for causal dose–response curves (CDRCs).

for some weight function $w(a, \mathbf{I})$. This is essentially a weighted average of the conditional dose–response curve $E(E(Y^a | \mathbf{L})w(a, \mathbf{L}))$.

To illustrate the construction of a meaningful weight function for (3), consider two extreme cases: $w(a, \mathbf{I}) = 1$, and $w(a, \mathbf{I}) = g(a | \mathbf{I})/g(a)$, where $g(a)$ is the marginal density of A . Under the first case, we have $m_w(a) = m(a)$, which is equal to the dose–response curve (Estimand 1) whenever the positivity assumption holds. Under $w(a, \mathbf{I}) = g(a | \mathbf{I})/g(a)$, we have $m_w(a) = E(Y | A = a)$, which is not a causal quantity (unless exchangeability is assumed) but it does not require the positivity assumption. We therefore propose to use a function $w(a, \mathbf{I})$ such that $w(a, \mathbf{I}) = 1$ in areas of $\{A, \mathcal{L}\}$ that have good support, and $w(a, \mathbf{I}) = g(a | \mathbf{I})/g(a)$ in areas of $\{A, \mathcal{L}\}$ that have low support. By “good support”, we mean that the positivity assumption is met in the sense of $g(a | \mathbf{I}) > c > 0$, that is, bounded away from zero. Thus, one possible function to use is

$$w(a, \mathbf{I}) = \begin{cases} 1 & \text{if } g(a | \mathbf{I}) > c \\ \frac{g(a | \mathbf{I})}{g(a)} & \text{otherwise} \end{cases} \quad (4)$$

Alternatively, one may use $g(a | \mathbf{I})/g(a) > c > 0$ to define good support. Using (4) in (3) targets the desired dose–response curve under enough support and avoids reliance on positivity otherwise. Formally,

Estimand 2 is then given by:

$$m_w(a) = E(Y^a | g(a | L) > c)P(g(a | L) > c) + E(Y^a w(a | L) | g(a | L) \leq c)P(g(a | L) \leq c) \quad \forall a \in \mathcal{A} \quad (5)$$

For fixed a , this is taking a weighted average of the counterfactuals Y^a , where the weights are equal to one for units in the population for whom positivity $g(a | L) > c$ holds, and the weights are $w(a | L) < c/g(a)$ for those for whom it does not. Whenever $c/g(a) < 1$, this simply downweights observations that rely more on extrapolation so that they are not as influential in the estimators, and imply an interpretation as outlined below otherwise

Interpretation. To understand the implication and interpretation of using this estimand, consider our motivating data example: first, note that both $m(a)$ and $m_w(a)$ are undefined outside the support region of A . For instance, in our study the CDRC is undefined for negative concentration values and biologically implausible concentration values of > 40 mg/L. Suppose we are interested in $E(Y^a)$ for $a = 0.5$ mg/L: for all patients with covariate regions \mathbf{I} that have $g(0.5 | \mathbf{I}) > c$ we stick to the CDRC $m(a)$, such that we obtain $E(Y^a)$; but for those where this does not hold, maybe because they are ultraslow metabolizers who will not clear the drug fast enough to ever achieve 0.5 mg/L under full adherence (i.e., $g(0.5 | \text{ultraslow, adherent}) = 0$), we target $E(Y | a)$. This means that we do not require positivity or rely on parametric extrapolations in poor support regions (where the intervention seems “unrealistic”); we rather use the present associations to allow individual concentration trajectories that lead to $E(Y | 0.5)$ for those patient groups. Thus, the proposed estimand (5) offers a tradeoff between identifiability and interpretability: that is, a

tradeoff between estimating the CDRC as closely as possible, at the risk of bias due to positivity violations because of the continuous intervention; and minimizing the risk of bias due to positivity violations, at the cost of redefining the estimand. More details on possible interpretations are given in Sections 3.2.2 and 3.5.

Choice of c . As in the case for binary interventions, when truncating the propensity score for inverse probability weighting or targeted maximum likelihood estimation, one may use rules of thumbs and simulation evidence to decide for c . This is because we want the conditional treatment density to be bounded away from zero, to avoid negative effects of (near-)positivity violations. Possible ad-hoc choices are 0.01 and $5/(\sqrt{(n) \ln n}/5)$ [29], though we argue below that multiple c 's may be selected and presented based on diagnostics, see Section 3.4.

Estimands under different weight functions. In principal, one could construct other weight functions too. An obvious choice would be to simply use the marginal treatment density $g(a)$ as a weight. Such a weight choice has been motivated in the context of estimands that are defined through parameters in a marginal structural working model for (continuous summaries of) longitudinal binary interventions [30]. In this case, greater weight is given to interventions with greater marginal support; and the more support there is for each possible intervention choice, the estimand will be closer to the CDRC. While intuitively it may make sense to rely more on interventions that are more often observed in the data, this approach does not directly address positivity violations; in the example above, there may be enough patients with concentrations close to 0.5 mg/L overall, but not among ultraslow metabolizers, which is the issue to be addressed.

3.2.2 | Estimands for Multiple Time Points

To illustrate our proposed concepts for multiple time points, consider data for two time points first: $(\mathbf{L}_0, A_0, Y_0, \mathbf{L}_1, A_1, Y_1)$. We are interested in $E(Y_1^{A_0=a_0, A_1=a_1})$. In principle, we may be interested in any (a_0, a_1) within the support region $\{\mathcal{A}_1 \times \mathcal{A}_2\}$. Practically, it may be possible that we only care about interventions (a_0, a_1) for which $a_0 = a_1 > 0$ (as in the motivating example), or $a_0 < a_1$: in this case, we may restrict the estimand to the respective region $\{\mathcal{A}_1^* \times \mathcal{A}_2^*\} \subset \{\mathcal{A}_1 \times \mathcal{A}_2\}$. Similar to a single time point there may be the situation where $g(a_1) > 0$, $g(a_1, a_0) > 0$ but $g(a_1 | \mathbf{I}_1, a_0, y_0, \mathbf{I}_0) \approx 0$. This corresponds to a situation where there is little support for the intervention value a_1 given that we already intervened with a_0 and given the covariate history; for regions, where this is the case, we address the respective positivity violations by redefining the estimand.

A. Causal Dose–Response–Curve. If there are multiple time points, the CDRC is

Estimand 1:

$$m_t : \bar{a}_t \mapsto E(Y_t^{\bar{a}_t}) \quad \forall \bar{a}_t \in \bar{\mathcal{A}}_t, \forall t \in \mathcal{T} \quad (6)$$

where $\mathcal{T} = \{0, 1, \dots, T\}$ and $\bar{\mathcal{A}}_t = \{\mathcal{A}_1 \times \dots \times \mathcal{A}_t\}$. The estimand (6) can, in principle, be identified through various ways. A possible option under the assumptions of sequential conditional

exchangeability, consistency and positivity [31] – as defined below – is the sequential g-formula (also known as the iterated conditional expectation representation [3]):

$$\begin{aligned} & \mathbb{E}(Y_t^{\bar{a}_t}) \\ &= \mathbb{E}(\dots \mathbb{E}(\mathbb{E}(Y_t | A_t = a_t, \mathbf{H}_t) | A_{t-1} = a_{t-1}, \mathbf{H}_{t-1}) \dots | A_0 = a_0, \mathbf{H}_0)) \end{aligned} \quad (7)$$

In the above expression, \bar{Y}_{t-1} is part of $\bar{\mathbf{L}}_t$, and thus \mathbf{H}_t . Because A_t is continuous, a *strong positivity* requirement corresponds to

$$\begin{aligned} & \inf_{a_t \in \mathcal{A}_t} g(a_t | \mathbf{h}_t) > 0 \quad \text{whenever} \\ & p_0(\mathbf{1}_t | A_{t-1} = a_{t-1}, \mathbf{H}_{t-1} = \mathbf{h}_{t-1}) > 0 \quad \forall t \in \mathcal{T} \end{aligned} \quad (8)$$

This means, we require within the support region $\bar{\mathcal{A}}_t$ a positive conditional treatment density for each a_t , given its past. This strong assumption may be relaxed either under additional parametric modeling assumptions or by a restriction to some $\bar{\mathcal{A}}_t^* \subset \bar{\mathcal{A}}_t$ or by a redefinition of the estimand. We continue with the latter strategy. Consistency in the multiple time-point case is the requirement that $Y_t^{\bar{a}_t} = Y_t$ if $\bar{A}_t = \bar{a}_t$ and $\bar{\mathbf{L}}_t^{\bar{a}_{t-1}} = \bar{\mathbf{L}}_t$ if $\bar{A}_{t-1} = \bar{a}_{t-1}$. With sequential conditional exchangeability we require the counterfactual outcome under the assigned treatment trajectory to be independent of the actually assigned treatment at time t , given the past: $Y_t^{\bar{a}_t} \perp\!\!\!\perp A_t | \mathbf{H}_t$ for $t = 0, \dots, T$.

Note: Sometimes, we may want to visualize the CDRC graphically, for example, by plotting $m_t(\bar{a}_t)$ for each t (and stratified by \mathbf{L}_0 , or a subset thereof, if meaningful) if the set of strategies is restricted to those that always assign the same value at each time point, see Figure 1b for an illustration. Alternatively, we may opt to plot the CDRC as a function of t , for some selected strategies \bar{a}_t , see Figure 1c. This may be useful if intervention values change over time, or for survival settings.

To link the above longitudinal g-computation formula (7) to our proposals, it may be written in terms of the following recursion. Let $\bar{m}_{t+1} = Y_t$. For $s = t, \dots, 0$ recursively define

$$\bar{m}_s : (a_s, \mathbf{h}_s) \mapsto \int \bar{m}_{s+1}(a_{s+1}, \mathbf{h}_{s+1}) dP_0(\mathbf{1}_{s+1} | A_s = a_s, \mathbf{H}_s = \mathbf{h}_s) \quad (9)$$

where $\int \bar{m}_{s+1}(a_{s+1}, \mathbf{h}_{s+1}) dP_0(\mathbf{1}_{s+1} | A_s = a_s, \mathbf{H}_s = \mathbf{h}_s) = E(\bar{m}_{s+1}(a_{s+1}, \mathbf{h}_{s+1}) | A_s = a_s, \mathbf{H}_s = \mathbf{h}_s)$.

Then, the counterfactual mean outcome $E(Y_t^{\bar{a}_t})$ is identified as $\mathbb{E}(\bar{m}_0(a_0, \mathbf{h}_0))$, which follows from the reexpression of (9) in terms of (7) by recursively evaluating the integral; see Appendix A.1 for details. The above recursive integral is well defined only if the positivity assumption (8) is met. As for the single time-point case, we propose to address violations of the positivity assumption by targeting a modified identifying expression, which we express through the following recursive integral:

$$\begin{aligned} & \bar{m}_{w,s} : (a_s, \mathbf{h}_s) \mapsto \int \bar{m}_{s+1} \\ & (a_{s+1}, \mathbf{h}_{s+1}) w_s(a_{s+1}, \mathbf{h}_{s+1}) dP_0(\mathbf{1}_{s+1} | A_s = a_s, \mathbf{H}_s = \mathbf{h}_s) \end{aligned} \quad (10)$$

As before, if the weight function $w_s(a_{s+1}, \mathbf{h}_{s+1})$ is equal to one, the above expression can be used to define the actual CDRC (6)

through recursive evaluation of (10). If, however, the weight function is equal to

$$w_s(a_{s+1}, \mathbf{h}_{s+1}) = \frac{g_s(a_{s+1} | \mathbf{h}_{s+1})}{g_s(a_{s+1} | a_s, \mathbf{h}_s)} \quad (11)$$

then the expression becomes

$$\begin{aligned} & \bar{m}_{w,s} : (a_s, \mathbf{h}_s) \mapsto \int \bar{m}_{s+1}(a_{s+1}, \mathbf{h}_{s+1}) \\ & dP_0(\mathbf{1}_{s+1} | A_{s+1} = a_{s+1}, A_s = a_s, \mathbf{H}_s = \mathbf{h}_s) \end{aligned} \quad (12)$$

by application of Bayes' rule. Intuitively, the above quantity does not remove confounding of the relation between A_{s+1} and Y_{s+1} , because it conditions on $A_{s+1} = a_{s+1}$ rather than fixing (i.e., intervening on) $A_{s+1} = a_{s+1}$. However, the above expression does not require the positivity assumption to be well defined, as motivated further below. Generally speaking, applying the above weights iteratively over all time points (and all units) leads to

$$E(Y_t | \bar{a}_t)$$

see Appendix A.2 for details.

For the longitudinal case, we follow a strategy similar to that considered for the single time-point case, and consider a compromise between satisfying the positivity assumption and adjusting for confounding. This compromise can be achieved by using weight functions that are 1 in areas of good support, and are equal to (11) in areas of poor support:

$$\begin{aligned} & w_s(a_{s+1}, \mathbf{h}_{s+1}, c) \\ &= \begin{cases} 1 & \text{if } g_s(a_{s+1} | \mathbf{h}_{s+1}) > c \\ \frac{g_s(a_{s+1} | \mathbf{h}_{s+1})}{g_s(a_{s+1} | a_s, \mathbf{h}_s)} & \text{if } g_s(a_{s+1} | \mathbf{h}_{s+1}) \leq c \text{ (and } g_s(a_{s+1} | a_s, \mathbf{h}_s) > c) \end{cases} \end{aligned} \quad (13)$$

Using the weights (13) in (10) has similar implications as for a single time point; that is, we obtain a weighted average of counterfactuals depending on the units' support for a given strategy \bar{a}_t . Formally, Estimand 2 equates to:

$$\begin{aligned} & m_{w,t}(\bar{a}_t) = E(Y_t^{\bar{a}_t} | \{g(a_s | \mathbf{H}_s) > c\}_{s=0}^t) P(\{g(a_s | \mathbf{H}_s) > c\}_{s=0}^t) + \\ & E(Y_t^{\bar{a}_t} \bar{w}_t | \bigvee_{s=0}^t \{g(a_s | \mathbf{H}_s) \leq c\}) \\ & P(\bigvee_{s=0}^t \{g(a_s | \mathbf{H}_s) \leq c\}) \quad \forall \bar{a}_t \in \bar{\mathcal{A}}_t, \forall t \in \mathcal{T} \end{aligned} \quad (14)$$

where $\{g(a_s | \mathbf{H}_s) > c\}_{s=0}^t = (g(a_t | \mathbf{H}_t) > c, \dots, g(a_0 | \mathbf{H}_0) > c)$, $\bigvee_{s=0}^t \{g(a_s | \mathbf{H}_s) \leq c\}$ indicates that at least one $g(a_s | \mathbf{H}_s) \leq c$, $w_t = w_t(a_{t+1}, \mathbf{h}_{t+1}, c)$, and the notation \bar{w}_t indicates that the weights have been applied up to and including time t . Practically, this means that

- i. For units where there is enough conditional support for a strategy \bar{a}_s in terms of $g_s(a_s | \mathbf{h}_s) > c$ for all time points s up to and including time t , the weights are 1. For those, we still target the CDRC, that is, $E(Y_t^{\bar{a}_t})$.
- ii. For units that do *not* have enough conditional support for a strategy at all time points, the weights in (13), second row, are used. With this, intuitively we target $E(Y_t | \bar{a}_t)$ as outlined above.

- iii. If there exist units for which there is support for a strategy at some points, but not at others, the weights will be 1 at some time points, but not for all. The implication is that we target the CDRC as much as this possible and deviate from it only at the time points where it is necessary.

Note that for the second estimand *no positivity assumption* as defined in (8) is required for identification because whenever a (near-)positivity violation is present in terms of $g_s(a_s | \mathbf{h}_s) \leq c$ the weights (13) redefine the estimand in a way such that the assumption is not needed.

If the denominator in (13) is very small, that is, $\leq c$, one could replace it with a very small ad-hoc value. Alternatively, one may want to compromise in the sense of evaluating the conditional density until previous time points:

$$w_s(a_{s+1}, \mathbf{h}_{s+1}, c) = \begin{cases} 1 & \text{if } g_s(a_{s+1} | \mathbf{h}_{s+1}) > c \\ \frac{g_s(a_{s+1} | \mathbf{h}_{s+1})}{g_s(a_{s+1} | a_s, \mathbf{h}_s)} & \text{if } g_s(a_{s+1} | \mathbf{h}_{s+1}) \leq c \text{ and } g_s(a_{s+1} | a_s, \mathbf{h}_s) > c \\ \frac{g_s(a_{s+1} | \mathbf{h}_{s+1})}{g_s(a_{s+1} | a_{s-1}, \mathbf{h}_{s-1})} & \text{if } g_s(a_{s+1} | \mathbf{h}_{s+1}) \leq c \text{ and } g_s(a_{s+1} | a_s, \mathbf{h}_s) \leq c \\ & \text{and } g_s(a_{s+1} | a_{s-1}, \mathbf{h}_{s-1}) > c \\ \vdots & \vdots \\ \frac{g_s(a_{s+1} | \mathbf{h}_{s+1})}{g_s(a_{s+1})} & \text{otherwise} \end{cases} \quad (15)$$

Interpretation of the weighted estimand. The weighted estimand can be interpreted as follows: those units which have a covariate trajectory that makes the intervention value of interest at t not unlikely to occur (under the desired intervention before t), receive the intervention at t ; all other units get different interventions that produce, on average, outcomes as we would expect among those who actually follow the intervention trajectory of interest (if the weight denominator is well-defined). Informally speaking, we calculate the CDRC in regions of enough support and stick to the actual research question, but make use of the marginal associations otherwise, if possible.

As a practical illustration for this interpretation, consider our motivating study: if, given the covariate and intervention history, it seems possible for a child to have a concentration level of a mg/L, then we do indeed calculate the counterfactual outcome under a . For some patients however, it may be unlikely (or even biologically impossible!) to actually observe this concentration level: for example, children who are adherent to their drug regimen and got an appropriate drug dose prescribed, but are slow metabolizers will likely never be able to have very low concentration values. In this case, we do not consider the intervention to be *feasible* and rather let those patients have individual concentration levels which generate outcomes that are typical for children with a mg/L. In short, we calculate the probability of failure at time t if the concentration level is set at a for all patients where this seems “feasible”, and otherwise to individual concentration trajectories that produce “typical” outcomes with a mg/L.

The more complex cases of the weighted estimand reflect the fact that certain strategies become unlikely only at certain time points: for example, patients who are adherent to their treatment and receive always the same dose will likely not have heavy varying concentrations in their body; thus, a strategy that

looks at the effect of a concentration jump during follow-up, for example, (3, 3, 3, 0) will only be unlikely at the fourth visit, but not before—and hence a deviation from the CDRC will only be needed at the fourth time point.

The interpretation of this weighted dose–response curve seems somewhat unusual; note however that this approach has the advantages that (i) compared to a naive approach it does not require the positivity assumption, and (ii) compared to LMTP’s it both sticks to the actual research question as close as “possible” and does not require a positivity assumption for the chosen policies of interest. We explain below, in Section 3.5, why we believe, however, that the weighted estimand (Estimand 2) should typically be presented together with the CDRC, that is, Estimand 1.

3.3 | Estimation

To develop an estimation strategy, it is useful to see that we can reexpress our target quantity (10) in terms of the iterated weighted conditional expectation representation

$$\begin{aligned} \bar{m}_{w,t} &= \mathbb{E}(\mathbb{E}(\dots \mathbb{E}(\mathbb{E}(Y_t w_t | \bar{A}_t = \bar{a}_t, \bar{\mathbf{L}}_t) w_{t-1} | \bar{A}_{t-1} \\ &= \bar{a}_{t-1}, \bar{\mathbf{L}}_{t-1}) \dots) w_0 | A_0 = a_0, \mathbf{L}_0)) \end{aligned} \quad (16)$$

where $w_s = w_s(a_{s+1}, \mathbf{h}_{s+1})$; see Appendix A.3 for details. This means that Estimand 2 can be written as a series of iterated weighted outcomes. An alternative representation, based on a weighted parametric g -formula is given in Appendix A.4.

For Estimand 2, one can build a substitution estimator based on either (16) or (A.3). That is, we use either a parametric or sequential g -formula type of approach where the (iterated) outcome is multiplied with the respective weight at each time point, regressed on its past, and recursively evaluated over time for a specific intervention \bar{a}_t . More specifically, a substitution estimator of (16) can be constructed as follows: (i) estimate the innermost expectation, that is, multiply the estimated weights (13) with the outcome at time t ; then, (ii) intervene with the first intervention trajectory of interest $\bar{a}_t^{(1)}$; (iii) predict the weighted outcome under this intervention; (iv) estimate the second innermost expectation by regressing the product of the prediction from (iii) and the respective weight on its past; (v) then, intervene with $\bar{a}_{t-1}^{(1)}$ and obtain the predicted weighted outcome under the intervention. These steps can be repeated until $t = 0$, and for each intervention trajectory of interest.

As the weighted (iterated) outcomes may possibly have a skewed, complex distribution (depending on the weights), we advocate for the use of (16) as a basis for estimation. This is because estimating the expectation only is typically easier than estimating the whole conditional distribution, as needed for (A.3). In many applications, a data-adaptive estimation approach may be a good choice for modeling the expectation of the weighted outcomes.

For the estimates of the weighted curve to be consistent, both the expected weighted (iterated) outcomes and the weights need to be estimated consistently.

Note that for $t = 0$, we suggest to calculate a weighted mean, instead of estimating the weighted expected outcome data-adaptively. This is because the former is typically more stable when Y_0 follows some standard distribution, but w_0 does not. It is a valid strategy as standardizing with respect to the pre-intervention variables L_0 and then calculating the weighted mean of the outcome is identical to standardizing the weighted outcome with respect to L_0 , and then calculating the mean. Another option for $t = 0$ is to facilitate step 6 before step 5 and then create a stacked dataset of the counterfactual outcomes for all intervention values a_0 , and then fit a weighted regression of Y_0 on A_0 , based on the stacked weight vectors. The fitted model can then be used to predict the expected outcome under all a_0 . This strategy is however not explored further in this paper.

The algorithm in Table 1 presents the detailed steps for a substitution estimator of (16). The algorithm is implemented

in the R-package mentioned below, and further illustrated in Appendix A.6.

3.3.1 | Multiple Interventions and Censoring

So far, we have described the case of 1 intervention variable A_t per time point. Of course, it is possible to employ the proposed methodology for multiple intervention variables $\mathbf{A}_t = (A_{t,1}, A_{t,2}, \dots, A_{t,p})$. In this case, one simply has to estimate the conditional densities for all those intervention variables, and taking their product, at each time point. This follows from replacing (11) with $g_s(\mathbf{a}_{s+1} | \mathbf{h}_{s+1}) / g_s(\mathbf{a}_{s+1} | \mathbf{a}_s, \mathbf{h}_s)$, and then factorizing the joint intervention distribution based on the assumed time ordering of $(L_s, A_{s,1}, A_{s,2}, \dots, A_{s,p}, Y_s)$. Intervention variables may include censoring variables: we may, for instance, construct estimands under the intervention of no censoring. A possible option to apply the proposed methodology is then to use

TABLE 1 | Algorithm for estimating Estimand 2 at time t .

Step 0a	Define a set of interventions $\bar{\mathcal{A}}_t$, with $ \bar{\mathcal{A}}_t = n_a$ and $\bar{A}_t^{(j)}$ is the j^{th} element of $\bar{\mathcal{A}}_t$, $j = 1, \dots, n_a$.
Step 0b	Set $\bar{Y}_t = Y_t$.
For $s = t, \dots, 1$	
Step 1a	Estimate the conditional density $g(A_s \mathbf{H}_s)$, see also Footnote 1.
Step 1b	Estimate the conditional density $g(A_s A_{s-1}, \mathbf{H}_{s-1})$, see also Footnote 1.
Step 2	Set $a_s = a_s^{(1)}$, where $a_s^{(1)}$ is the s^{th} element of intervention $\bar{\mathcal{A}}_t^{(1)} \in \bar{\mathcal{A}}_t$.
Step 3a	Plug in $a_s^{(1)}$ into the estimated densities from step 1, to calculate $\hat{g}(a_s^{(1)} \mathbf{h}_s)$ and $\hat{g}(a_s^{(1)} a_{s-1}^{(1)}, \mathbf{h}_{s-1})$.
Step 3b	Calculate the weights $w_s(a_s^{(1)}, c)$ from (15) based on the estimates from 3a.
Step 4	If $\hat{g}(a_s^{(1)} a_{s-1}^{(1)}, \mathbf{h}_{s-1}) \leq c$, estimate $g(a_s^{(1)} a_{s^*}^{(1)}, \mathbf{h}_{s^*})$ as required by the definition of (15) for $s^* = s - 2, \dots, 0$.
Step 5	Estimate $\mathbb{E}(w_s(a_s^{(1)} c) \bar{Y}_s A_s, \mathbf{H}_s)$, see also Footnote 2.
Step 6	Predict $\bar{Y}_{s-1} = \hat{\mathbb{E}}(\hat{w}_s(a_s^{(1)} c) \bar{Y}_s \bar{A}_s = \bar{a}_s^{(1)}, \mathbf{I}_s)$ based on the fitted model from step 4 and the given intervention $\bar{a}_s^{(1)}$.
For $t = 0$	
Step 1a	Estimate the conditional density $g(A_0 \mathbf{L}_0)$, see also Footnote 1.
Step 1b	Estimate the conditional density $g(A_0)$, see also Footnote 1.
Step 2	Set $a_0 = a_0^{(1)}$.
Step 3a	Calculate $\hat{g}(a_0^{(1)} \mathbf{I}_0)$ and $\hat{g}(a_0^{(1)})$.
Step 3b	Calculate the weights $w_0(a_0^{(1)}, c)$. If $\hat{g}(a_0^{(1)}) \leq c$, then $m_{w,t}(\bar{a}_t^{(1)}, c)$ is undefined.
Step 4	Estimate $\mathbb{E}(\bar{Y}_0 A_0, \mathbf{L}_0)$.
Step 5	Calculate $\hat{m}_{w,t}(\bar{a}_t^{(1)}, c) = \hat{\mathbb{E}}_{w(c)}(Y_t^{\bar{a}_t^{(1)}}) = (\sum_{i=1}^n w_{0,i})^{-1} (\hat{w}_0(a_0^{(1)}, c))^T \bar{Y}_{-1}$; that is, obtain the estimate of Estimand 2 at $\bar{a}_t^{(1)}$ through calculating the weighted mean of the iterated outcome under the respective intervention $a_0^{(1)}$ at $t=0$.
Then	
Step 6	Repeat steps 2–5 for the other interventions $\bar{A}_t^{(j)}$, $j = 2, \dots, n_a$. This yields an estimate of estimand 2 at t .
Step 7	Repeat steps 1–6 on B bootstrap samples to obtain confidence intervals.

Note: ¹The conditional treatment densities can be estimated with (i) parametric models, if appropriate, like the linear model, (ii) nonparametric flexible estimators, like highly-adaptive LASSO density estimation [32], (iii) a “binning strategy” where a logistic regression model models the probability of approximately observing the intervention of interest at time t , given one has followed the strategy so far and given the covariates, (iv) other options, like transformation models or generalized additive models of location, shape and scale [33, 34]. Items (i)–(iii) are implemented in our package mentioned below. ²The iterated weighted outcome regressions are recommended to be estimated data-adaptively, because the weighted outcomes are often non-symmetric. We recommend super learning for it [4], and this is what is implemented in the package mentioned below.

weights to target $E\left(Y_t^{\bar{C}_t=0} | A_t = a_t, \dots, A_0 = a_0\right)$ for units that do not have enough conditional support for the intervention strategy of interest. Appendix A.5 lists the required modifications of the estimation procedure in this case. Alternatively, one may use weights that lead to $E(Y_t | C_t = c_t, A_t = a_t, \dots, C_0 = c_0, A_0 = a_0)$ under positivity violations. For this, one would require estimating the conditional censoring mechanisms in steps 1a and 1b too. There are however many subtleties and dangers related to the interpretation and estimation under censoring, and possibly competing events; for example, conditioning on the censoring indicators may lead to collider bias, intervening on censoring mechanisms to direct effect estimands [35], identification assumptions need to be refined and generic time-orderings where treatment variables are separated by different blocks of confounders may lead to more iterated expectations that have to be fitted. Those details go beyond the scope of this paper.

3.4 | Violations of the Positivity Assumption, Diagnostics and the Choice of c

Suppose we are interested in a subset of interventions $\bar{\mathcal{A}}_t^*$ that are part of the support region $\{\mathcal{A}_1 \times \dots \times \mathcal{A}_t\}$. A violation of the strong positivity assumption exists if for any particular intervention of interest $(\bar{a}_0, \dots, \bar{a}_t) \in \bar{\mathcal{A}}_t^*$ it happens that $g(\bar{a}_t | \mathbf{h}_t) = 0$. To diagnose the extent of *practical* positivity violations, one needs estimates of the conditional treatment densities, at each time point, and evaluated for each treatment strategy that is part of $\bar{\mathcal{A}}_t^*$ — and summarize them in a meaningful way. We propose two options for facilitating this:

1. Calculate the proportion of weights (13) that are different from 1 (and thus indicate positivity violations) for each time point, for a range of c 's and the interventions of interest. Figure 4c gives an example on how to visualize those proportions: it shows that for intervention trajectories close to $(0, \dots, 0)$ a high proportion of children have weights $\neq 1$ and thus for these interventions practical positivity violations exist (which we may want to address with the weighted estimand).
2. Estimate the conditional treatment densities, under the intervention trajectories of interest, with the following *binning strategy*: suppose $\bar{A}_t^{(j)} = (a_t^{(j)}, \dots, a_t^{(j)})$ is the j^{th} element of $\bar{\mathcal{A}}_t^*$ and the intervention values at t are already ordered such that $a_t^{(1)} < a_t^{(2)} < \dots < a_t^{(n_a)}$; then calculate

$$P\left(a_t^{(j)} \in \left[a_t^{(j)} - \frac{1}{2}(a_t^{(j)} - a_t^{(j-1)}); a_t^{(j)} + \frac{1}{2}(a_t^{(j)} - a_t^{(j+1)})\right] \mid \bar{a}_{t-1}^{(j)} \in \left[a_{t-1}^{(j)} - \frac{1}{2}(a_{t-1}^{(j)} - a_{t-1}^{(j-1)}); a_{t-1}^{(j)} + \frac{1}{2}(a_{t-1}^{(j)} - a_{t-1}^{(j+1)})\right], \bar{l}_t\right) \quad (17)$$

that is, we want to estimate the probability to approximately observe the intervention value $a_t^{(j)}$ under strategy j , given that one has followed the same strategy of interest j so far (until $t - 1$), and irrespective of the covariate history. Alternatively, instead of defining the bin widths through the intervention values of interest, we may calculate them data-adaptively [36]. Estimating

the mean of those probabilities over all observed \bar{l}_t in a particular data set serves as diagnostic tool to measure the support for each rule j , at each time point. One could estimate (17) with standard regression techniques among the subset of those units who followed the respective strategy until $t - 1$, and present a summary of those probabilities as a rough measure of support for each intervention trajectory of interest. This approach gives a sense of actual units following the strategies of interest in the data, given the covariates. Examples are given in Figure B2.

Choice of c : The first diagnostic, possibly visualized as in Figure 4c, can be used to get a sense which set of c 's are informative to present weighted curves (as argued in Section 3.5). More specifically, those c 's which are close enough to zero to detect positivity violations, but do not alter the intervention values of interest for a too large proportion of units (and for too many intervention values), are good candidates. Alternatively, as indicated above, similar to the case when truncating the propensity score for inverse probability of treatment weighting or TMLE, ad-hoc choices could be used to decide when a treatment density is considered to be bounded away from zero (e.g., 0.01).

3.5 | The Case for Presenting Both Estimands

Consider again our running example where we are interested in the probability of failure at 84 weeks under concentration levels of $(0, 0, \dots, 0), \dots, (6, 6, \dots, 6)$. As indicated above, adherent patients who are ultraslow metabolizers may not be able to have low concentration levels, for example, $g(0 | \text{ultraslow, adherent}) = 0$. In such a case, there is a strong case to redefine the estimand and not enforce to calculate $E\left(Y_t^{(0,0,\dots,0)}\right)$. Now, suppose we have $g_{84}(6 | \text{prior concentrations} = 6, \text{past covariates}) > 0$: obviously, there is no need to redefine the estimand. However, due to the finite sample the actual estimate may be $\hat{g}_{84}(6 | \text{prior concentrations} = 6, \text{past covariates}) \approx 0$. Obviously, we would in this region more likely rely on extrapolation than changing the estimand.

Now, practically, whenever an estimate of the conditional treatment density is close to zero, we do not necessarily know whether this is a finite sample (or estimation) issue, or due to “infeasibility”, that is, an illogical intervention given the history of a patient. Presenting Estimand 1 shows the case where nothing is infeasible and we rely on extrapolations; presenting Estimand 2, possibly shown for multiple c 's, shows the change in the curve as we consider more intervention trajectories to be infeasible. Both estimands together give us a sense if main conclusions could change when positivity violations are addressed differently.

4 | Monte-Carlo Simulations

In this section, we evaluate the proposed methods in three different simulation settings. Simulation 1 considers a very simple setting, as a basic reference for all approaches presented. In Simulation 2, a survival setting is considered, to explore the stability of standard, unweighted g-computation in more sophisticated setups. The third simulation is complex, and inspired by

the data-generating process of the data analysis. It serves as the most realistic setup for all method evaluations.

4.1 | Data Generating Processes and Estimands

Simulation 1: We simulated both a binary and normally distributed confounder, a continuous (normally distributed) intervention and a normally distributed outcome—for 3 time points and a sample size of $n = 10,000$. The exact model specifications are given in Appendix C.1. The intervention strategies of interest comprised intervention values in the interval $[2, 11]$ which were constant over time; that is, $\overline{A}_t^1 = \{(2,2,2), \dots, (11,11,11)\}$. The primary estimand ψ^1 is the CDRC (6) for $t = 1,2,3$ and all $\overline{a}_t \in \overline{A}_t^1$. The secondary estimand ψ_w^1 is the weighted curve defined through (14), with the weights (13), for the same intervention strategies.

Simulation 2: We simulated a continuous (normally distributed) intervention, two covariates (one of which is a confounder), an event outcome and a censoring indicator—for 5 time points and varying sample sizes of $n_1 = 500, n_2 = 1,000, n_3 = 10,000$ and $n_4 = 50,000$. The exact model specifications are given in Appendix C.2. The intervention strategies of interest comprised intervention values in the interval $[-7, 13]$ which were constant over time; that is, $\overline{A}_t^2 = \{(-7, -7, -7, -7, -7), \dots, (13,13,13,13,13)\}$. The estimand ψ^2 is the CDRC (6) for $t = 1,2,3,4,5$, all $\overline{a}_t \in \overline{A}_t^2$ and under no censoring ($\overline{c}_t = (0, \dots, 0)$).

Simulation 3: We simulated data inspired by the data generating process of the data example outlined in Sections 2 and 4.3 as well as Figure 3. The continuous intervention refers to drug concentration (of evavirenz), modeled through a truncated normal distribution. The binary outcome of interest is viral failure. Time-varying confounders, affected by prior interventions, are weight and adherence. Other variables include co-morbidities and drug dose (time-varying) as well as sex, genotype, age and NRTI regimen. We considered 5 clinic visits and a sample size of $n = 1,000$. The exact model specifications are given in Appendix C.3. The intervention strategies of interest comprised concentration values in the interval $[0, 10]$ which were constant over time; that is, $\overline{A}_t^3 = \{(0,0,0,0,0), \dots, (10,10,10,10,10)\}$. The primary estimand ψ_3 is the CDRC (6) for $t = 1,2,3,4,5$ and all $\overline{a}_t \in \overline{A}_t^3$. The secondary estimand ψ_w^3 is the weighted curve (14) with the weights (13) for the same intervention strategies.

4.2 | Estimation and Evaluation

All simulations were evaluated based on the results of $\mathcal{R} = 1,000$ simulation runs. We evaluated the bias of estimating ψ^1, ψ^2 and ψ^3 with standard parametric g-computation with respect to the true CDRC. In Simulations 1 and 3 model specification for g-computation was based on variable screening with LASSO, Simulation 2 explored the idealized setting of using correct model specifications. We further evaluated the bias of the estimated weighted curve with the true weighted curve, for $c = 1$; that is, we looked whether $E(Y_t | A_t = a_t, \dots, A_0 = a_0)$ could be recovered for every intervention strategy of interest. Estimation was based

on the algorithm of Table 1. The density estimates, which are required for estimating the weights, were based on both appropriate parametric regression models (i.e., linear regression) and binning the continuous intervention into intervals; a computationally more sophisticated data-adaptive approach for density estimation is considered in Section 4.3. We then compared the estimated CDRC and weighted curve ($c = 1$) with other weighted curves where $c = 0.01$ and $c = 0.001$. To estimate the iterated conditional expectations of the weighted outcome, we used super learning, that is, a data adaptive approach [4]. Our learner sets included different generalized linear (additive) regression models (with penalized splines, and optionally interactions), multivariate adaptive regression splines, LASSO estimates and regression trees—after prior variable screening with LASSO and Cramer's V [37]. Lastly, we visualize the conditional support of all considered intervention strategies, in all simulations, as suggested in (17).


4.3 | Results

The results of the simulations are summarized in Figures 2, B1, and B2.

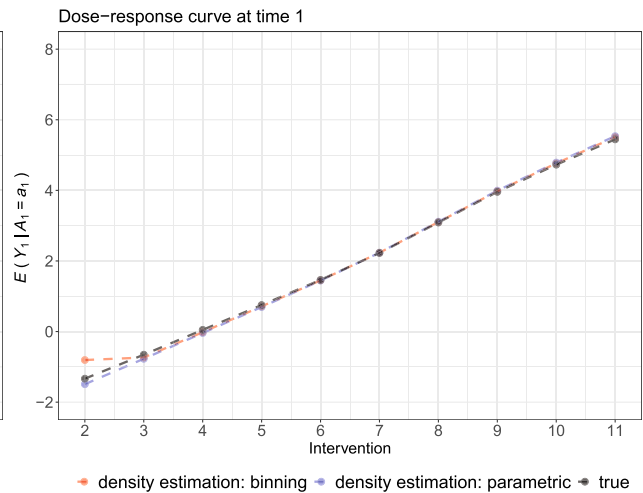
In Simulation 1, applying standard g-computation to the continuous intervention led to approximately unbiased estimates for all intervention values, through all time points (Figure 2a); and independent of the level of support (Figure B2a). The weighted curve, with $c = 1$, recovered the association $E(Y_1 | A_1 = a_1)$ perfectly for the first time point, independent of the density estimation strategy (Figure 2b). For later time points, some bias can be observed in regions of lower support (i.e., for intervention values between 9 and 11), see Figure B1a.

Simulation 2 shows a setting in which g-computation estimates for the CDRC are not always unbiased, despite correct model specifications! Figure 2c visualizes this finding for the fifth time point: it can be seen that a sample size of $\geq 10,000$ is needed for approximately unbiased estimation of the CDRC over the range of all intervention strategies. Note, however, that for some time points lower sample sizes are sufficient to guarantee unbiased estimation (Figure B1b).


The most sophisticated simulation setting 3 reveals some more features of the proposed methods. First, we can see that in areas of lowest support, toward concentration values close to zero (Figure B2c), there is relevant bias of a standard g-computation analysis (Figure 2d); whereas in areas of reasonable to good support the CDRC estimates are approximately unbiased. The weighted curve with $c = 1$ can recover the associations $E(Y_t | A_t = a_t, \dots, A_1 = a_1)$ (Figure 2e) but there is some bias for later time points as the most critical area close to zero is approached, independent of the method used for estimating the weights (Figure 5d). Both Figures 2f and B1e highlight the behavior of Estimand 2 for $c = 0.01$ and $c = 0.001$: it can be clearly seen that the curves represent a compromise between the CDRC and the association represented by the weighted curve (with $c = 1$). Evaluating the results for intervention values of zero, shows that weighting the curve in areas of low support yields to a compromise that moves the CDRC away from the estimated high probabilities of viral failure to more moderate values informed by the




(a) Simulation 1: true and estimated CDRC at $t = 2$




(b) Simulation 1: weighted estimand, $c = 1$, $t = 1$




(c) Simulation 2: true and estimated CDRC at $t = 5$



(d) Simulation 3: true and estimated CDRC at $t = 5$



(e) Simulation 3: weighted estimand, $c = 1$, $t = 1$



(f) Simulation 3: CDRC's, $t = 5$, $c = 0.001, 0.01, 1$

FIGURE 2 | Results of the Monte-Carlo simulations.




FIGURE 3 | Directed acyclic graph for the data analysis. The intervention variable is shown in green (efavirenz concentration), the outcome in blue (viral load at the end of follow-up). Unmeasured variables are colored in gray. Both MEMS and weight are time-dependent confounders which are affected by prior treatment nodes.

observed association. Knowledge of this behavior may be informative for the data analysis below.

5 | Data Analysis

We now illustrate the ideas based on the data from Bienczak et al. [25], introduced in Section 2. We consider the trial visits at $t = 0, 6, 36, 48, 60, 84$ weeks of 125 children on an efavirenz-based treatment regime. The intervention of interest is efavirenz mid-dose interval concentration (A_t), defined as plasma concentration (in mg/L) 12 h after dose; the outcome is viral failure (Y_t , defined as > 100 copies/mL). Measured baseline variables, which we included in the analysis, are $\mathbf{L}_0 = \{\text{sex, genotype, age, the nucleoside reverse transcriptase inhibitors drug (NRTI), weight}\}$. Genotype refers to the metabolism status (slow, intermediate, extensive) related to the single nucleotide polymorphisms in the CYP2B6 gene, which is relevant for metabolizing efavirenz and directly affects its concentration in the body. Measured follow-up variables are $\mathbf{L}_t = \{\text{weight, adherence (measured through memory caps MEMS), dose}\}$.

The assumed data generating process is visualized in the DAG in Figure 3, and explained in more detail in Appendix D. Briefly, both weight and adherence are time-dependent confounders, potentially affected by prior concentration trajectories, which are needed for identification of the CDRC.

Our target estimands are the CDRC (6) and the weighted curve (14) at $t = 84$ weeks. The intervention strategies of interest are $\bar{\mathcal{A}}_t = \{(0,0,0,0), \dots, (6,6,6,6)\}$.

The analysis illustrates the ideas based on a complete case analysis of all measured variables represented in the DAG ($n = 58$), but excluding dose (not needed for identification) and MEMS (due to the high proportion of missingness) [38].

We estimated the CDRC both with sequential and parametric g-computation using the intervention strategies $\bar{\mathcal{A}}_t$. The

estimation of the weighted curve followed the algorithm of Table 1. The conditional treatment densities, which are needed for the construction of the weights, were estimated both parametrically based on the implied distributions from linear models (for concentrations under 5 mg/L) and with highly-adaptive LASSO conditional density estimation (for concentrations ≥ 5 mg/L) [32, 39]. This is because the density in the lower concentration regions were approximately normally distributed, but more complex for higher values.

The conditional expectation of the weighted outcome (Step 4 of the algorithm) was estimated data-adaptively with super learning using the following learning algorithms: multivariate adaptive regression splines, generalized linear models (also with interactions), ordinal cumulative probability models, generalized additive models as well as the mean and median. Prior variable screening, which was essential given the small sample size, was based on both the LASSO and Cramer's V .

We estimated the weighted curve for $c = 0.001, 0.01, 0.025, 0.2, 1$. We also calculate the support of the continuous intervention strategies $\bar{\mathcal{A}}_t$ as described in Section 3.4.

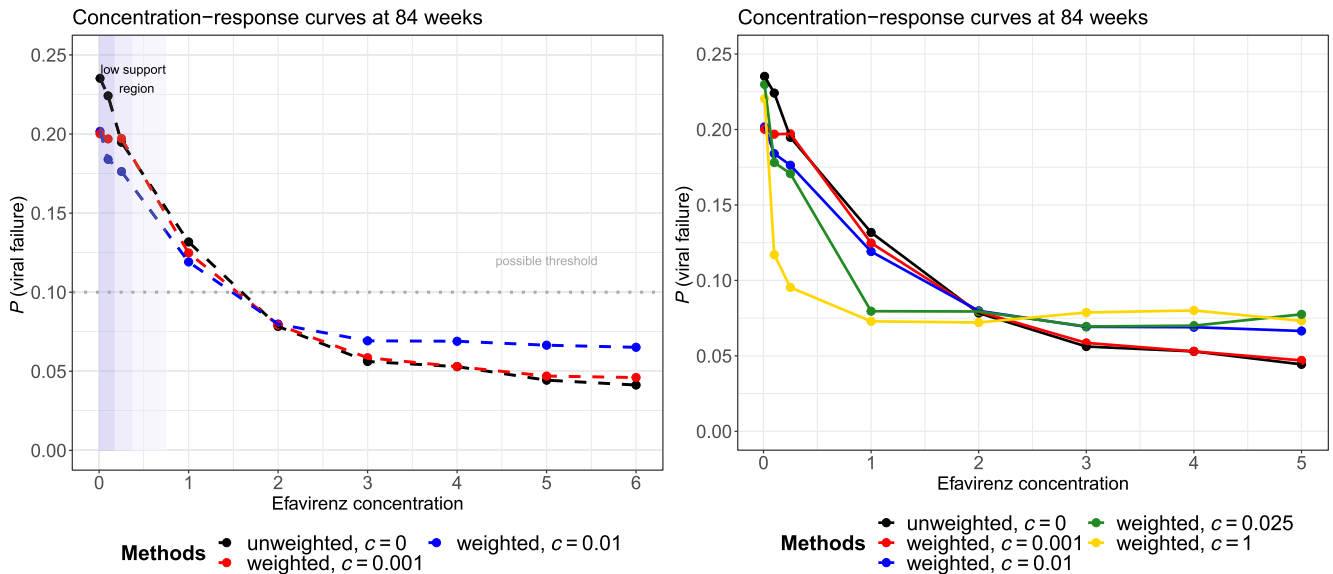
5.1 | Results

The main results of the analyses are given in Figure 4.

Figure 4a,b show that the estimated CDRC (black solid line) suggests higher probabilities of failure with lower concentration values. The curve is steep in the region from 0 to 2 mg, which is even more pronounced for parametric g-computation when compared to sequential g-computation (Figure B1f).

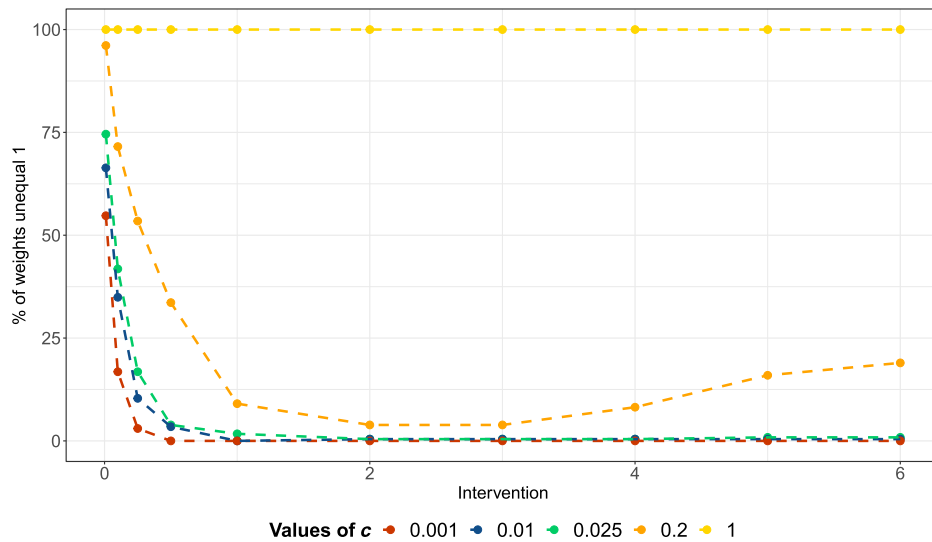
Both Figures 4a and B2d illustrate the low level of (conditional support) support for extremely small concentration values close to 0 mg/L.

Note that in Figure 4a the area shaded in dark blue relates to low conditional support, defined as a proportion $> 50\%$ of weights being unequal from 1; that is, more than 50% of observations



(a) CDRC, weighted curve and conditional support, represented through background shading. Support is defined as proportion of weights, which are unequal to 1, see text for details.

(b) CDRC and weighted curves



(c) Conditional support and weight behaviour: percentage of weights which are unequal 1, as a function of c and concentration values, averaged over all visits.

FIGURE 4 | Results of the data analysis.

having an estimated conditional treatment density < 0.01 under the respective intervention values. Intermediate, light and no blue background shadings refer to percentages of 15%–50%, 5%–15% and $< 5\%$, respectively. More details on how the conditional support defined through weight summaries varies as a function of c and the intervention values is illustrated in Figure 4c. It can be seen that with a choice of $c \leq 0.025$, for most patients the actual intervention of interest seems feasible for concentrations values above 0.5 mg/L; however, concentrations ≤ 0.5 mg/L seem unrealistic for many patients and thus weighting is then used more often. As the proportion of weights, which are unequal 1, do not change much between $c = 0.01$ and $c = 0.025$, we report only curves for $c = 0.001$ and $c = 0.01$ in Figure 4a.

The weighted curves in Figure 4b are less steep in the crucial areas close to zero and –as in the simulation studies– we observe a tendency of the curves with $c = 0.001, 0.01, 0.025$ to provide a visual compromise between the CDRC and the estimand recovering the relevant association ($c = 1$).

The weighted curves show the probability of failure under different concentration levels, at least if the respective concentration level is likely possible (realistic, probable) to be observed given the patient’s covariate history; for patients where this seems unlikely, patients rather have individual concentrations that lead to outcomes comparable to the patient population with the respective concentration trajectory. For example, if patients are adherent

and slow metabolizers, it may be unlikely that they achieve low concentration levels and we do not enforce the intervention of interest. We let these patients rather behave individually in line with outcomes typically seen for low concentration values.

The practical implication is that if we do not want to extrapolate in regions of low support and/or do not consider specific concentration values to be feasible for some patients, the concentration-response curve is flatter. The current lower recommended target concentration limit is 1 mg/L, at which we still estimate high failure probabilities under standard g -computation approaches. *If, for example, we do not want to accept failure probabilities < 10%—but are worried about positivity violations—, then we see that main conclusions with respect to a lower recommended target concentration limit would not change when using Estimand 2 in addition to Estimand 1* (Figure 4a).

We recommend to report results similar to Figure 4a, that is presenting a figure that shows both Estimand 1 and Estimand 2 (possibly based on several c 's, informed by a graph as in Figure 4c).

6 | Discussion

We have introduced and evaluated various methods for causal effect estimation with continuous multiple time-point interventions. As an obvious first step, we investigated standard (parametric or sequential) g -formula based approaches for identification and estimation of the actual causal dose-response curve. We emphasized that this has the advantage of sticking to the estimand of interest, but with a relatively strong positivity assumption which will likely be violated in most data analyses. Our simulations were designed to explore how well such a standard approach may extrapolate in sparse data regions and how the approach performs in both simple and complex settings. We found that in simple scenarios the CDRC estimates were approximately unbiased, but that in more complex settings sometimes large sample sizes were needed for good performance and that in regions of very poor support a relatively large bias could be observed. These findings suggest that the toolkit for causal effect estimation with longitudinal continuous interventions should therefore ideally be broader.

We therefore proposed the use of a weight function which returns the CDRC under enough conditional support, and makes use of the crude associations between the outcome at time t and the treatment history otherwise. This choice ensures that the estimand is always well-defined and one does not require positivity. Our simulations suggest that the provided estimation algorithm can indeed successfully recover the weighted estimand and illustrate the compromise that is practically achieved. As the proposed weights may have skewed distributions, we highlighted the importance of using a data-adaptive estimation approach, though appropriate interval estimation, beyond using bootstrapping, remains to be investigated.

We hope that our manuscript has shown that for continuous multiple time point interventions one has to ultimately make a trade-off between estimating the CDRC as closely as possible, at the risk of bias due to positivity violations and minimizing the risk of

bias due to positivity violations, at the cost of redefining the estimand. If the scientific question of interest allows a redefinition of the estimand in terms of longitudinal modified treatment policies [23, 40, 41] then this may be a great choice. Otherwise, it may be helpful to use g -computation algorithms for continuous interventions, visualize the estimates in appropriate curves, calculate the suggested diagnostics by estimating the conditional support and then estimate the proposed weighted curves as a magnifying glass for the curves behavior in areas of low support (where the intervention of interest may not be feasible or realistic).

Our suggestions can be extended in several directions: it may be possible to design different weight functions that keep the spirit of modifying estimands under continuous interventions only in areas of low support; the extent of the compromise, controlled by the tuning parameter c , may be chosen data-adaptively and our basic considerations for time-to-event data may be extended to cover more estimands, particular in the presence of competing risks.

6.1 | Software

All approaches considered in this paper have been implemented in the R -packages CICI and CICIplus available at <https://cran.r-project.org/web/packages/CICI/index.html> and <https://github.com/MichaelSchomaker/CICIplus>.

Acknowledgments

We are grateful for the support of the CHAPAS-3 trial team, their advice regarding the illustrative data analysis and making their data available to us. We would like to particularly thank David Burger, Sarah Walker, Di Gibb, Andrzej Bienczak, and Elizabeth Kaudha. We would also like to acknowledge Daniel Saggau, who has contributed to the data-generating processes of our simulation setups and Igor Stojkov for contributing substantive knowledge to pharmacological angles relevant to the data example. Computations were performed using facilities provided by the University of Cape Town's ICTS High Performance Computing team. Michael Schomaker is supported by the German Research Foundations (DFG) Heisenberg Programm (grants 465412241 and 465412441).

Conflicts of Interest

The authors declare no conflicts of interest.

Data Availability Statement

The authors have nothing to report.

References

1. J. M. Robins, M. A. Hernan, and B. Brumback, "Marginal Structural Models and Causal Inference in Epidemiology," *Epidemiology* 11, no. 5 (2000): 550–560.
2. J. Robins, "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, no. 9–12 (1986): 1393–1512.
3. H. Bang and J. M. Robins, "Doubly Robust Estimation in Missing Data and Causal Inference Models," *Biometrics* 64, no. 2 (2005): 962–972.
4. M. Van der Laan and S. Rose, *Targeted Learning* (New York: Springer, 2011).
5. M. E. Schnitzer, M. J. van der Laan, E. E. Moodie, and R. W. Platt, "Effect of Breastfeeding on Gastrointestinal Infection in Infants:

- A Targeted Maximum Likelihood Approach for Clustered Longitudinal Data,” *Annals of Applied Statistics* 8, no. 2 (2014): 703–725.
6. M. Schomaker, M. A. Luque Fernandez, V. Leroy, and M. A. Davies, “Using Longitudinal Targeted Maximum Likelihood Estimation in Complex Settings With Dynamic Interventions,” *Statistics in Medicine* 38 (2019): 4888–4911.
7. P. Baumann, M. Schomaker, and E. Rossi, “Estimating the Effect of Central Bank Independence on Inflation Using Longitudinal Targeted Maximum Likelihood Estimation,” *Journal of Causal Inference* 9, no. 1 (2021): 109–146.
8. H. Bell-Gorrod, M. P. Fox, A. Boule, et al., “The Impact of Delayed Switch to Second-Line Antiretroviral Therapy on Mortality, Depending on Failure Time Definition and CD4 Count at Failure,” *American Journal of Epidemiology* 189 (2020): 811–819.
9. L. E. Cain, M. S. Saag, M. Petersen, et al., “Using Observational Data to Emulate a Randomized Trial of Dynamic Treatment-Switching Strategies: An Application to Antiretroviral Therapy,” *International Journal of Epidemiology* 45, no. 6 (2016): 2038–2049.
10. N. Kreif, L. Tran, R. Grieve, B. De Stavola, R. C. Tasker, and M. Petersen, “Estimating the Comparative Effectiveness of Feeding Interventions in the Pediatric Intensive Care Unit: A Demonstration of Longitudinal Targeted Maximum Likelihood Estimation,” *American Journal of Epidemiology* 186, no. 12 (2017): 1370–1379.
11. K. Hirano and G. W. Imbens, *The Propensity Score With Continuous Treatments*, chapter 7 (Hoboken, NJ: Wiley, 2004), 73–84.
12. N. Kreif, R. Grieve, I. Díaz, and D. Harrison, “Evaluation of the Effect of a Continuous Treatment: A Machine Learning Approach With an Application to Treatment for Traumatic Brain Injury,” *Health Economics* 24, no. 9 (2015): 1213–1228.
13. E. H. Kennedy, Z. Ma, M. McHugh, and D. S. Small, “Nonparametric Methods for Doubly Robust Estimation of Continuous Treatment Effects,” *Journal of the Royal Statistical Society, Series B: Statistical Methodology* 79 (2017): 1229–1245.
14. T. Westling, P. Gilbert, and M. Carone, “Causal isotonic regression,” *Journal of the Royal Statistical Society, Series B* 82 (2020): 719–747.
15. Z. Zhang, J. Zhou, W. Cao, and J. Zhang, “Causal Inference With a Quantitative Exposure,” *Statistical Methods in Medical Research* 25, no. 1 (2016): 315–335.
16. A. F. Galvao and L. Wang, “Uniformly Semiparametric Efficient Estimation of Treatment Effects With a Continuous Treatment,” *Journal of the American Statistical Association* 110, no. 512 (2015): 1528–1542.
17. T. J. VanderWeele, Y. Chen, and H. Ahsan, “Inference for Causal Interactions for Continuous Exposures Under Dichotomization,” *Biometrics* 67, no. 4 (2011): 1414–1421.
18. I. Díaz and M. J. van der Laan, “Targeted Data Adaptive Estimation of the Causal Dose–Response Curve,” *Journal of Causal Inference* 1, no. 2 (2013): 171–192.
19. R. Neugebauer and M. van der Laan, “Nonparametric Causal Effects Based on Marginal Structural Models,” *Journal of Statistical Planning and Inference* 137, no. 2 (2007): 419–434.
20. S. Goetgeluk, S. Vansteelandt, and E. Goetghebeur, “Estimation of Controlled Direct Effects,” *Journal of the Royal Statistical Society, Series B: Statistical Methodology* 70 (2008): 1049–1066.
21. J. G. Young, M. A. Hernan, and J. M. Robins, “Identification, Estimation and Approximation of Risk Under Interventions That Depend on the Natural Value of Treatment Using Observational Data,” *Epidemiological Methods* 3, no. 1 (2014): 1–19.
22. J. G. Young, R. W. Logan, J. M. Robins, and M. A. Hernán, “Inverse Probability Weighted Estimation of Risk Under Representative Interventions in Observational Studies,” *Journal of the American Statistical Association* 114, no. 526 (2019): 938–947.
23. I. Díaz, N. Williams, K. L. Hoffman, and E. J. Schenck, “Non-Parametric Causal Effects Based on Longitudinal Modified Treatment Policies,” *Journal of the American Statistical Association* 118, no. 542 (2021): 1–16.
24. V. Mulenga, V. Musiime, A. Kekitiinwa, et al., “Abacavir, Zidovudine, or Stavudine as Paediatric Tablets for African HIV-Infected Children (Chapas-3): An Open-Label, Parallel-Group, Randomised Controlled Trial,” *Lancet Infectious Diseases* 16, no. 2 (2016): 169–179.
25. A. Bienczak, P. Denti, A. Cook, et al., “Plasma Efavirenz Exposure, Sex, and Age Predict Virological Response in HIV-Infected African Children,” *Journal of Acquired Immune Deficiency Syndromes* 73, no. 2 (2016): 161–168.
26. A. Bienczak, P. Denti, A. Cook, et al., “Determinants of Virological Outcome and Adverse Events in African Children Treated With Paediatric Nevirapine Fixed-Dose-Combination Tablets,” *AIDS* 31, no. 7 (2017): 905–915.
27. M. Hernan and J. Robins, *Causal Inference* (Boca Raton, USA: Chapman & Hall/CRC, 2020), <https://www.hsph.harvard.edu/miguel-hernan/causal-inference-book/>.
28. M. L. Petersen, K. E. Porter, S. Gruber, Y. Wang, and M. J. van der Laan, “Diagnosing and Responding to Violations in the Positivity Assumption,” *Statistical Methods in Medical Research* 21, no. 1 (2012): 31–54.
29. S. Gruber, R. V. Phillips, H. Lee, and M. J. van der Laan, “Data-Adaptive Selection of the Propensity Score Truncation Level for Inverse-Probability—Weighted and Targeted Maximum Likelihood Estimators of Marginal Point Treatment Effects,” *American Journal of Epidemiology* 191, no. 9 (2022): 1640–1651.
30. M. Petersen, J. Schwab, S. Gruber, N. Blaser, M. Schomaker, and M. van der Laan, “Targeted Maximum Likelihood Estimation for Dynamic and Static Longitudinal Marginal Structural Working Models,” *Journal of Causal Inference* 2, no. 2 (2014): 147–185.
31. M. J. van der Laan and S. Gruber, “Targeted Minimum Loss Based Estimation of Causal Effects of Multiple Time Point Interventions,” *International Journal of Biostatistics* 8, no. 1 (2012): 9.
32. N. S. Hejazi, M. J. van der Laan, and D. Benkeser, “Haldensify’: Highly Adaptive Lasso Conditional Density Estimation in R,” *Journal of Open Source Software* 7, no. 77 (2022): 4522, <https://doi.org/10.21105/joss.04522>.
33. T. Hothorn, T. Kneib, and P. Bühlmann, “Conditional Transformation Models,” *Journal of the Royal Statistical Society, Series B* 76, no. 1 (2014): 3–27.
34. M. D. Stasinopoulos, R. A. Rigby, G. Z. Heller, V. Voudouris, and F. D. Bastiani, *Flexible Regression and Smoothing: Using GAMLSS in R* (New York: Chapman and Hall/CRC, 2017).
35. J. G. Young, M. J. Stensrud, E. J. T. Tchetgen, and M. A. Hernan, “A Causal Framework for Classical Statistical Estimands in Failure-Time Settings With Competing Events,” *Statistics in Medicine* 39, no. 8 (2020): 1199–1236.
36. I. Díaz and M. van der Laan, “Super Learner Based Conditional Density Estimation With Application to Marginal Structural Models,” *International Journal of Biostatistics* 7, no. 1 (2011): 38, <https://doi.org/10.2202/1557-4679.1356>.
37. C. Heumann, M. Schomaker, and Shalabh, *Introduction to Statistics and Data Analysis—With Exercises, Solutions and Applications in R* (Heidelberg, Germany: Springer, 2023).
38. A. Holovchak, H. McIlleron, P. Denti, and M. Schomaker, “Recoverability of Causal Effects in a Longitudinal Study Under Presence of Missing Data,” arXiv Eprints, 2024, <https://arxiv.org/abs/2402.14562>.

39. S. Nima, D. B. Hejazi, and M. J. van der Laan, “*haldensify: Highly Adaptive Lasso Conditional Density Estimation*,” R package version 0.2.3, 2022, <https://github.com/nhejazi/haldensify>.

40. I. Díaz, K. L. Hoffman, and N. S. Hejazi, “Causal Survival Analysis Under Competing Risks Using Longitudinal Modified Treatment Policies,” *Lifetime Data Analysis* 30, no. 1 (2024): 213–236.

41. K. L. Hoffman, D. Salazar-Barreto, N. T. Williams, K. E. Rudolph, and I. Díaz, “Studying Continuous, Time-Varying, and/or Complex Exposures Using Longitudinal Modified Treatment Policies,” *Epidemiology* 35, no. 5 (2024): 667–675.

42. M. L. Petersen, “Commentary: Applying a Causal Road Map in Settings With Time-Dependent Confounding,” *Epidemiology* 25, no. 6 (2014): 898–901.

43. A. Bienczak, A. Cook, L. Wiesner, et al., “The Impact of Genetic Polymorphisms on the Pharmacokinetics of Efavirenz in African Children,” *British Journal of Clinical Pharmacology* 82, no. 1 (2016): 185–198.

Supporting Information

Additional supporting information can be found online in the Supporting Information section.

Appendix A

Additional Theory Details

A.1 | Iterated Nested Expectation Representation of Estimand 1

Using (9) corresponds to a recursive definition of Estimand 1. For $s = t$, the integral corresponds to

$$m_t(a_t, \mathbf{h}_t) = E(Y_t | A_t = a_t, \mathbf{H}_t = \mathbf{h}_t)$$

Moving to $s = t - 1$ yields

$$\begin{aligned} m_{t-1}(a_{t-1}, \mathbf{h}_{t-1}) &= E(m_t(a_t, \mathbf{h}_t) | A_{t-1} = a_{t-1}, \mathbf{H}_{t-1} = \mathbf{h}_{t-1}) \\ &= E\left(E\left(Y_t | \bar{A}_t = \bar{a}_t, \mathbf{H}_t = \mathbf{h}_t\right) | \bar{A}_{t-1} = \bar{a}_{t-1}, \mathbf{H}_{t-1} = \mathbf{h}_{t-1}\right) \end{aligned}$$

Reevaluating the integral from $s = t - 2, \dots, 0$ then leads to

$$\begin{aligned} m_0(a_0, \mathbf{h}_0) &= \mathbb{E}\left(\dots \mathbb{E}\left(E\left(Y_t | \bar{A}_t = \bar{a}_t, \bar{\mathbf{L}}_t\right) | \bar{A}_{t-1} = \bar{a}_{t-1}, \bar{\mathbf{L}}_{t-1}\right) \dots | A_0 = a_0, \mathbf{L}_0\right) \end{aligned}$$

This expression is identical to $\mathbb{E}\left(Y_t^{\bar{a}_t}\right)$ under conditional sequential exchangeability, positivity and consistency, which yields the sequential g -formula listed in (7). The result is known from Bang and Robins [3]. Briefly, first recall that positivity can be defined as in (8), consistency is the requirement that

$$Y_t^{\bar{a}_t} = Y_t \text{ if } \bar{A}_t = \bar{a}_t \quad (\text{A.1})$$

and sequential conditional exchangeability is defined as

$$Y_t^{\bar{a}_t} \prod_{r=0}^{t-1} A_r | H_t \quad (\text{A.2})$$

Suppose $t = 2$. Then,

$$\begin{aligned} E(Y_1^{a_1, a_0}) &= E(E(Y_1^{a_1, a_0} | L_0)) \\ &= E(E(Y_1^{a_1, a_0} | L_0, A_0 = a_0)) \\ &= E(E(E(Y_1^{a_1, a_0} | A_0 = a_0, L_0, L_1)) | L_0, A_0 = a_0) \\ &= E(E(E(Y_1^{a_1, a_0} | A_0 = a_0, L_0, L_1, A_1 = a_1)) | L_0, A_0 = a_0) \end{aligned}$$

$$= E\left(E\left(E\left(Y_1^{a_1, a_0} | \bar{A}_1 = \bar{a}_1, \bar{L}_1\right)\right) | L_0, A_0 = a_0\right)$$

$$= E\left(E\left(E\left(Y_1 | \bar{L}_1, \bar{A}_1 = \bar{a}_1\right)\right) | L_0, A_0 = a_0\right)$$

where the first and third equality follow from the law of iterated expectation, the second and fourth by conditional exchangeability (i.e., $Y^{a_1, a_0} \prod A_0 | L_0$ and $Y^{a_1, a_0} \prod A_1 | L_0, L_1, A_0$), the fifth by definition and the sixth by consistency. Similarly, we can derive (7) generically for any $t \in \mathcal{T}$. For the conditional expectations to be well-defined, one needs the positivity assumption.

A.2 | On Why the Weights (11) Lead to $E(Y_t | \bar{a}_t)$

Consider the recursive evaluation of the integral (12) from $s = t$ to $s = 0$, which shows how the weights recover the association:

$$\begin{aligned} \tilde{m}_{w,t}(a_t, \mathbf{h}_t) &= E(Y_t | A_t = a_t, \mathbf{H}_t = \mathbf{h}_t) \\ \tilde{m}_{w,t-1}(a_{t-1}, \mathbf{h}_{t-1}) &= \int E(Y_t | A_t = a_t, \mathbf{H}_t = \mathbf{h}_t) \\ &\quad dP_0(\mathbf{I}_t | A_t = a_t, A_{t-1} = a_{t-1}, \mathbf{H}_{t-1} = \mathbf{h}_{t-1}) \\ &= E[E(Y_t | A_t, \mathbf{H}_t) | A_t = a_t, A_{t-1} = a_{t-1}, \mathbf{H}_{t-1} = \mathbf{h}_{t-1}]) \\ &= E[Y_t | A_t = a_t, A_{t-1} = a_{t-1}, \mathbf{H}_{t-1} = \mathbf{h}_{t-1}]) \\ &\quad \dots \\ \tilde{m}_{w,0}(a_0, \mathbf{h}_0) &= E[Y_t | A_t = a_t, A_{t-1} = a_{t-1}, \dots, A_0 = a_0] = E(Y_t | \bar{a}_t) \end{aligned}$$

A.3 | Iterated Nested Expectation Representation of Estimand 2

Using the definition of (10), $\tilde{m}_{t+1} = Y_t$ and observing that $\int \tilde{m}_{t+1}(a_{t+1}, \mathbf{h}_{t+1}) w_t(a_{t+1}, \mathbf{h}_{t+1}) dP_0(\mathbf{I}_{t+1} | A_t = a_t, \mathbf{H}_t = \mathbf{h}_t) = E(\tilde{m}_{t+1}(a_{t+1}, \mathbf{h}_{t+1}) w_t(a_{t+1}, \mathbf{h}_{t+1}) | A_t = a_t, \mathbf{H}_t = \mathbf{h}_t)$, we have for $s = t$:

$$\tilde{m}_t(a_t, \mathbf{h}_t) = E(Y_t w_t | A_t = a_t, \mathbf{H}_t = \mathbf{h}_t)$$

where $w_t = w_t(a_{t+1}, \mathbf{h}_{t+1})$. For $s = t - 1$, we then get

$$\begin{aligned} \tilde{m}_{t-1}(a_{t-1}, \mathbf{h}_{t-1}) &= E(\tilde{m}_t(a_t, \mathbf{h}_t) w_{t-1} | A_{t-1} = a_{t-1}, \mathbf{H}_{t-1} = \mathbf{h}_{t-1}) \\ &= E\left(E\left(Y_t w_t | \bar{A}_t = \bar{a}_t, \bar{L}_t = \bar{l}_t\right) w_{t-1} | \bar{A}_{t-1} = \bar{a}_{t-1}, \mathbf{H}_{t-1} = \mathbf{h}_{t-1}\right) \end{aligned}$$

Reevaluating the integral from $s = t - 2, \dots, 0$ then leads to

$$\begin{aligned} \tilde{m}_0(a_0, \mathbf{h}_0) &= \mathbb{E}(\mathbb{E}(\dots \mathbb{E}(E(Y_t w_t | \bar{A}_t = \bar{a}_t, \bar{L}_t) w_{t-1} | \\ &\quad \bar{A}_{t-1} = \bar{a}_{t-1}, \bar{\mathbf{L}}_{t-1}) \dots) w_0 | A_0 = a_0, \mathbf{L}_0)) \end{aligned}$$

A.4 | Parametric g -Formula Representation

The weighted estimand, shown in the iterated weighted outcome representation in (16) can, alternatively, also be rewritten as

$$\tilde{m}_{w,t} = \int_{\bar{l}_t} \left\{ \mathbb{E}\left(w_t Y_t | \bar{A}_t = \bar{a}_t, \bar{L}_t = \bar{l}_t\right) \times \prod_{r=0}^t f\left(\mathbf{I}_r | \bar{\mathbf{A}}_{r-1} = \bar{a}_{r-1}, \bar{\mathbf{L}}_{r-1} = \bar{l}_{r-1}\right) \right\} d\bar{\mathbf{l}} \quad (\text{A.3})$$

where the second term in (A.3) is

$$\begin{aligned} \prod_{s=1}^{q-1} f\left(l_r^s | \bar{\mathbf{A}}_{r-1} = \bar{a}_{r-1}, \bar{\mathbf{L}}_{r-1} = \bar{l}_{r-1}, L_r^1 = l_r^1, \dots, L_r^{s-1} = l_r^{s-1}\right) \\ \times f\left(w_r Y_r | \bar{\mathbf{A}}_r = \bar{a}_r, \bar{\mathbf{L}}_r = \bar{l}_r\right) \end{aligned}$$

Note that in the above notation \bar{L}_s includes past outcomes, multiplied with the respective weights.

The equality follows from the knowledge that iterated nested expectation representations of the g-formula can be reexpressed as a traditional (parametric) g-formula factorization because both representations essentially standardize with respect to the post intervention distribution of the time-dependent confounders [3, 6, 42]. In our case, the outcome is $w_t Y_t$.

A.5 | Time-to-Event Considerations

Suppose we have a time-ordering of $(L_s, A_{s,1}, A_{s,2}, \dots, A_{s,p}, C_s, Y_s)$, $s = 0, \dots, t$. If we use weights to target $E(Y_t^{C_t=0} | A_t = a_t, \dots, A_0 = a_0)$ for units that do not have enough conditional support for the intervention strategy of interest, that is, $\bar{A}_t^{(j)}$, we have to modify the estimation algorithm of Table 1 as follows:

Step 1a/b	Estimate the densities among those uncensored, and without prior event ($C_t = 0, Y_{t-1} = 0$).
Step 2	Also, set $c_t = 0$.
Step 4	Fit the model for the weighted outcome among those uncensored and without prior event, that is, estimate $E(w_t (a_t^{(1)} c) \bar{Y}_t \bar{A}_t, C_t = 0, Y_{t-1} = 0, \mathbf{H}_t)$.
Step 5	Additionally, set $\bar{Y}_{t-1,i} = 1$, if $Y_{t-2,i} = 1$.

A.6 | Example Code for the Implementation of the Algorithm of Table 1

```
# load example efavirenz data from package CICI, select 2 time points for illustration
# variable order: sex metabolic log_age NRTI weight.0 efv.0 VL.0 adherence.1 weight.1 efv.1 VL.1
library(CICI); library(CICIpplus); data(EFV); EFV.2 <- EFV[,1:11]

# manual implementation of estimand 2

# Step 0a: interventions are (efv.0, efv.1) = [(0,0), (0.5,0.5), (1,1)]
# Step 0b: define \tilde{Y}_1 = Y_1 = VL.1
Y.tilde.1 <- EFV.2$VL.1

### iteration: t=1
# Step 1a: estimate numerator density g(efv.1 | past)
#           (assuming normality for illustration)
fitted.n.1 <- lm(efv.1 ~ ., data=subset(EFV.2, select=-VL.1))
# Step 1b: estimate denominator density g(efv.1 | past up to efv.0)
fitted.d.1 <- lm(efv.1 ~ sex+metabolic+log_age+NRTI+weight.0+efv.0, data=EFV.2)
# Step 2: set a_1=a_0=0
EFV.2.A <- EFV.2; EFV.2.A[,c("efv.0","efv.1")] <- 0
# Step 3a: evaluate densities at a=0
eval.n.1 <- dnorm(0, mean=predict(fitted.n.1, newdata=EFV.2.A), sd=summary(fitted.n.1)$sigma)
eval.d.1 <- dnorm(0, mean=predict(fitted.d.1, newdata=EFV.2.A), sd=summary(fitted.d.1)$sigma)
# Step 3b: estimate weights with c=0.01
wf <- function(num, den, c) as.numeric(num > c) + as.numeric(num <= c) * (num/den)
w.1 <- wf(eval.n.1, eval.d.1, c=0.01)
# Step 4: estimate iterated weighted outcome regression
# (binomial as all wY are 0/1 in this specific case, otherwise gaussian)
mY.1 <- glm(I(Y.tilde.1*w.1) ~ ., data=subset(EFV.2, select=-VL.1), family="binomial")
# Step 5: predict iterated weighted outcome under intervention
Y.tilde.0 <- predict(mY.1, newdata=EFV.2.A, type="response")

### iteration: t=0
# Step 1a: estimate numerator density g(efv.0 | past)
fitted.n.0 <- lm(efv.0 ~ sex+metabolic+log_age+NRTI+weight.0, data=EFV.2)
# Step 1b: estimate denominator density g(efv.0)
fitted.d.0 <- lm(efv.0 ~ 1, data=EFV.2)
# Step 2: set a_0=0 (done above already)
# Step 3a: evaluate densities at a=0
eval.n.0 <- dnorm(0, mean=predict(fitted.n.0, newdata=EFV.2.A), sd=summary(fitted.n.0)$sigma)
eval.d.0 <- dnorm(0, mean=predict(fitted.d.0, newdata=EFV.2.A), sd=summary(fitted.d.0)$sigma)
# Step 3b: estimate weights with c=0.01
w.0 <- wf(eval.n.0, eval.d.0, c=0.01)
# Step 4: estimate iterated (unweighted) outcome regression
mY.0 <- glm(Y.tilde.0 ~ sex+metabolic+log_age+NRTI+weight.0+efv.0, data=EFV.2)
# Step 5: final estimate under intervention
estimate <- weighted.mean(predict(mY.0, newdata=EFV.2.A), w=w.0)


# Step 6: Repeat steps 1-5 for a=(0.5,0.5) and a=(1,1)
# Step 7: Bootstrapping

# Note: the online material contains code on how to arrive at exactly the same result with the package
# Note: typically super learning is used to estimate the iterated outcome regressions, and the
#       densities are preferably estimated non-parametrically. The package offers these options.
```


Appendix B

Additional Results


B.1 | Simulation and Analysis Results



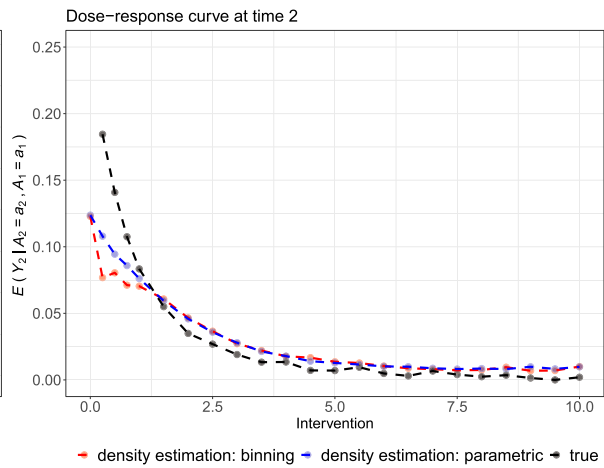
(a) Simulation 1: weighted curve, $c = 1, t = 2$



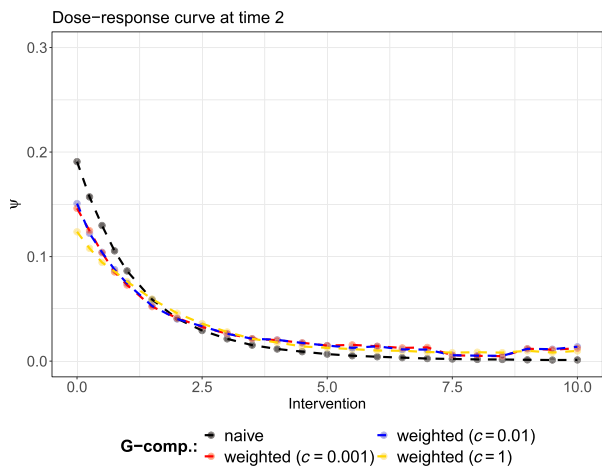
(b) Simulation 2: true and estimated CDRC at $t = 2$



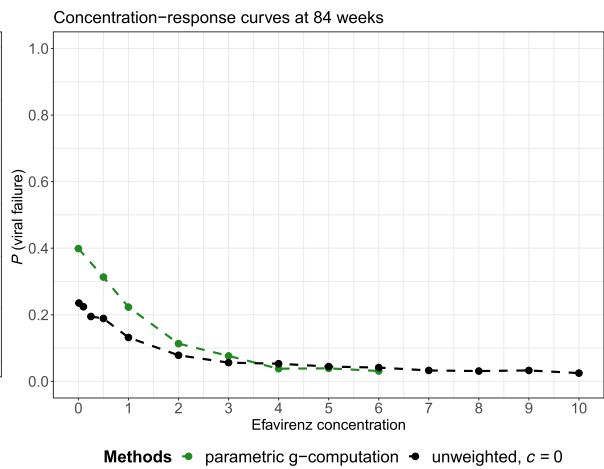
(c) Simulation 3: true and estimated CDRC at $t = 2$



(d) Simulation 3: weighted curve, $c = 1, t = 2$



(e) Simulation 3: CDRC's, $t = 2, c = 0.001, 0.01, 1$



(f) Data analysis: estimated CDRC

FIGURE B1 | Additional simulation and analysis results.

B.2 | Intervention Support

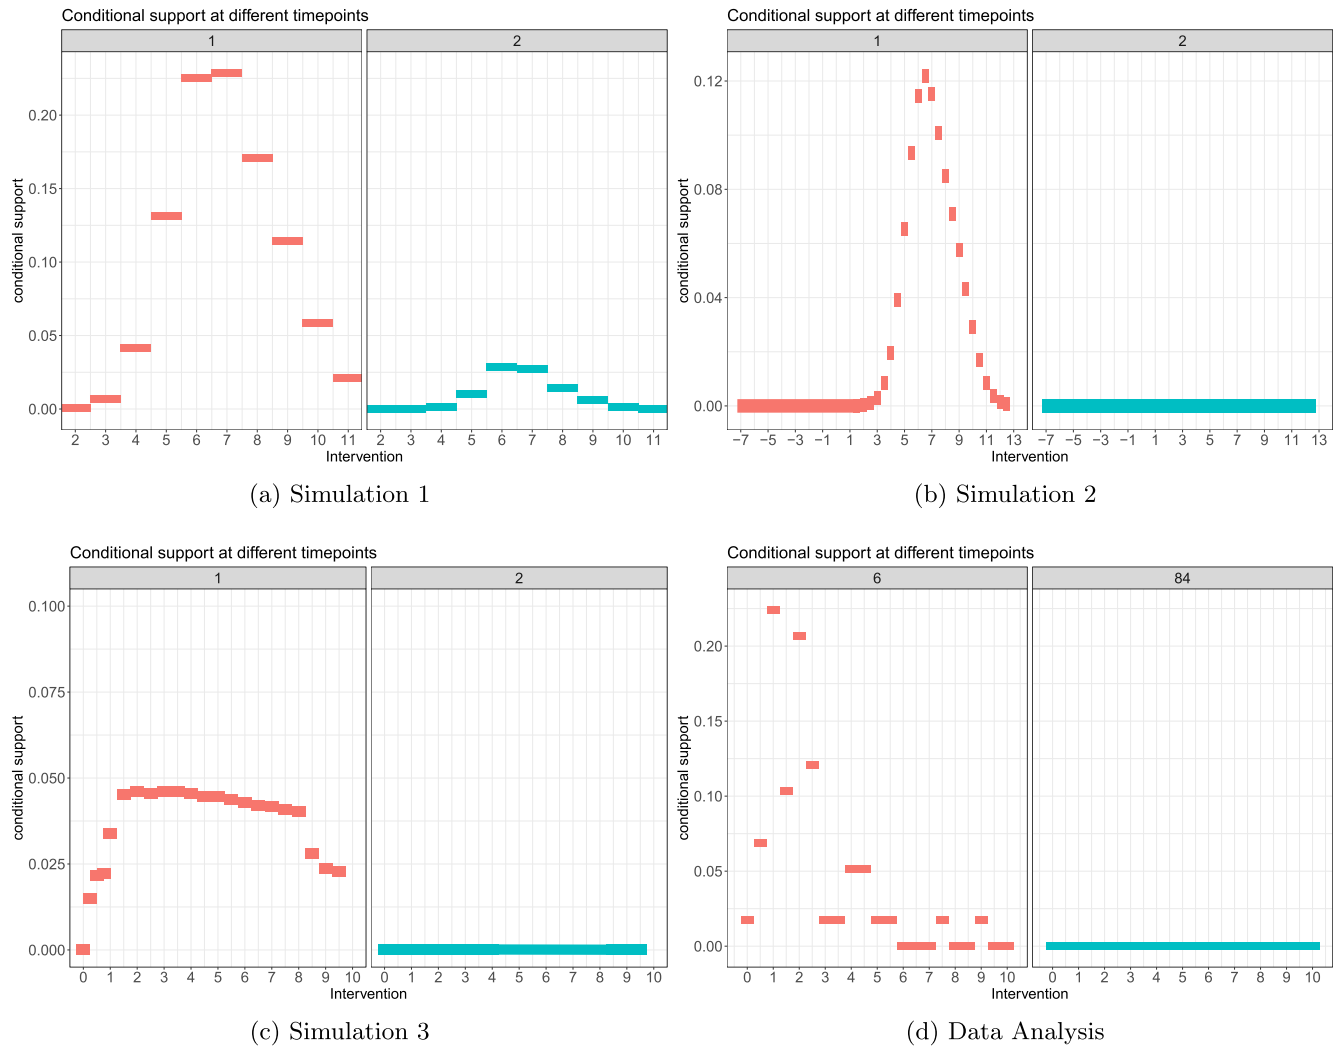


FIGURE B2 | Conditional support for intervention strategies of interest, for both the simulation settings and the data analysis.

Appendix C

Data-Generating Processes

C.1 | DGP for Simulation 1

For $t = 0$:

$$L_0^1 \sim B(p = 0.3)$$

$$L_0^2 \sim N(\mu = -1 + 2 \times L_0^1, \sigma = 1)$$

$$A_0 \sim N(\mu = 7 + L_0^1 + 0.7 \times L_0^2, \sigma = 1)$$

$$Y_0 \sim N(\mu = -1 + 0.5 \times A_0 + 0.5 \times L_0^2, \sigma = 1)$$

For $t = 1, 2$:

$$L_t^1 \sim B(p = 1/(1 + \exp(-(-4 + L_{t-1}^1 + 0.15 \times L_{t-1}^2 + 0.15 \times A_{t-1}))))$$

$$L_t^2 \sim N(\mu = 0.5 \times L_t^1 + 0.25 \times L_{t-1}^2 + 0.5 \times A_{t-1}, \sigma = 1)$$

$$A_t \sim N(\mu = A_{t-1} + L_t^1 - 0.1 \times L_t^2, \sigma = 0.5)$$

$$Y_t \sim N(\mu = -2 + 0.25 \times A_t - 0.2 \times L_t^1 + L_t^2, \sigma = 0.5)$$

C.2 | DGP for Simulation 2

For $t = 0$:

$$L_0^1 \sim B(p = 0.3)$$

$$L_0^2 \sim N(\mu = -1 + 2 \times L_0^1, \sigma = 1)$$

$$A_0 \sim N(\mu = 7 + L_0^1 + 0.7 \times L_0^2, \sigma = 1)$$

$$Y_0 \sim B(p = 1/(1 + \exp(-(-4 + 0.2 \times A_0 + 0.5 \times L_0^2))))$$

For $t = 1, 2, 3, 4$:

$$L_t^1 \sim B(p = 1/(1 + \exp(-(-4 + L_{t-1}^1 + 0.15 \times L_{t-1}^2 + 0.15 \times A_{t-1}))))$$

$$L_t^2 \sim N(\mu = 0.5 \times L_t^1 + 0.25 \times L_{t-1}^2 + 0.5 \times A_{t-1}, \sigma = 1)$$

$$A_t \sim N(\mu = -2 + 0.5 \times A_{t-1} + 0.75 \times L_t^1 + 0.35 \times L_t^2, \sigma = 0.5)$$

$$C_t \sim B(p = 1/(1 + \exp(-(-2 + 0.5 \times L_t^1 + 0.2 \times A_t))))$$

$$Y_t \sim B(p = 1/(1 + \exp(-(-4 + 0.2 \times A_t + 0.5 \times L_t^2))))$$

C.3 | DGP for Simulation 3

Both baseline data ($t = 0$) and follow-up data ($t = 1, \dots, 4$) were created using structural equations using the *R*-package *simcausal*. The below listed distributions, listed in temporal order, describe the data-generating process. Our baseline data consists of sex, genotype, log(age), log(weight) and the respective Nucleoside Reverse Transcriptase Inhibitor (NRTI). Time-varying variables are co-morbidities (CM), dose, efavirenz mid-dose concentration (EFV), elevated viral load (= viral failure, VL) and adherence (measured through memory caps, MEMS), respectively. In addition to Bernoulli (*B*), Multinomial (*MN*) and Normal (*N*) distributions, we also use truncated normal distributions; they are denoted by

$N_{[a,a_1,a_2,b,b_1,b_2]}$, where a and b are the truncation levels. Values which are smaller than a are replaced by a random draw from a $U(a_1, a_2)$ distribution and values greater than b are drawn from a $U(b_1, b_2)$ distribution, where U refers to a continuous uniform distribution. For the specified multinomial distributions, probabilities are normalized, if required, such that they add up to 1. The data-generating process reflects the following considerations: more complexity than the first two simulation settings in terms of distribution shape and variety, as well as non-linearities; similarity to the assumed DGP in the data analysis; generation of both areas of poor and good conditional intervention support, such that the proposed weighting scheme can be evaluated in all its breadth.

For $t = 0$:

$$\begin{aligned} \text{Sex}_0 &\sim B(p = 0.5) \\ \text{Genotype}_0 &\sim \text{MN} \left(\begin{array}{l} p1 = 1/(1 + \exp(-(-0.103 + I(\text{Sex}_0 = 1) \times 0.223 + I(\text{Sex}_0 = 0) \times 0.173))) \\ p2 = 1/(1 + \exp(-(-0.086 + I(\text{Sex}_0 = 1) \times 0.198 + I(\text{Sex}_0 = 0) \times 0.214))) \\ p3 = 1/(1 + \exp(-(-0.090 + I(\text{Sex}_0 = 1) \times 0.082 + I(\text{Sex}_0 = 0) \times 1.070))) \end{array} \right) \\ \text{Age}_0 &\sim N_{[0.693,0.693,1.2,8,2,7,2,8]}(\mu = 1.501, \sigma = 0.369) \\ \text{Weight}_0 &\sim N_{[2.26,2.26,2.67,3.37,3.02,3,37]}(\mu = (1.5 + 0.2 \times \text{Sex} + 0.774 \times \text{Age}) \times 0.94, \sigma = 0.369) \\ \text{NRTI}_0 &\sim \text{MN} \left(\begin{array}{l} p1 = 1/(1 + \exp(-(-0.006 + I(\text{Age}_0 > 1.4563) \times \text{Age}_0 \times 0.1735 + I(\text{Age}_0 \leq 1.4563) \times \text{Age}_0 \times 0.1570))) \\ p2 = 1/(1 + \exp(-(-0.006 + I(\text{Age}_0 > 1.4563) \times \text{Age}_0 \times 0.1735 + I(\text{Age}_0 \leq 1.4563) \times \text{Age}_0 \times 0.1570))) \\ p3 = 1/(1 + \exp(-(-0.006 + I(\text{Age}_0 > 1.4563) \times \text{Age}_0 \times 0.1570 + I(\text{Age}_0 \leq 1.4563) \times \text{Age}_0 \times 0.1818))) \end{array} \right) \\ \text{CM}_0 &\sim B(p = 0.15) \\ \text{Dose}_0 &\sim \text{MN} \left(\begin{array}{l} p1 = 1/\left(1 + \exp\left(-\left(5 + \sqrt{(\text{Weight}_0)} \times 8 - \text{Age}_0 \times 10\right)\right)\right) \\ p2 = 1/\left(1 + \exp\left(-\left(4 + \sqrt{(\text{Weight}_0)} \times 8.768 - \text{Age}_0 \times 9.06\right)\right)\right) \\ p3 = 1/\left(1 + \exp\left(-\left(3 + \sqrt{(\text{Weight}_0)} \times 6.562 - \text{Age}_0 \times 8.325\right)\right)\right) \\ p4 = 1 - (p1 + p2 + p3) \end{array} \right) \\ \text{EFV}_0 &\sim N_{[0.2032,0.2032,0.88,21,8,376,21]}(\mu = -8 + \text{Age}_0 \times 0.1 + \text{Genotype}_0 \times 4.66 + \text{Dose}_0 \times 0.1 + \\ &\quad I(\text{Genotype}_0 \leq 2) \times 2.66 + I(\text{Genotype}_0 = 3) \times 4.6, \sigma = 4.06) \\ \text{VL}_0 &\sim B\left(p = 1 - \left(1/\left(1 + \exp\left(-\left(0.4 + 1.9 \times \sqrt{\text{EFV}_0}\right)\right)\right)\right)\right) \end{aligned}$$

For $t \geq 1$:

$$\begin{aligned} \text{MEMS}_t &\sim B(p = 1/(1 + \exp(-0.71 + \text{CM}_{t-1} \times 0.31 + \text{MEMS}_{t-1} \times I(t \geq 2) \times 0.31))) \\ &\quad \text{with MEMS}_{t-1} \text{ used as of } t = 2 \\ \text{Weight}_t &\sim N_{[2.26,2.26,2.473,3.37,3,2,3,37]}(\mu = \text{Weight}_{t-1} \times 1.04 - 0.05 \times I(\text{CM}_{t-1} = 1), \sigma = 0.4) \\ \text{CM}_t &\sim B(p = 1 - (1/(1 + \exp(-0.5 \times I(\text{CM}_{t-1} = 1) + \text{Age}_0 \times 0.1 + \text{Weight}_{t-1} \times 0.1)))) \\ \text{Dose}_t &\sim \text{MN} \left(\begin{array}{l} p1 = (1/(1 + \exp(-\left(4 + \text{Dose}_{t-1} \times 0.5 + \sqrt{\text{Weight}_t} \times 4 - \text{Age}_0 \times 10\right)))) \\ p2 = (1/(1 + \exp(-\left(-8 + \text{Dose}_{t-1} \times 0.5 + \sqrt{\text{Weight}_t} \times 8.568 - \text{Age}_0 \times 9.06\right)))) \\ p3 = (1/(1 + \exp(-\left(20 + \text{Dose}_{t-1} \times 0.5 + \sqrt{\text{Weight}_t} \times 6.562 - \text{Age}_0 \times 18.325\right)))) \\ p4 = 1 - (p1 + p2 + p3) \end{array} \right) \\ \text{EFV}_t &\sim N_{[0.2032,0.2032,0.88,21,84,8,37,21,84]}(\mu = 0.1 \times \text{DOSE}_t + 0.1 \times \text{MEMS}_t + I(\text{Genotype}_0 \leq 2) \times 2.66 + I(\text{Genotype}_0 = 3) \times 4.6, \sigma = 4.06) \\ \text{VL}_t &\sim B\left(p = 1 - \left(1/\left(1 + \exp\left(-\left(1 - 0.6 \times I(t = 1) - 1.2 \times I(t = 4) + 0.1 \times \text{CM}_{t-1} + (2 - 0.2 \times I(t = 3)) \times \sqrt{\text{EFV}_t}\right)\right)\right)\right) \end{aligned}$$

Appendix D

More on the DAG

The measured efavirenz concentration depends on the following factors: the dose itself (which is recommended to be assigned based on the weight-bands), adherence (without regular drug intake, the concentration becomes lower), and the metabolism characterized in the gene CYP2B6, through the 516G and 983T polymorphisms [43]. Given the short half-life of the drug relative to the measurement interval, no arrow from EFV_t to EFV_{t+1} is required. Viral failure is essentially caused if there is not enough drug concentration in the body; and there might be interactions with co-morbidities and co-medications. Note also that co-morbidities, which are reflected in the DAG, are less frequent in the given data analysis as trial inclusion criteria did not allow children with active infections, treated for tuberculosis and laboratory abnormalities to be enrolled into the study. Both weight and MEMS (= adherence) are assumed to be time-dependent confounders affected by prior treatments (= concentrations). Weight affects the concentration indirectly through the dosing, whereas adherence affects it directly. Adherence itself is affected by prior concentrations, as too high concentration values can cause nightmares and other central nervous system side effects, or strong discomfort, that might affect adherence patterns. Weight is affected from prior concentration trajectories through the pathway of viral load and co-morbidities. Finally, both weight and adherence affect viral outcomes not only through EFV concentrations, but potentially also through co-morbidities such as malnutrition, pneumonia and others.