



Published in final edited form as:

Biom J. 2023 June ; 65(5): e2100359. doi:10.1002/bimj.202100359.

Monte Carlo sensitivity analysis for unmeasured confounding in dynamic treatment regimes

Eric J. Rose^{1,2}, Erica E. M. Moodie¹, Susan M. Shortreed^{3,4}

¹Department of Epidemiology and Biostatistics, McGill University, Montreal, QC, Canada

²Department of Epidemiology and Biostatistics, University at Albany, Rensselaer, New York, USA

³Kaiser Permanente Washington Health Research Institute, Seattle, Washington, USA

⁴Department of Biostatistics, University of Washington, Seattle, Washington, USA

Abstract

Data-driven methods for personalizing treatment assignment have garnered much attention from clinicians and researchers. Dynamic treatment regimes formalize this through a sequence of decision rules that map individual patient characteristics to a recommended treatment. Observational studies are commonly used for estimating dynamic treatment regimes due to the potentially prohibitive costs of conducting sequential multiple assignment randomized trials. However, estimating a dynamic treatment regime from observational data can lead to bias in the estimated regime due to unmeasured confounding. Sensitivity analyses are useful for assessing how robust the conclusions of the study are to a potential unmeasured confounder. A Monte Carlo sensitivity analysis is a probabilistic approach that involves positing and sampling from distributions for the parameters governing the bias. We propose a method for performing a Monte Carlo sensitivity analysis of the bias due to unmeasured confounding in the estimation of dynamic treatment regimes. We demonstrate the performance of the proposed procedure with a simulation study and apply it to an observational study examining tailoring the use of antidepressant medication for reducing symptoms of depression using data from Kaiser Permanente Washington.

Keywords

adaptive treatment strategies; bias; precision medicine

1 | INTRODUCTION

Precision medicine focuses on data-driven methods for personalizing treatment assignment to improve health care. Dynamic treatment regimes (DTRs) operationalize clinical decision-making through a sequence of functions that map individual patient characteristics to a

Correspondence Eric J. Rose, Department of Epidemiology and Biostatistics, University at Albany, Rensselaer, NY, 12144, USA. ejrose@albany.edu.

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

CONFLICT OF INTEREST STATEMENT

The authors have declared no conflict of interest.

recommended treatment (Chakraborty and Moodie, 2013; Tsiatis et al., 2019). An optimal DTR is one that maximizes the mean of some desirable measure of clinical outcome when applied to select treatment for all patients in the population of interest (Murphy, 2003; Robins, 2004). Optimal DTRs have been estimated to improve care for many different medical conditions including HIV (Cain et al., 2010; van der Laan et al., 2005), bipolar disorder (Wu et al., 2015; Zhang et al., 2018), diabetes (Ertefaie and Strawderman, 2018; Lockett et al., 2018), and cancer (Murray et al., 2018; Thall et al., 2000).

Data for estimating an optimal DTR are ideally collected through the use of a sequential multiple assignment randomized trial (SMART) (Lavori and Dawson, 2000, 2004; Murphy, 2005). However, longitudinal observational studies are frequently used due to the availability of data collected from electronic health records (EHRs) and the potentially prohibitive cost of conducting an SMART. These analyses rely on assuming that there is no unmeasured confounding; this assumption is not verifiable from the observed data. Not adjusting for all of the confounding variables can result in a biased estimate of the optimal DTR.

Sensitivity analyses have been used as a way to assess how the estimated effect would change if there was an unmeasured confounder. There has been much work on conducting sensitivity analyses for unmeasured confounding when estimating average causal effects from observational data beginning with Cornfield et al. (1959), in which they studied whether the causal link found between smoking and lung cancer could be due to unmeasured confounding. A wide array of different approaches have been proposed since. Rosenbaum and Rubin (1983) proposed a method for examining how sensitive the conclusions of a study with a binary outcome are to a binary unmeasured confounder by assessing what odds ratios between the unmeasured confounder and each of the treatment and the outcome would cause the effect to be no longer significant. Lin et al. (1998) proposed positing a regression model for the outcome containing the effect of the unmeasured confounder and probability distributions for the unmeasured confounder in each treatment group. Under certain assumptions, this allows for simple expressions of the bias in the treatment effect that can then be used to calculate treatment effects for a wide range of different effect sizes of the unmeasured confounder on the outcome and exposure.

An alternative to formulaic approaches are probabilistic sensitivity analyses, which posit probability distributions for the bias parameters and average over the distributions to obtain bias-adjusted estimates. There are two different existing approaches to implementing a probabilistic sensitivity analysis, Bayesian sensitivity analysis, and Monte Carlo sensitivity analysis (McCandless and Gustafson, 2017). A Bayesian sensitivity analysis (McCandless et al., 2007) works by positing prior distributions for the bias parameters and uses Bayes theorem to generate a posterior distribution for the causal effect that accounts for the uncertainty due to the unmeasured confounder. In a Monte Carlo sensitivity analysis (Greenland, 2005; Phillips, 2003; Steenland and Greenland, 2004), we posit prior distributions for the bias parameters in the same way and then draw values from the prior distributions and calculate a bias-adjusted estimate for the parameter of interest for each Monte Carlo repetition. This allows for a straightforward implementation that corrects bias while incorporating uncertainty due to the unmeasured confounding within a frequentist framework.

There has been some work on estimating DTRs in the presence of unmeasured confounding. This has mainly focused on using instrumental variables to estimate an optimal DTR (Chen and Zhang, 2021; Cui and Tchetgen Tchetgen, 2021; Liao et al., 2021; Qiu et al., 2021). An instrumental variable is a pretreatment variable that is correlated with treatment, independent of all unmeasured confounders, and has no direct effect on the outcome of interest. In some applications, there is not an obvious choice for an instrumental variable based on domain knowledge and it is impossible to empirically check whether a variable satisfies the assumptions for being an instrumental variable (Hernan and Robins, 2006). Zhang et al. (2021) proposed a method for ranking individualized treatment rules when there is unmeasured confounding by creating a partial order under the framework proposed by Rosenbaum (1987). Kallus and Zhou (2019) examined estimating a DTR that maximizes the value for the worst-case scenario of an uncertainty set that quantifies the degree of confounding that is unmeasured.

We propose using a Monte Carlo sensitivity analysis for the estimation of DTRs. Section 2 details setup and notation. In Section 3, we provide an overview of our proposed sensitivity analysis procedure. In Section 4, we examine the performance of our proposed method using a simulation study. Section 5 demonstrates applying the proposed procedure to EHRs when estimating DTRs to reduce symptoms of depression using data from Kaiser Permanente Washington (KPWA). A discussion of the proposed procedure is contained in Section 6.

2 | SETUP AND NOTATION

Suppose we have data from an observational study in which patients are treated throughout the course of K stages. We collect data of the form $\mathcal{D}_n = \{(X_{1,i}, A_{1,i}, \dots, X_{K,i}, A_{K,i}, Y_i)\}_{i=1}^n$, which consists of n *i.i.d.* replicates of $(X_1, A_1, \dots, X_K, A_K, Y)$ such that $X_1 \in \mathbb{R}^{p_1}$ denotes baseline patient characteristics, $A_k \in \{0, 1\}$ denotes the treatment assigned at stage k for $k = 1, \dots, K$, $X_k \in \mathbb{R}^{p_k}$ for $k = 2, \dots, K$ denotes patient information recorded during the course of the $(k-1)^{\text{st}}$ treatment, and $Y \in \mathbb{R}$ denotes the continuous patient outcome coded such that higher values are better. We let U denote an unmeasured time-fixed confounding variable that can be continuous or binary. This unmeasured confounder is measured at baseline and does not change over time, but influences treatment choices at any of the K stages as well as the outcome Y . Let $H_1 = X_1$ and $H_k = (X_1, A_1, \dots, A_{k-1}, X_k)$ for $k = 2, \dots, K$ denote the history of a patient such that H_k is all of the information available to a clinical decision maker at stage k . Let $d = (d_1, \dots, d_K)$ denote a DTR such that $d_k: \text{dom } H_k \rightarrow \text{dom } A_k$ for $k = 1, \dots, K$ is a function that maps a patient's history to a recommended treatment, where dom indicates the domain or set of possible treatments. Therefore, DTR d would recommend treatment $d_k(h_k)$ at stage k for a patient with history $H_k = h_k$.

To allow us to reference components of the history and regimes, we will use overbar notation to indicate up to stage k such that $\bar{\mathbf{x}}_k = (\mathbf{x}_1, \dots, \mathbf{x}_k)$, $\bar{\mathbf{a}}_k = (a_1, \dots, a_k)$, and $\bar{\mathbf{d}}_k = (d_1, \dots, d_k)$. We will suppress the subscript when denoting the entire sequence, for example, $\bar{\mathbf{x}} = \bar{\mathbf{x}}_K$. Similarly, we will use an underbar to denote treatments, covariates, and regimes after and including stage k such that $\underline{\mathbf{x}}_k = (\mathbf{x}_k, \dots, \mathbf{x}_K)$ and likewise for $\underline{\mathbf{a}}_k$ and $\underline{\mathbf{d}}_k$.

We will use the potential outcomes framework to define an optimal DTR (Rubin, 1978). Let $Y^*(\bar{a})$ denote the potential outcome under the treatment sequence $\bar{A} = \bar{a}$ and $H_k^*(\bar{a}_{k-1})$ denote the potential history at stage k under treatments $\bar{A}_{k-1} = \bar{a}_{k-1}$. Denote the set of all potential outcomes as $\mathcal{O}^* = \{H_2^*(a_1), H_3^*(\bar{a}_2), \dots, H_K^*(\bar{a}_{K-1}), Y^*(\bar{a}) : \bar{a} \in \{0, 1\}^K\}$. The potential outcome of a DTR, d , is then defined as.

$$Y^*(d) = \sum_{\bar{a} \in \{0, 1\}^K} Y^*(\bar{a}) \mathbb{1}\{d_1(\mathbf{h}_1) = a_1\} \prod_{k=2}^K \mathbb{1}[d_k\{H_k^*(a_{k-1})\} = a_k],$$

where $\mathbb{1}$ is the indicator function (Murphy, 2003). Define the value of a DTR by $V(d) = E\{Y^*(d)\}$. An optimal regime, d^{opt} , is defined as a regime that satisfies $V(d^{opt}) \geq V(d)$ for all regimes d .

When estimating DTRs, it is standard to make the following causal assumptions (Robins, 2004): (i) the stable unit treatment value assumption (SUTVA), $Y = Y^*(\bar{A})$ and $H_k = H_k^*(\bar{A}_{k-1})$ for $k = 2, \dots, K$ (ii) positivity, $P(A_k = a_k | H_k = \mathbf{h}_k) > 0$ with probability one for each $a_k \in \{0, 1\}$ for $k = 1, \dots, K$; and (iii) sequential ignorability, $\mathcal{O}^* \perp A_k | H_k$ for $k = 1, \dots, K$. We will assume SUTVA and positivity holds, but sequential ignorability does not hold because we have an unmeasured confounder, U . This is a key assumption and if it is violated, standard methods for estimating DTRs will lead to biased estimates of the optimal DTR. In addition, sequential ignorability is unverifiable from data and requires domain expertise to assess whether it is a reasonable assumption. Our proposed Monte Carlo sensitivity analysis requires making the following additional assumptions. We will assume that U is the only unmeasured confounder and $\mathcal{O}^* \perp A_k | U, H_k$ for $k = 1, \dots, K$. We will also assume that the unmeasured confounder has an additive effect on the outcome Y and is conditionally independent of later-stage covariates, so we have that $X_k \perp U | H_{k-1}, A_{k-1}$ for $k = 2, \dots, K$. These last two assumptions can be relaxed, though that would require more complex assumptions about the distribution and effect of the unmeasured confounder as well as a more complex expression for the bias.

2.1 | Dynamic weighted ordinary least squares

Dynamic weighted ordinary least squares (dWOLS) is a regression-based method for estimating DTRs (Wallace and Moodie, 2015). This method estimates an optimal DTR by estimating a contrast function that represents the difference in expected outcome of receiving a given treatment at stage k when compared to the reference treatment, which we take to be treatment $A_k = 0$, assuming that treatment is assigned optimally after stage k . This contrast function is referred to as the optimal blip-to-zero function, $\gamma(\mathbf{h}_k, a_k)$, and is defined as

$$\begin{aligned} \gamma_K(\mathbf{h}_K, a_K) &= E\{Y^*(\bar{a}_{K-1}, a_K) - Y^*(\bar{a}_{K-1}, 0) | H_K = \mathbf{h}_K\}, \\ \gamma_k(\mathbf{h}_k, a_k) &= E\{Y^*(\bar{a}_{k-1}, a_k, \underline{d}_{k+1}^{opt}) - Y^*(\bar{a}_{k-1}, 0, \underline{d}_{k+1}^{opt}) | H_k = \mathbf{h}_k\} \text{ for } k = 2, \dots, K-1, \\ \gamma_1(\mathbf{h}_1, a_1) &= E\{Y^*(a_1, \underline{d}_2^{opt}) - Y^*(0, \underline{d}_2^{opt}) | H_1 = \mathbf{h}_1\}. \end{aligned}$$

The optimal blip-to-zero function characterizes the optimal DTR (Robins, 2004). The optimal treatment rule at stage k is given by recommend treatment $A_k = 1$ if $\gamma_k(\mathbf{h}_k, 1) > 0$, otherwise recommend $A_k = 0$. In general, when considering more than two treatment options, the optimal DTR at stage k is given by $d_k^{\text{opt}}(\mathbf{h}_k) = \text{argmax}_{a_k} \gamma_k(\mathbf{h}_k, a_k)$. dWOLS estimates an optimal DTR by positing models for the blip functions and estimating their parameters through a series of weighted ordinary least squares regressions. Let $\gamma_k(\mathbf{h}_k, a_k; \psi_k)$ denote our posited model for the blip function. We will assume that the blip-to-zero function is correctly specified, so we have that $\gamma_k(\mathbf{h}_k, a_k; \psi_k^*) = \gamma_k(\mathbf{h}_k, a_k)$ for all k . The estimated optimal DTR is then given by $\hat{\mathbf{d}}^{\text{opt}} = (\hat{d}_1^{\text{opt}}, \dots, \hat{d}_K^{\text{opt}})$ such that $\hat{d}_k^{\text{opt}}(\mathbf{h}_k) = \text{argmax}_{a_k} \gamma_k(\mathbf{h}_k, a_k; \hat{\psi}_k)$.

Denote the treatment-free outcome at stage k by $G_k(\underline{\psi}_k)$ that represents a patient's observed outcome adjusted by the expected difference in a patient's outcome had they received treatment 0 at stage k and then treated optimally via the blip models with parameters $\underline{\psi}_k$ after stage k . Therefore, $G_k(\underline{\psi}_k) = Y - \gamma_k(\mathbf{h}_k, a_k; \psi_k) + \sum_{j=k+1}^K [\gamma_j\{\mathbf{h}_j, d_j^{\text{opt}}(\mathbf{h}_j); \psi_j\} - \gamma_j(\mathbf{h}_j, a_j; \psi_j)]$. We posit a model for the treatment-free outcome which we denote by $g_k(\mathbf{h}_k, a_k; \beta_k)$. We assume linear models for the treatment-free and blip models, so a model for the outcome at stage K is given by

$$\begin{aligned} E(Y \mid \mathbf{H}_K = \mathbf{h}_K, A_K = a_K; \beta_K, \psi_K) &= g_K(\mathbf{h}_K, a_K; \beta_K) + \gamma_K(\mathbf{h}_K, a_K; \psi_K) \\ &= \mathbf{h}_{K,\beta}^T \beta_K + a_K \mathbf{h}_{K,\psi}^T \psi_K, \end{aligned}$$

where $\mathbf{h}_{K,\beta}$ and $\mathbf{h}_{K,\psi}$ are components of \mathbf{h}_K with a leading one for the intercept and $\mathbf{h}_{K,\psi} \subset \mathbf{h}_{K,\beta}$. We also model the propensity score, $\pi_k(\mathbf{h}_k) = E(A_k \mid \mathbf{H}_k = \mathbf{h}_k)$, which we denote by $\pi_k(\mathbf{h}_k; \alpha_k)$ for $k = 1, \dots, K$. We estimate the blip model parameters using weighted ordinary least squares using weights $w_k(a_k, \mathbf{h}_k; \alpha_k)$ that satisfy $\pi_k(\mathbf{h}_k; \alpha_k) w_k(1, \mathbf{h}_k; \alpha_k) = \{1 - \pi_k(\mathbf{h}_k; \alpha_k)\} w_k(0, \mathbf{h}_k; \alpha_k)$. Examples of weights that satisfy this equality are given by the inverse probability of treatment weights (IPTWs), $a_k \{\pi(\mathbf{h}_k; \alpha_k)\}^{-1} + (1 - a_k) \{1 - \pi(\mathbf{h}_k; \alpha_k)\}^{-1}$ and overlap weights, $|a_k - \pi(\mathbf{h}_k; \alpha_k)|$. First, we estimate the blip parameters for the final stage K using our posited model for the outcome. Then, form the pseudo-outcome which we define as

$$\tilde{Y}_k = Y + \sum_{j=k+1}^K [\gamma_j\{\mathbf{h}_j, \hat{d}_j^{\text{opt}}(\mathbf{h}_j); \hat{\psi}_j\} - \gamma_j(\mathbf{h}_j, a_j; \hat{\psi}_j)].$$

At each stage starting with $k = K - 1$ and moving backwards, we model the pseudo-outcome by

$$E(\tilde{Y}_k \mid \mathbf{H}_k = \mathbf{h}_k, A_k = a_k; \beta_k, \psi_k) = g_k(\mathbf{h}_k, a_k; \beta_k) + \gamma_k(\mathbf{h}_k, a_k; \psi_k)$$

and fit the model using weighted least squares. The resulting estimate of ψ is consistent as long as the blip model is correctly specified and at least one of the treatment free and propensity score models is correctly specified. When we have an unmeasured confounder, U ,

both the treatment free and propensity score model will be misspecified leading to bias in the estimator of ψ .

3 | MONTE CARLO SENSITIVITY ANALYSIS

A Monte Carlo sensitivity analysis is a probabilistic approach to sensitivity analysis in which probability distributions are posited for the parameters governing the bias due to the unmeasured confounder. The posited probability distributions are repeatedly sampled from and used to calculate a series of bias corrected effect estimates (Greenland, 2005; McCandless and Gustafson, 2017). In what follows, we detail a Monte Carlo approach that seeks to posit models (and associated parameters) to describe the confounding due to the unmeasured variable U through its relationship to the outcome, the treatment, and other covariates, and to use that to impute the unmeasured confounder and hence assess and correct bias. Ideally, a secondary data set that contains data on the unmeasured confounder as well as the confounders, treatment, and outcomes measured in the current study with the potential unmeasured confounding is used to learn about the probability distributions for the bias parameters that must be posited. This allows us to fit the postulated unmeasured confounder models to the data from the secondary study and use the estimated parameters and standard errors to posit distributions that capture the uncertainty related to the unmeasured confounder. These data could be historical data or gathered from a concurrent or follow-up study and can be significantly smaller than the full study. If a secondary data set is unavailable, we can instead use subject matter expertise to posit probability distributions for the bias parameters with wide distributions to capture the additional uncertainty. For example, if we only know that the unmeasured confounder is positively associated with the outcome, we could posit a uniform distribution from 0 to a conservative estimate of the maximum for the presumed effect of the unmeasured confounder on the outcome. If there are bias parameters for which we have no information, we could posit a uniform distribution that is centered at 0 with conservative estimates for the bounds of the distribution. When noninformative distributions for the bias parameters are used, the bias-adjusted point estimate of the parameters indexing the optimal regime will be closer to the unadjusted estimate, but the confidence intervals will still capture the uncertainty in the optimal regime due to both sampling variability and unmeasured confounding.

Alternative approaches to sensitivity analyses for unmeasured confounding typically make assumptions about the association between the unmeasured confounder and treatment as well as the unmeasured confounder and outcome, often capturing these by a single sensitivity parameter (Groenwold et al., 2010; Hernan and Robins, 2020; Lash et al., 2009). When the bias due to unmeasured confounding is controlled by one or two parameters, a conditional analysis in which we examine the estimated effect for a fixed value of the bias parameters is appropriate. This type of analysis then allows for examining what values of the bias parameters would significantly change the conclusions of the study (Rosenbaum & Rubin, 1983). For the estimation of DTRs, we are interested in treatment interaction effects in addition to the main effect of treatment. To better quantify the bias in the interaction effects, we make additional assumptions about the association between the unmeasured confounder and the *measured covariates*. These assumptions lead to the bias being dictated

by many more parameters. The additional parameters make it very difficult to succinctly describe the potential bounds of the bias parameters with a conditional analysis, as the regions for which the results would be consistent could be very complex in addition to being multidimensional. By instead positing distributions for the bias parameters, we can quantify the uncertainty in the estimated DTR while accounting for the uncertainty in the posited models for the unmeasured confounder and can integrate over the posited distributions to calculate a bias-adjusted estimate.

Recall that we assume that the unmeasured confounder has an additive effect on the outcome Y in addition to linear models for the treatment-free and blip models. Therefore, a model for the outcome Y is given by

$$E(Y | \mathbf{H}_K = \mathbf{h}_K, A_K = a_K, U = u; \beta_K, \psi_K, \beta_u) = \mathbf{h}_{K,\beta}^T \beta_K + \beta_u u + a_K \mathbf{h}_{K,\psi}^T \psi_K.$$

We will also posit a model for the conditional mean of the unmeasured confounder given by $E(U | \mathbf{H}_K = \mathbf{h}_K, A_K = a_K; \zeta)$ such that ζ denotes the parameters of the model. This will typically be given by a linear or logistic regression model depending on whether U is continuous or binary, but there are no restrictions on the form of this model. We then have that β_u and ζ are the parameters that dictate the bias due to the unmeasured confounding with ζ indexing the model for the conditional mean of U , and β_u controlling the effect of the unmeasured confounder on the outcome Y .

Let $\mathbf{B}_k = (\mathbf{H}_{k,\beta}^T, A_k \mathbf{H}_{k,\psi}^T)^T$ and $\hat{w}_k = w_k(A_k, \mathbf{H}_k; \hat{\alpha}_k)$ for $k = 1, \dots, K$. If the unmeasured confounder U is not included in the regression, and then the estimated coefficients are given by

$$\begin{pmatrix} \hat{\beta}_K \\ \hat{\psi}_K \end{pmatrix} = \{\mathbb{P}_n(\hat{w}_K \mathbf{B}_K \mathbf{B}_K^T)\}^{-1} \mathbb{P}_n(\hat{w}_K \mathbf{B}_K Y)$$

such that \mathbb{P}_n denotes the empirical expectation. Recall we assume that the unmeasured confounder, U , has an additive effect on the outcome and let β_u denote the coefficient for the unmeasured confounder. In Section A of the Supporting Information, we show that the bias is given by

$$\begin{aligned} \text{bias}(\hat{\psi}_K) &= E(\hat{\psi}_K) - \psi_K \\ &= \left[-E(\hat{w}_K A_K \mathbf{H}_{K,\psi} \mathbf{H}_{K,\psi}^T) \{E(\hat{w}_K \mathbf{H}_{K,\beta} \mathbf{H}_{K,\beta}^T)\}^{-1} E(\hat{w}_K A_K \mathbf{H}_{K,\beta} \mathbf{H}_{K,\psi}^T) \right. \\ &\quad \left. + E(\hat{w}_K A_K \mathbf{H}_{K,\psi} \mathbf{H}_{K,\psi}^T) \right]^{-1} \left[E(\hat{w}_K A_K \mathbf{H}_{K,\psi} \beta_u U) \right. \\ &\quad \left. - E(\hat{w}_K A_K \mathbf{H}_{K,\psi} \mathbf{H}_{K,\beta}^T) \{E(\hat{w}_K \mathbf{H}_{K,\beta} \mathbf{H}_{K,\beta}^T)\}^{-1} E(\hat{w}_K \mathbf{H}_{K,\beta} \beta_u U) \right]. \end{aligned} \tag{1}$$

We can estimate the bias in finite samples by replacing the expectations with empirical expectations and imputing an estimate of $E(U | \mathbf{H}_K, A_K)$ for U which we will denote by $\widehat{\text{bias}}(\hat{\psi}_K)$. Note that this is then a function of the data and the parameters of the unmeasured confounder models.

If we proceeded with dWOLS, the bias in the blip parameters would bias the pseudo-outcome given by $\tilde{Y}_{K-1} = Y + \gamma_K \{ \mathbf{h}_K, \hat{d}_K^{\text{opt}}(\mathbf{h}_K); \hat{\psi}_K \} - \gamma_K(\mathbf{h}_K, a_K; \hat{\psi}_K)$. Additionally, estimating the blip model at stage $K - 1$ without accounting for the unmeasured confounder would introduce additional bias that further compounds as we move backward through the stages. We can use the bias-corrected estimate of ψ to calculate the pseudo-outcome and, if we assume that $\mathbf{X}_k \perp U \mid \mathbf{H}_{k-1}, A_{k-1}$ for $k = 2, \dots, K$, we have that Equation (1) for $k = K$ holds for $k = 1, \dots, K$. This assumption can be relaxed at the expense of additional assumptions about the effect of U on \mathbf{X}_k and a more complex expression for the bias. Being able to accurately posit models and parameter distributions for U when a more complex relationship exists could be potentially difficult without data. Therefore, a different approach to sensitivity analysis may be more appropriate when secondary data are unavailable if this assumption is unreasonable.

ALGORITHM 1

Monte Carlo sensitivity analysis

-
- 1: Posit a model for $E(U \mid \mathbf{H}_K = \mathbf{h}_K, A_K = a_K)$ given by $E(U \mid \mathbf{H}_K = \mathbf{h}_K, A_K = a_K; \zeta)$
 - 2: Posit distributions $f(\zeta)$ and $f(\beta_u)$ for the bias parameters ζ and β_u
 - 3: **for** $b \in \{1, \dots, B\}$ Monte Carlo/bootstrap repetitions **do**
 - 4: Draw a bootstrap sample from \mathcal{D}_n
 - 5: Sample $\zeta^{(b)}, \beta_u^{(b)}$ from the distributions posited in step 2
 - 6: Impute $\hat{U}^{(b)}$ from $E(U \mid \mathbf{H}_K = \mathbf{h}_K, A_K = a_K; \zeta^{(b)})$
 - 7: Estimate $\hat{\alpha}_K^{(b)}$ with the propensity score model and calculate weights $\hat{W}_K^{(b)}$
 - 8: Estimate $\hat{\psi}_K^{(b)}$ using weighted least squares, i.e. perform a dWOLS estimation using the bootstrap sample
 - 9: Calculate bias in $\hat{\psi}_K^{(b)}$ with equation 1 using $\hat{U}^{(b)}$ and $\beta_u^{(b)}$ as estimates for U and β_u
 - 10: Calculate bias-adjusted estimate $\hat{\psi}_K^{\text{adj},(b)} = \hat{\psi}_K^{(b)} - \widehat{\text{bias}}(\hat{\psi}_K^{(b)})$
 - 11: **for** $k = K - 1, \dots, 1$ **do**
 - 12: Calculate pseudo-outcome

$$\tilde{Y}_k^{(b)} = Y + \sum_{j=k+1}^K [\gamma_j \{ \mathbf{h}_j, \hat{d}_j^{\text{opt}}(\mathbf{h}_j); \hat{\psi}_j^{\text{adj},(b)} \} - \gamma_j(\mathbf{h}_j, a_j; \hat{\psi}_j^{\text{adj},(b)})]$$
 - 13: Estimate $\hat{\alpha}_k^{(b)}$ with the propensity score model and calculate weights $\hat{W}_k^{(b)}$
 - 14: Estimate $\hat{\psi}_k^{(b)}$ using weighted least squares
 - 15: Calculate bias in $\hat{\psi}_k^{(b)}$ with equation 1 using $\hat{U}^{(b)}$ and $\beta_u^{(b)}$ as estimates for U and β_u
 - 16: Calculate bias-adjusted estimate $\hat{\psi}_k^{\text{adj},(b)} = \hat{\psi}_k^{(b)} - \widehat{\text{bias}}(\hat{\psi}_k^{(b)})$
 - 17: **end for**
 - 18: **end for**
 - 19: Calculate bias-adjusted estimate $\hat{\psi}_k^{\text{adj}} = B^{-1} \sum_{b=1}^B \hat{\psi}_k^{\text{adj},(b)}$ for $k = 1, \dots, K$
-

The primary objective of conducting sensitivity analyses in this context is to quantify the uncertainty in the estimated DTR when there is bias due to unmeasured confounding. If we

generated confidence intervals for the bias by taking the percentiles of our adjusted results across the Monte Carlo samples, these interval estimates would only capture the variability due to the uncertainty in our bias parameters and would not account for the uncertainty in our estimate of the parameters of interest. Therefore, we take a bootstrap sample for each Monte Carlo repetition and confidence intervals are formed by taking percentiles of the bias-adjusted values of ψ . To account for nonregularity in estimators resulting from the nonsmoothness in the pseudo-outcome when estimating a multistage DTR, we will adapt the m -out-of- n bootstrap (Chakraborty et al., 2013).

We perform a Monte Carlo sensitivity analysis by first positing a model for the conditional mean of the unmeasured confounder given by $E(U | \mathbf{H}_K = \mathbf{h}_K, A_K = a_K; \zeta)$. We also posit distributions $f(\zeta)$ and $f(\beta_u)$ for the bias parameters ζ and β_u . We then draw B bootstrap samples which we will denote by $\mathcal{D}_n^{(b)}$ for $b = 1, \dots, B$. For each bootstrap sample, we sample $\zeta^{(b)}$ and $\beta_u^{(b)}$ from the posited distributions, $f(\zeta)$ and $f(\beta_u)$. We use the sampled value of $\zeta^{(b)}$ and the bootstrap data to impute values for the unmeasured confounder, U , which we denote by $\hat{U}^{(b)}$. Then use the bootstrap data, the sampled bias parameter $\beta_u^{(b)}$, and the imputed values of U to calculate a bias-adjusted estimate of $\psi^{(b)}$ which we will denote by $\hat{\psi}^{\text{adj},(b)}$. This is done by implementing dWOLS to estimate $\hat{\psi}_K^{(b)}$ and then using Equation (1) to calculate a bias-adjusted estimate, $\hat{\psi}_K^{\text{adj},(b)}$. This bias-adjusted estimate is then used to calculate the pseudo-outcome $\tilde{Y}_{K-1}^{(b)}$. We continue implementing dWOLS to estimate $\hat{\psi}_K^{\text{adj},(b)}$ at each stage by estimating $\hat{\psi}_K^{(b)}$ and using Equation (1) to adjust for the bias due to the unmeasured confounder. A bias-adjusted estimate for ψ_k is then given by $\hat{\psi}_k^{\text{adj}} = B^{-1} \sum_{b=1}^B \hat{\psi}_k^{\text{adj},(b)}$ for $k = 1, \dots, K$. A step-by-step procedure for generating bias-adjusted estimates for the parameters indexing the optimal regime, ψ , is given in Algorithm 1.

To construct a confidence interval for ψ for a single-stage DTR, we calculate the percentiles of the bias-adjusted estimates from each bootstrap sample. When the number of stages is greater than 1, we have that the pseudo-outcome is a nonsmooth function of the generative model, and therefore, the estimator for ψ_k for $k < K$ is nonregular (Laber et al., 2014; Moodie and Richardson, 2009; Robins, 2004). Therefore, we adapt the m -out-of- n bootstrap that is a method for producing confidence intervals for nonsmooth functionals in which we take bootstrap samples with a smaller resample size of $m < n$ (Bickel et al., 1997; Dumbgen, 1993; Shao, 1994; Swanepoel, 1986). Chakraborty et al. (2013) proposed using a resample size of $m = n \frac{1 + \kappa(1-p)}{1 + \kappa}$ such that κ is a tuning parameter and $p = P(\mathbf{H}_2^T \psi_2 = 0)$ to construct a confidence interval for ψ_1 when $K = 2$. The tuning parameter controls the minimum resample size because, for a fixed n , the resample size takes values within the interval $\left[n \frac{1}{1+x}, n \right]$.

Chakraborty et al. (2013) proposed using a double bootstrap procedure to choose the tuning parameter κ . The nonsmoothness in the pseudo-outcome occurs at and near the point(s) where $\mathbf{H}_2^T \psi_2 = 0$, so large values of p indicate a high degree of nonregularity. They propose using a plug-in estimator for p given by

$$\hat{p} = \mathbb{P}_n \mathbb{1} \left\{ n(\mathbf{H}_{2,\psi}^T \hat{\psi}_2)^2 \leq \mathbf{H}_{2,\psi}^T \hat{\Sigma}_{\psi_2} \mathbf{H}_{2,\psi} \chi_{1,1-\nu}^2 \right\},$$

where $\hat{\Sigma}_{\psi_2}$ denotes an estimate of $n\text{Cov}(\hat{\psi}_2, \hat{\psi}_2)$ and $\chi_{1,1-\nu}^2$ denotes the $(1-\nu) \times 100$ percentile of a chi-squared distribution with 1 degree of freedom. To generalize this to $K > 2$ stages, we define $p_k = P(\mathbf{H}_k^T \psi_k = 0)$ and construct an estimate, \hat{p}_k , for each stage (Rose et al., 2022). We first calculate a resample size $\hat{m}_k = n \frac{1 + \kappa(1 - \hat{p}_k)}{1 + \kappa}$ such that

$$\hat{p}_k = \mathbb{P}_n \mathbb{1} \left\{ n(\mathbf{H}_{k,\psi}^T \hat{\psi}_k^{\text{adj}})^2 \leq \mathbf{H}_{k,\psi}^T \hat{\Sigma}_{\psi_k} \mathbf{H}_{k,\psi} \chi_{1,1-\nu}^2 \right\}.$$

We then take bootstrap samples of size \hat{m}_k to construct a confidence interval for $\mathbf{h}_{K-1,\psi}^T \psi_{K-1}^{\text{adj}}$ for each $\mathbf{H}_{K-1,\psi} = \mathbf{h}_{K-1,\psi}$ in the sample. This confidence interval is created by estimating $\hat{\psi}_{K-1,\hat{m}_{K-1}}^{\text{adj},(b)}$ for each bootstrap sample by drawing from the posited distributions for the bias parameters and adjusting the estimated parameters using Equation (1) as before. This allows us to obtain an estimate, \hat{p}_{K-1} , of $p_{K-1} = P(\mathbf{H}_{K-1}^T \psi_{K-1} = 0)$ by calculating the proportion of individuals in the study for which the confidence interval for $\mathbf{h}_{K-1,\psi}^T \psi_{K-1}^{\text{adj}}$ contains zero. This procedure is repeated for $k = K-2, \dots, 1$. We then calculate $\hat{p} = \max_k \hat{p}_k$ and the corresponding resample size $\hat{m} = n \frac{1 + \kappa(1 - \hat{p})}{1 + \kappa}$. Finally, we take B bootstrap samples of size \hat{m} and estimate $\hat{\psi}_{k,\hat{m}}^{\text{adj},(b)}$ for each bootstrap sample by again sampling from the posited distributions for the bias parameters, β_u and ζ , and adjusting for the bias using Equation (1). To construct a $(1 - \vartheta_k) \times 100\%$ confidence region for ψ_k , calculate $\vartheta_k/2 \times 100$ and $(1 - \vartheta_k/2) \times 100$ percentiles of $\sqrt{\hat{m}}(\hat{\psi}_{k,\hat{m}}^{\text{adj},(b)} - \hat{\psi}_k^{\text{adj}})$, denoted as \hat{l}_k and \hat{u}_k , respectively, for $k = 1, \dots, K$. A $(1 - \vartheta_k) \times 100\%$ confidence region for ψ_k is then given by $(\hat{\psi}_k^{\text{adj}} - \hat{u}_k/\sqrt{\hat{m}}, \hat{\psi}_k^{\text{adj}} + \hat{l}_k/\sqrt{\hat{m}})$ for $k = 1, \dots, K$. A step-by-step procedure for constructing a confidence interval for ψ that accounts for uncertainty in the unmeasured confounder and random sampling variability is outlined in Algorithm 2.

4 | SIMULATION EXPERIMENTS

We conducted a series of simulation experiments to evaluate the effectiveness of the proposed method. We applied the proposed method to two different data-generating models, the first being a one-stage study and the second a two-stage study. Additional simulations are contained in Sections B and C of the Supporting Information with B containing simulations with a binary unmeasured confounder and C containing a different generative model in which the bias does not meaningfully affect the performance of the estimated regime. The data-generating model for the one-stage study was given by:

$$\begin{aligned}
U &\sim N(0, \sigma_u^2), \\
X_1 &= \phi_{10} + \phi_{11}U + \epsilon_{x_1}, & \epsilon_{x_1} &\sim N(0, \sigma_{x_1}^2), \\
X_2 &= \phi_{20} + \phi_{21}U + \epsilon_{x_2}, & \epsilon_{x_2} &\sim N(0, \sigma_{x_2}^2), \\
P(A_1 = 1 \mid \mathbf{X} = \mathbf{x}, U = u) &= [1 + \exp\{-(\alpha_0 + \alpha_1 x_1 + \alpha_2 x_2 + \alpha_3 u)\}]^{-1}, \\
Y &= \beta_0 + \beta_1 X_1 + \beta_2 X_2 + \beta_u U + A_1(\psi_0 + \psi_1 X_1 + \psi_2 X_2) + \epsilon_y, & \epsilon_y &\sim N(0, \sigma_y^2).
\end{aligned}$$

ALGORITHM 2

Monte Carlo sensitivity analysis: confidence intervals

-
- 1: Calculate $\hat{p}_K = \mathbb{P}_n \left\{ n(\mathbf{H}_{K,\psi}^T \hat{\psi}_K^{\text{adj}})^2 \leq \mathbf{H}_{K,\psi}^T \hat{\Sigma}_{\psi_K} \mathbf{H}_{K,\psi} \chi_{1,1-\nu}^2 \right\}$ where $\hat{\Sigma}_{\psi_K}$ denotes an estimate of $n\text{Cov}(\hat{\psi}_K, \hat{\psi}_K)$
 - 2: Calculate resample size $\hat{m}_K = n \frac{1 + \kappa(1 - \hat{p}_K)}{1 + \kappa}$
 - 3: **for** $k = K - 1, \dots, 1$ **do**
 - 4: Draw B bootstrap samples of size \hat{m}_k
 - 5: Estimate $\hat{\psi}_{k,\hat{m}_k}^{\text{adj},(b)}$ for $b = 1, \dots, B$
 - 6: For each given $\mathbf{H}_{k,\psi} = \mathbf{h}_{k,\psi}$, construct a confidence interval for $\mathbf{h}_{k,\psi}^T \psi_k^{\text{adj}}$
 - 7: Calculate \hat{p}_k by the proportion of confidence intervals for $\mathbf{h}_{k,\psi}^T \psi_k^{\text{adj}}$ that contain zero
 - 8: Calculate the resample size $\hat{m}_k = n \frac{1 + \kappa(1 - \hat{p}_k)}{1 + \kappa}$
 - 9: **end for**
 - 10: Calculate $\hat{p} = \max_k \hat{p}_k$ and $\hat{m} = n \frac{1 + \kappa(1 - \hat{p})}{1 + \kappa}$
 - 11: Draw B bootstrap sample of size \hat{m}
 - 12: Estimate $\hat{\psi}_{k,\hat{m}}^{\text{adj},(b)}$ for $k = K, \dots, 1, b = 1, \dots, B$
 - 13: Calculate $\vartheta_k/2 \times 100$ and $(1 - \vartheta_k/2) \times 100$ percentiles of $\sqrt{\hat{m}}(\hat{\psi}_{k,\hat{m}}^{\text{adj},(b)} - \hat{\psi}_k^{\text{adj}})$ which we denote by \hat{l}_k and \hat{u}_k respectively for $k = 1, \dots, K$
 - 14: Calculate a confidence interval for ψ_k by $(\hat{\psi}_k^{\text{adj}} - \hat{u}_k/\sqrt{\hat{m}}, \hat{\psi}_k^{\text{adj}} + \hat{l}_k/\sqrt{\hat{m}})$ for $k = K, \dots, 1$
-

We conducted 1000 repetitions for the simulation study. We let the sample size be given by $n = 1000$ and conducted $B = 500$ Monte Carlo repetitions for the sensitivity analysis. The parameter values were given by:

$$\begin{aligned}
\psi &= (\psi_0, \psi_1, \psi_2) = (-1, 0.5, 0.5), & \alpha &= (\alpha_0, \alpha_1, \alpha_2, \alpha_3) = (0, 1, 1, 2), \\
\phi_1 &= (\phi_{10}, \phi_{11}) = (0, 1), & \beta &= (\beta_0, \beta_1, \beta_2, \beta_u) = (1, 1, 1, 2), \\
\phi_2 &= (\phi_{20}, \phi_{21}) = (0, -1), & \sigma_u^2 &= \sigma_{x_1}^2 = \sigma_{x_2}^2 = \sigma_y^2 = 1.
\end{aligned}$$

We posited four different sets of distributions for the parameters in the unmeasured confounder models. This allowed us to assess how sensitive our proposed method is to bias in the distributions for these parameters. The four different simulation scenarios were given by: (i) narrow normal, centered properly; (ii) wide normal, centered properly; (iii) narrow normal, off-center; (iv) wide normal, off-center. For scenario (i),

we posited models given by $\beta_u^{(b)} \sim N(\beta_u, 0.1)$ and $\zeta_j^{(b)} \sim N(\zeta_j^*, 0.1)$ for $j = 0, \dots, 3$ such that $\zeta_j^* = \min_{\zeta} E\left\{(U - \zeta_0 - \zeta_1 X_1 - \zeta_2 X_2 - \zeta_3 A_1)^2\right\}$. For the wide distribution settings, the variance was increased to 0.5 and for the off-center scenarios, the distribution was centered at the true mean plus 0.1.

Figure 1 shows boxplots of the point estimates of ψ_0 , ψ_1 , and ψ_2 across 1000 repetitions when not adjusting and adjusting for the unmeasured confounder using a Monte Carlo sensitivity analysis with different distributions for the bias parameters. We can see that when we did not adjust for bias due to the unmeasured confounder, the estimate of ψ_0 is biased with a root mean squared error (rMSE) of 1.094. When the parameter distributions were centered on the true value, the adjusted estimate of ψ_0 was unbiased. There was some bias in the adjusted values of ψ_0 when the distribution is off-center as expected. We can see that for the narrow distribution, the rMSE increased from 0.125 to 0.259 and for the wide distribution, rMSE increased from 0.134 to 0.264. The unmeasured confounder did not cause bias in the estimation of ψ_1 and ψ_2 , so the adjusted estimates using a Monte Carlo sensitivity analysis were very similar to the unadjusted results.

Bias in the blip model parameters that index the optimal DTR can cause the estimated regime to differ from the true optimal regime. For each repetition of the simulation study, we simulated 10,000 new patients and assessed what proportion of the patients' treatment would match the recommended treatment of the true optimal DTR when following each one of the adjusted and unadjusted estimates of the optimal regime. Table 1 displays the proportion of patients who would receive the treatment recommend by the true optimal regime across all 1000 repetitions of the simulation study. We can see that the bias in ψ_0 causes the treatment recommended by the estimated regime to match the optimal regime 52.8% of the time. When adjusting for the unmeasured confounder, this increased to 95.6% in the ideal scenario of a narrow, centered parameter distribution and to 95.8% when we used a wide, off-centered parameter distributions.

In practice, if we do not have external data to help posit the mean model for the unmeasured confounder and distributions for the parameters, we will have less belief that the bias-adjusted DTR is close to the true optimal regime. The proposed method should be used to examine confidence intervals to assess whether the unmeasured confounder could be introducing bias to our estimated regimes. Table 2 displays the coverage and width of 95% confidence intervals generated using a Monte Carlo sensitivity analysis and from the unadjusted analysis that assumes no unmeasured confounding. For the unadjusted analysis, the empirical coverage of the confidence interval for ψ_0 was 0%. Even for ψ_1 and ψ_2 , which did not appear to have any bias in the estimation, the coverage was still below the nominal rate of 95%. For all of the different posited parameter distributions, Monte Carlo sensitivity analysis produced confidence intervals that attained or were close to the nominal coverage probability. The width of the confidence intervals for ψ_0 increased significantly when conducting sensitivity analysis, indicating that this parameter is sensitive to the unmeasured confounding. Note that the width of the confidence intervals increased when the variability in the parameter distributions was increased.

To compare our proposed procedure to an alternative method for sensitivity analysis, we conducted a simulation study for estimating an optimal DTR with G-estimation using the same single-stage generative model (Hernan and Robins, 2020; Robins et al., 1999). G-estimation for a single-stage study proceeds by defining $H(\psi) = Y - \gamma_1(\mathbf{H}_1, A_1; \psi)$ (Robins, 2004). $H(\psi)$ then represents a candidate potential outcome for $Y^*(0)$ for a given value of ψ . Consider the model

$$\text{logit}[P\{A_1 = 1 \mid H(\psi), \mathbf{H}_1 = \mathbf{h}_1\}] = \xi_0 + H(\psi)(\mathbf{h}_1^T \boldsymbol{\xi}_1) + \mathbf{h}_1^T \boldsymbol{\xi}_2. \quad (2)$$

We assume sequential ignorability holds, which implies that $\boldsymbol{\xi}_1 = \mathbf{0}$ for the true value of ψ where $\mathbf{0}$ denotes the zero vector. Therefore, to estimate the true value of ψ , we find the value of ψ in $H(\psi)$ that results in $\hat{\boldsymbol{\xi}}_1 = \mathbf{0}$ when fitting the logistic regression model in Equation (2). When there is unmeasured confounding, we have that $\boldsymbol{\xi}_1 \neq \mathbf{0}$. Therefore, we can conduct a sensitivity analysis by setting $\boldsymbol{\xi}_1$ to a range of different fixed values and estimating ψ for each to examine how sensitive the estimated DTR is to an unmeasured confounder.

This approach leads to bias parameters that are far more difficult to interpret than the bias parameters in our proposed Monte Carlo sensitivity analysis. A sensitivity analysis could conclude that a given value of $\boldsymbol{\xi}_1$ would significantly alter the conclusions of the study, but it is unclear what degree of confounding would result in that value of $\boldsymbol{\xi}_1$. To avoid this problem, we can use a similar approach to our proposed Monte Carlo sensitivity analysis by collecting data from a small, secondary study that includes the unmeasured confounder to learn about the true value of the bias parameters $\boldsymbol{\xi}_1$. First, we use the secondary study to estimate $\hat{\psi}$ using G-estimation with the unmeasured confounder included in the logistic regression model given in Equation (2). We refit the logistic regression model with $H(\hat{\psi})$ included and the unmeasured confounder removed to estimate $\boldsymbol{\xi}_1$. We then proceed in the same manner as the proposed Monte Carlo procedure. We posit distributions for $\boldsymbol{\xi}_1$ using the point estimate and standard error from the model fit to the secondary study. We then use the bootstrap, sample $\boldsymbol{\xi}_1$ from the posited distribution for each repetition, and calculate a bias-adjusted estimate of ψ for each Monte Carlo repetition. We then take the mean of the estimates for ψ to get a point estimate and take percentiles to construct a confidence interval.

We conducted a simulation study of this procedure using comparable distributions for the bias parameters to those used for the previous simulations. For this generative model, $\boldsymbol{\xi}_1 = (\xi_{10}, \xi_{11}, \xi_{12})$ is given by (0.52, 0, 0). For the simulation study, we will assume $\xi_{11} = 0$ and $\xi_{12} = 0$ and only posit a distribution for ξ_{10} . The four different simulation settings will again be given by (i) narrow normal, centered properly; (ii) wide normal, centered properly; (iii) narrow normal, off-center; and (iv) wide normal, off-center. For scenario (i), we posited a distribution for ξ_{10} given by $\xi_{10}^{(b)} \sim N(\xi_{10}, 0.1)$. For the wide distributions settings, we increased the variance to 0.5, and for the off-center simulations, we centered the distribution at the true value of ξ_{10} plus 0.1.

Figure 2 shows boxplots of the point estimates of ψ across the 1000 repetitions. For ψ_0 , adjusting for the unmeasured confounder using the narrow, correctly centered distribution reduced the rMSE from 1.104 to 0.188. This was worse than our proposed procedure that reduced the rMSE to 0.125. The G-estimation sensitivity analysis was also more sensitive to misspecification of the bias parameter distribution with the incorrectly centered distribution resulting in an rMSE of 0.383 and 0.488 for the narrow and wide distributions, respectively. For the interaction effect parameters, ψ_1 and ψ_2 , the bias-adjusted estimates had a greater rMSE than the unadjusted estimates with the narrow, correctly centered bias distribution increasing the rMSE from 0.204 to 0.316 for ψ_1 and from 0.125 to 0.184 for ψ_2 . Section D of the Supporting Information contains additional results on the coverage of confidence intervals and the proportion of new patients whose recommended treatment from the estimated regimes matched the optimal DTR for this simulation study. Section D also contains an additional simulation study in which we do not assume $\xi_{11} = 0$ and $\xi_{12} = 0$. The results for this were very poor with all four of the posited bias parameter distributions resulting in an increase in the rMSE in the estimates of ψ_0 , ψ_1 , and ψ_2 . These simulation results demonstrate how simpler sensitivity analysis methods that avoid positing a model for the relationship between the unmeasured confounder and the other covariates can reduce the number of models and bias parameters distributions that need to be posited, but can produce considerably less reliable results about the effect of the unmeasured confounding on the estimated DTR.

We conducted a similar simulation study with a two-stage study to assess the effectiveness of a Monte Carlo sensitivity analysis for the estimation of multistage DTRs. We let the data-generating model be given by:

$$\begin{aligned}
 U &\sim N(0, \sigma_u^2), \\
 X_{11} &= \phi_{10} + \phi_{11}U + \epsilon_{x_{11}}, & \epsilon_{x_{11}} &\sim N(0, \sigma_{x_{11}}^2), \\
 X_{12} &= \phi_{20} + \phi_{21}U + \epsilon_{x_{12}}, & \epsilon_{x_{12}} &\sim N(0, \sigma_{x_{12}}^2), \\
 X_2 &= \omega_0 + \omega_1 X_{11} + \omega_2 X_{12} + \epsilon_{x_2}, & \epsilon_{x_2} &\sim N(0, \sigma_{x_2}^2), \\
 P(A_1 = 1 \mid \mathbf{H}_1 = \mathbf{h}_1, U = u) &= [1 + \exp\{-(\alpha_{10} + \alpha_{11}x_{11} + \alpha_{12}x_{12} + \alpha_{13}u)\}]^{-1}, \\
 P(A_2 = 1 \mid \mathbf{H}_2 = \mathbf{h}_2, U = u) &= [1 + \exp\{-(\alpha_{20} + \alpha_{21}x_{11} + \alpha_{22}x_{12} + \alpha_{23}a_1 + \alpha_{24}x_2 + \alpha_{25}u)\}]^{-1}, \\
 Y &= \beta_{20} + \beta_{21}X_{11} + \beta_{22}X_{12} + \beta_{23}A_1 + \beta_{24}A_1X_{11} + \beta_{25}A_1X_{12} \\
 &\quad + \beta_{26}X_2 + \beta_u U + A_2(\psi_{20} + \psi_{21}X_{11} + \psi_{22}X_{12} + \psi_{23}X_2) + \epsilon_y, & \epsilon_y &\sim N(0, \sigma_y^2).
 \end{aligned}$$

As before, we conducted 1000 repetitions for the simulation study, set the sample size to $n = 1000$, and conducted the sensitivity analysis using $B = 500$ Monte Carlo repetitions. The parameter values for this data-generating process were given by:

$$\begin{aligned}
 \beta_2 &= (\beta_{20}, \beta_{21}, \beta_{22}, \beta_{23}, \beta_{24}, \beta_{25}, \beta_{26}, \beta_u) = (1, -1, 1, -1, 1, 1, 1, 2), \\
 \psi_2 &= (\psi_{20}, \psi_{21}, \psi_{22}, \psi_{23}) = (-1, 0.5, 0.5, 0.5), \\
 \alpha_1 &= (\alpha_{10}, \alpha_{11}, \alpha_{12}, \alpha_{13}) = (0, 1, 1, 2), & \phi_1 &= (\phi_{10}, \phi_{11}) = (0, 1), \\
 \alpha_2 &= (\alpha_{20}, \alpha_{21}, \alpha_{22}, \alpha_{23}, \alpha_{24}, \alpha_{25}) = (0, 1, 1, 1, 1, 3), & \phi_2 &= (\phi_{10}, \phi_{11}) = (0, -1), \\
 \sigma_u^2 &= \sigma_{x_{11}}^2 = \sigma_{x_{12}}^2 = \sigma_{x_2}^2 = \sigma_y^2 = 1, & \omega &= (\omega_0, \omega_1, \omega_2) = (0, 1, 1).
 \end{aligned}$$

We again varied the posited distributions for the parameters of the unmeasured confounder model and the effect of the unmeasured confounder using the same posited distributions as the single-stage simulation study.

Figures 3 and 4 show boxplots of the point estimates for ψ_2 and ψ_1 , respectively, across the 1000 repetitions. The results for the two-stage study were very similar to the one-stage results. In the second-stage models, the unmeasured confounder caused significant bias in ψ_{20} in the unadjusted analysis leading to an rMSE of 1.201. Applying our Monte Carlo sensitivity analysis reduced the bias significantly with all scenarios having an rMSE of 0.242 or less. We found that the unmeasured confounder did not cause bias in the estimates of ψ_{21} , ψ_{22} , and ψ_{23} , leading to similar results between the unadjusted and adjusted estimates. There was significant bias in the unadjusted estimate of ψ_{10} for the first-stage blip model. The rMSE of the unadjusted estimate is 0.745, whereas the rMSE of the bias-adjusted estimate under the narrow, centered parameter distribution is 0.165. The rMSE increased to 0.25 when using a wide parameter distribution that is not centered on the true parameter value.

As before, we simulated 10,000 additional patients and assessed whether the treatment recommended by each of the estimated regimes matched the recommendation of the true optimal regime for each of the 1000 repetitions. Table 3 displays the proportion of patients who were recommended the same treatment as the true optimal regime when following each of the different estimated regimes. The bias from the unmeasured confounder caused the unadjusted estimated regime to recommend the same treatment as the true optimal regime 81.4% of the time at the first stage and 70.2% of the at the second stage. After performing Monte Carlo sensitivity analysis using the narrow, centered posited distribution, the proportion increased to 95.1% at the first stage and 93.7% at the second stage. The alternative parameter distributions produced similar results.

We also examined the coverage and width of the confidence intervals generated using a Monte Carlo sensitivity analysis for the two-stage study. Table 4 contains the empirical coverage and width of 95% confidence intervals for the parameters of the second-stage blip model. The empirical coverage of ψ_{20} for the unadjusted analysis was 0% as it was significantly biased by the unmeasured confounder. Monte Carlo sensitivity analysis produced intervals for ψ_{20} with 100% coverage for all of the simulation scenarios at the expense of being much wider. The remaining second-stage blip parameters, ψ_{21} , ψ_{22} , and ψ_{23} , were not biased in the unadjusted analysis, but the unmeasured confounder led to confidence intervals that did not achieve the nominal coverage probability. Monte Carlo sensitivity analysis produced confidence intervals that achieved the nominal coverage for all sets of posited parameter distributions.

Table 5 contains the empirical coverage and average width of 95% confidence intervals for the parameters of the first-stage blip model. The unadjusted confidence interval for ψ_{10} had a coverage of only 0.1%, whereas the sensitivity analysis resulted in conservative confidence intervals that had coverage of 100%. For ψ_{11} and ψ_{12} , the unadjusted analysis resulted in confidence intervals that did not attain nominal coverage with empirical coverages of 86.1% and 86.8%, respectively, despite the unmeasured confounder not resulting in bias

in the parameters estimation. The sensitivity analysis intervals were instead conservative with coverages of 100% and 98.9% for ψ_{11} and ψ_{12} , respectively, for the narrow, centered parameter distributions. The confidence intervals resulting from the sensitivity analysis for the first-stage blip model were considerably wider than those for the second-stage blip model parameters.

5 | ILLUSTRATION USING KPWA EHR DATA

We applied the proposed method to data from EHRs. KPWA is a health system that provides health care and insurance to its members. This study used data from EHRs and health insurance claims for all KPWA clients to study the use of antidepressants for treating depression. The data include information on demographics, prescription fills, and depressive symptom severity for 82,691 patients with depression who received antidepressant medication from 2008 through 2018. Severity of depression symptoms was assessed through the use of the Patient Health Questionnaire-8 (PHQ) (Kroenke et al., 2001). This is a self-report questionnaire that produces a score ranging from 0 to 24, such that higher values indicate more severe symptoms of depression. The inclusion criteria for the study were that patients must have been 13 years or older, have been enrolled in KPWA insurance for at least a year, have been diagnosed with a depressive disorder in the year before or 15 days after treatment initiation, had no prescription fills for antidepressant medications in the past year, and had no diagnosis for personality, bipolar, or psychotic disorders in the past year. We additionally required, for this analysis, that patients did not have missing information on obesity, baseline PHQ, or follow-up PHQ at 1 year after baseline. Our analysis focuses on demonstrating the impact of our sensitivity analysis but does not adjust for missing information; therefore, the estimated DTRs must be interpreted as potentially biased; any clinical findings should be viewed in this light.

Research has suggested that obesity increases the risk of depression (Luppino et al., 2010). Weight gain has been found to be a side effect of taking antidepressants (Fava, 2000), with different classes of antidepressants affecting weight differently; treatment guidelines for clinicians recommend using the current weight of a patient as a consideration when prescribing antidepressants (Santasieri and Schwartz, 2015). Obesity is therefore a potentially important confounder that, if unmeasured, could lead to bias in observational studies of their effects, whether average or individualized. We considered a patient to be obese if they had a body mass index (BMI) of 30 or larger. We used this data in two different ways. We assumed that obesity was unavailable and estimated a DTR to minimize depression symptoms without adjusting for obesity in the analysis. We then applied the proposed procedure to assess how sensitive the estimated regime was to obesity being unmeasured and compared this to the estimated regime with obesity measured. We also conducted a plasmode simulation study in which we used the real data with simulated obesity and patient outcomes. This allowed us to know the true optimal DTR and distribution of obesity.

5.1 | Application to the study of first-line choice of antidepressant

Our outcome is the negative of the PHQ score after 1 year so that higher values correspond with better patient outcomes. Equally, we could use the PHQ score and find the regime that minimizes instead of maximizes the mean outcome. The PHQ score after 1 year was given by the PHQ score recorded between 305 and 425 days after initiating treatment that was measured closest to 365 days after medication initiation. Patients were treated with 1 of 17 different antidepressants initially. Selective serotonin reuptake inhibitors (SSRIs) are a class of antidepressants that increase serotonin levels in the brain by decreasing reabsorption by nerve cells and are commonly prescribed due to having generally milder side effects. We classified the initial treatment received as an antidepressant from either the SSRI class or an alternative class of antidepressants.

We let Y denote the negative of the PHQ score after 1 year. We denote treatment by A , such that $A = 1$ if assigned an SSRI and $A = 0$ if assigned a non-SSRI. We considered sex, age, baseline PHQ, and obesity as confounders and examined the bias that occurs if obesity was unmeasured. We centered the age variable to make the intercept more interpretable. Our outcome model was given by:

$$E(Y | X, A; \beta, \psi) = \beta_0 + \beta_1 \text{SEX} + \beta_2 \text{AGE} + \beta_3 \text{PHQ} + \beta_u \text{OBESE} + A(\psi_0 + \psi_1 \text{SEX} + \psi_2 \text{AGE} + \psi_3 \text{PHQ}) \quad (3)$$

such that the optimal DTR is then given by recommend an SSRI if $(\psi_0 + \psi_1 \text{SEX} + \psi_2 \text{AGE} + \psi_3 \text{PHQ})$ is greater than zero.

To posit the models for the unmeasured confounder and the distributions for the bias parameters in practice, we would ideally conduct a small, secondary study. To replicate this process for this example, we took a random subsample of 250 patients from our data and treated this as a secondary data set. We found that obesity was correlated with age, baseline PHQ, race, census block education level (EDU), and diagnosis of an anxiety disorder in the prior year (ANX). Race was categorized as Asian, Black or African American, Hispanic, Native Hawaiian/Pacific Islander, American Indian/Native Alaskan, White, other, or unknown and was coded using dummy coding with Asian as the reference group. EDU was given by an indicator for whether less than 25% of people living in the patient's census block had a college degree. Therefore, to conduct a Monte Carlo sensitivity analysis for the bias if obesity is unmeasured, we posited a model for obesity given by:

$$P(\text{OBESE} = 1 | X; \zeta) = [1 + \exp\{-\zeta_0 + \zeta_1 \text{AGE} + \zeta_2 \text{PHQ} + \zeta_3^T \text{RACE} + \zeta_4 \text{EDU} + \zeta_5 \text{ANX} + \zeta_6 A\}]^{-1}. \quad (4)$$

We first estimated β_u and ζ using the secondary data set. We then posited normal distributions for β_u and ζ that were centered at the estimated values of the parameters with standard deviations equal to the estimated standard errors.

We also examined using our proposed procedure when we do not have a secondary data set to help specify distributions for the bias parameters and instead need to use subject matter expertise. Recall that research has found that obesity increases the risk of depression that implies that obesity is negatively correlated with the negative of the PHQ score. Therefore, we will let the distribution for β_u be given by a uniform distribution from -2 to 0 . Since we centered the age variable, $\text{expit}(\zeta_0)$ can be interpreted as the probability of being obese, given that an individual is 45 years old, Asian, has a baseline PHQ score of 0, has no diagnosis of an anxiety disorder, is not from a low education census block, and who received a non-SSRI. We used the overall prevalence of obesity in Asian people over 18 years old in the United States obtained from the 2017 National Health Interview Survey to center the distribution for ζ_0 (Blackwell and Villarroel, 2018). Since the overall prevalence was given by 0.119, the distribution for ζ_0 was selected to be a uniform distribution with support $[\text{logit}(0.119) - 0.5, \text{logit}(0.119) + 0.5]$. We chose to use a uniform distribution as using the prevalence in the entire Asian population in the United States as a surrogate for the conditional probability given by $\text{expit}(\zeta_0)$ is an imperfect approach. We chose the distributions for each component of ζ_3 in the same manner with a uniform distribution centered using the overall prevalence in the United States by race with a width of 1. The remaining parameters in ζ were given by a uniform distribution from -1 to 1 because we are unsure of how they are related to obesity.

Table 6 contains estimates and confidence intervals for ψ from the full model, the model with obesity unmeasured, and after adjusting for the bias with and without a secondary data set to posit the bias parameter distributions. The adjusted estimates of ψ using the secondary data set with obesity unmeasured were close to the estimates from the full model with obesity included; however, uncertainty was increased reflecting the lack of information available on the unmeasured confounder. The adjusted estimates using uniform distributions posited with subject matter expertise resulted in estimates close to the unadjusted estimate as expected. Both sensitivity analyses also produced confidence intervals that were wider for the parameters that were biased due to the unmeasured confounding. Section E of the Supporting Information contains a similar sensitivity analysis in which we used the full data set with obesity included to posit distributions for the bias parameters.

5.2 | Plasmode simulations

We also conducted a plasmode simulation study in which we used the real data with simulated obesity and patient outcomes. This allowed us to know the true optimal DTR and unmeasured confounder distribution. The data were simulated using models 3 and 4 that we used for the real data analysis with the outcome having a random error given by $\epsilon \sim N(0, 1)$.

We let the value of β , ψ , and ζ be given by the estimated values from the data that can be found in Section F of the Supporting Information. For conducting the sensitivity analysis, we posited the same normal distributions for β_u and ζ that were posited using the small, secondary study for the empirical evaluation in addition to the uniform distributions posited using subject matter expertise. We simulated 1000 data sets and estimated the adjusted parameter estimates for each data set. Table 7 contains the estimates of ψ averaged over the 1000 simulated data sets in addition to the empirical coverage and average width of

95% confidence intervals for ψ . The unadjusted estimate for ψ_0 was biased with an average estimate of -1.432 as opposed to the true value of -1.556 . Adjusting for obesity using a secondary study reduced the bias and resulted in an estimate of -1.538 . Positing the bias parameter distributions with subject matter expertise resulted in an estimate of -1.446 that was closer to the unadjusted estimate than the true value. The 95% confidence intervals from the unadjusted analysis did not achieve the nominal rate for any of the parameters. Monte Carlo sensitivity analysis resulted in slightly wider intervals that had the correct coverage when we posited the bias parameter distributions using a secondary data set or subject matter expertise.

6 | DISCUSSION

We proposed a method for conducting sensitivity analysis for bias due to unmeasured confounding in the estimation of DTRs. We used a Monte Carlo sensitivity analysis to estimate a bias-adjusted DTR and construct confidence intervals for the parameters indexing the optimal regime that account for the uncertainty due to unmeasured confounding. This procedure is straightforward to implement and can be used for continuous or binary unmeasured confounders. This approach was found to perform well for both simulated and real data.

This method requires positing parametric models for the relationship between the unmeasured confounder(s) and the outcome, treatment, and measured confounders. Moreover, the model for the conditional mean of the unmeasured confounder needs to be of high quality, which can be challenging without external data or domain expertise. We also must posit probability distributions for the parameters indexing these models. These distributions do not need to be centered on the true value of the parameter to construct confidence intervals, and they only need the true value to not be in the tails of the distribution. If there are high levels of uncertainty about these parameters, wider distributions can be posited resulting in wider confidence intervals for the parameters indexing the estimated optimal DTR. We posited normal distributions for these parameters, but any probability distribution can be used that accurately reflect the prior information about the true value of these parameters. Here, we focused on DTRs that are estimated using dWOLS, though this approach can be applied to any regression-based method of estimation. Direct-search or value-search estimators are an alternative approach to estimating DTRs that our proposed method could not be directly applied to (Laber and Zhao, 2015; Orellana et al., 2010). We leave extensions to direct-search estimators as future work.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

ACKNOWLEDGMENTS

This work was supported by the National Institute of Mental Health of the National Institutes of Health under Award Number R01 MH114873. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health. EEMM is a Canada Research Chair (Tier 1) and acknowledges the support of a chercheur de mérite career award from the Fonds de Recherche du Québec, Santé.

Funding information

Fonds de Recherche du Québec, Santé; National Institute of Mental Health, Grant/Award Number: R01 MH114873

DATA AVAILABILITY STATEMENT

Research data are not shared.

REFERENCES

- Bickel PJ, Gotze F, & van Zwet WR (1997). Resampling fewer than n observations: Gains, losses, and remedies for losses. *Statistica Sinica*, 7, 1–31.
- Blackwell DL, & Villarroel MA (2018). Tables of summary health statistics for U.S. adults: 2017 National Health Interview Survey. <http://www.cdc.gov/nchs/nhis/SHS/tables.htm>
- Cain LE, Robins JM, Lanoy E, Logan R, Costagliola D, & Hernán MA (2010). When to start treatment? A systematic approach to the comparison of dynamic regimes using observational data. *The International Journal of Biostatistics*, 6(2), 1557–4679.
- Chakraborty B, Laber EB, & Zhao Y. (2013). Inference for optimal dynamic treatment regimes using an adaptive m -out-of- n bootstrap scheme. *Biometrics*, 69(3), 714–723. [PubMed: 23845276]
- Chakraborty B, & Moodie EEM (2013). *Statistical methods for dynamic treatment regimes*. Springer.
- Chen S, & Zhang B. (2021). Estimating and improving dynamic treatment regimens with a time-varying instrumental variable. arXiv: 2104.07822.
- Cornfield J, Haenszel W, Hammond EC, Lilienfeld AM, Shimkin MB, & Wynder EL (1959). Smoking and lung cancer: Recent evidence and a discussion of some questions. *Journal of the National Cancer Institute*, 22(1), 173–203. [PubMed: 13621204]
- Cui Y, & Tchetgen Tchetgen E. (2021). A semiparametric instrumental variable approach to optimal treatment regimes under endogeneity. *Journal of the American Statistical Association*, 116(533), 162–173. [PubMed: 33994604]
- Dumbgen L. (1993). On nondifferentiable functions and the bootstrap. *Probability Theory and Related Fields*, 95, 125–140.
- Ertefaie A, & Strawderman RL (2018). Constructing dynamic treatment regimes in infinite-horizon settings. *Biometrika*, 105(4), 963–977.
- Fava M. (2000). Weight gain and antidepressants. *Journal of Clinical Psychiatry*, 61(11), 37–41. [PubMed: 10926053]
- Greenland S. (2005). Multiple-bias modelling for analysis of observational data. *Journal of the Royal Statistical Society: Series A*, 168(2), 267–306.
- Groenwold RHH, Nelson DB, Nichol KL, Hoes AW, & Hak E. (2010). Sensitivity analyses to estimate the potential impact of unmeasured confounding in causal research. *International Journal of Epidemiology*, 39, 107–117. [PubMed: 19948779]
- Hernan MA, & Robins JM (2006). Instruments for causal inference: An epidemiologist's dream. *Epidemiology*, 17(4), 360–372. [PubMed: 16755261]
- Hernan MA, & Robins JM (2020). *Causal inference: What if*. Chapman & Hall/CRC.
- Kallus N, & Zhou A. (2019). Confounding-robust policy improvement. arXiv: 1805.08593.
- Kroenke K, Spitzer RL, & Williams JBW (2001). The PHQ-9: Validity of a brief depression severity measure. *Journal of General Internal Medicine*, 16(9), 606–613. [PubMed: 11556941]
- Laber EB, Lizotte DJ, Qian M, Pelham WE, & Murphy SA (2014). Dynamic treatment regimes: Technical challenges and applications. *Electronic Journal of Statistics*, 8(1), 1225–1272. [PubMed: 25356091]
- Laber EB, & Zhao YQ (2015). Tree-based methods for estimating individualized treatment regimes. *Biometrika*, 102(3), 501–514. [PubMed: 26893526]
- Lash TL, Fox MP, & Fink AK (2009). *Applying quantitative bias analysis to epidemiologic data*. Springer.

- Lavori P, & Dawson R. (2000). A design for testing clinical strategies: Biased adaptive within-subject randomization. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 163(1), 29–38.
- Lavori PW, & Dawson R. (2004). Dynamic treatment regimes: Practical design considerations. *Clinical Trials*, 1(1), 9–20. [PubMed: 16281458]
- Liao L, Fu Z, Yang Z, Kolar M, & Wang Z. (2021). Instrumental variable value iteration for causal offline reinforcement learning. arXiv: 2104.09907.
- Lin DY, Psaty BM, & Kronmal RA (1998). Assessing the sensitivity of regression results to unmeasured confounders in observational studies. *Biometrics*, 54(3), 948–963. [PubMed: 9750244]
- Luckett DJ, Laber EB, Kahkoska AR, Maahs DM, Mayer-Davis E, & Kosorok MR (2018). Estimating dynamic treatment regimes in mobile health using V-learning. *Journal of the American Statistical Association*, 115(530), 692–706.
- Luppino FS, de Wit LM, Bouvy PF, Stijnen T, Cuijpers P, Penninx BWJH, & Zitman FG (2010). Overweight, obesity, and depression: A systematic review and meta-analysis of longitudinal studies. *Archives of General Psychiatry*, 67(3), 220–229. [PubMed: 20194822]
- McCandless LC, Gustafson P, & Levy A. (2007). Bayesian sensitivity analysis for unmeasured confounding in observational studies. *Statistics in Medicine*, 26(11), 2331–2347. [PubMed: 16998821]
- McCandless LC, & Gustafson P. (2017). A comparison of Bayesian and Monte Carlo sensitivity analysis for unmeasured confounding. *Statistics in Medicine*, 36(18), 2887–2901. [PubMed: 28386994]
- Moodie EEM, & Richardson TS (2009). Estimating optimal dynamic regimes: Correcting bias under the null. *The Scandinavian Journal of Statistics*, 37, 126–146.
- Murphy SA. (2003). Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 65(2), 331–355.
- Murphy SA (2005). An experimental design for the development of adaptive treatment strategies. *Statistics in Medicine*, 24(10), 1455–1481. [PubMed: 15586395]
- Murray T, Yuan Y, & Thall P. (2018). A Bayesian machine learning approach for optimizing dynamic treatment regimes. *Journal of the American Statistical Association*, 113(523), 1255–1267. [PubMed: 30739965]
- Orellana L, Rotnitzky A, & Robins J. (2010). Dynamic regime marginal structural mean models for estimation of optimal dynamic treatment regimes, part I: Main content. *International Journal of Biostatistics*, 6(2), 1–49.
- Phillips CV (2003). Quantifying and reporting uncertainty from systemic errors. *Epidemiology*, 14(4), 459–466. [PubMed: 12843772]
- Qiu H, Carone M, Sadikova E, Petukhova M, Kessler RC, & Luedtke A. (2021). Optimal individualized decision rule using instrumental variable methods. *Journal of the American Statistical Association*, 116(533), 174–191. [PubMed: 33731969]
- Robins J. (2004). Optimal structural nested models for optimal sequential decisions. In Lin DY & Heagerty P. (Eds.), *Proceedings of the Second Seattle Symposium in Biostatistics: Analysis of Correlated Data* (pp. 189–326). Springer.
- Robins JM, Rotnitzky A, & Scharfstein DO (1999). Sensitivity analysis for selection bias and unmeasured confounding in missing data and causal inference models. In Halloran ME & Berry D. (Eds.), *Statistical models in epidemiology: The environment and clinical trials* (pp. 1–92). Springer-Verlag.
- Rose EJ, Moodie EEM, & Shortreed S. (2022). Using pilot data to size observational studies for the estimation of dynamic treatment regimes. arXiv: 2202.09451.
- Rosenbaum PR (1987). Sensitivity analysis for certain permutation inferences in matched observational studies. *Biometrika*, 74(1), 13–26.
- Rosenbaum PR, & Rubin DB (1983). Assessing sensitivity to an unobserved binary covariate in an observational study with binary outcome. *Journal of the Royal Statistical Society: Series B*, 45(2), 212–218.

- Rubin D. (1978). Bayesian inference for causal effects: The role of randomization. *The Annals of Statistics*, 6(1), 34–58.
- Santasieri D, & Schwartz TL (2015). Antidepressant efficacy and side-effect burden: A quick guide for clinicians. *Drugs in Context*, 4.
- Shao J. (1994). Bootstrap sample size in nonregular cases. *Proceedings of the American Mathematical Society*, 122(4), 1251–1262.
- Steenland K, & Greenland S. (2004). Monte Carlo sensitivity analysis and Bayesian analysis of smoking as an unmeasured confounder in a study of silica and lung cancer. *American Journal of Epidemiology*, 160(4), 384–392. [PubMed: 15286024]
- Swanepoel JWH (1986). A note on proving that the (modified) bootstrap works. *Communications in Statistics-Theory and Methods*, 15, 3193–3203.
- Thall PF, Millikan RE, & Sung H-G (2000). Evaluating multiple treatment courses in clinical trials. *Statistics in Medicine*, 19(8), 1011–1028. [PubMed: 10790677]
- Tsiatis AA, Davidian M, Holloway ST, & Laber EB (2019). *Dynamic treatment regimes: Statistical methods for precision medicine*. CRC Press.
- van der Laan MJ, Petersen ML, & Joffe MM (2005). History-adjusted marginal structural models and statically-optimal dynamic treatment regimens. *The International Journal of Biostatistics*, 1(1).
- Wallace MP, & Moodie EEM (2015). Doubly-robust dynamic treatment regimen estimation via weighted least squares. *Biometrics*, 71(3), 636–644. [PubMed: 25854539]
- Wu F, Laber EB, Lipkovich IA, & Severus E. (2015). Who will benefit from antidepressants in the acute treatment of bipolar depression? A reanalysis of the STEP-BD study by Sachs et al. 2007, using Q-learning. *International Journal of Bipolar Disorders*, 3(1), 7. [PubMed: 25844303]
- Zhang B, Small DS, & Zhao Q. (2021). Selecting and ranking individualized treatment rules with unmeasured confounding. *Journal of the American Statistical Association*, 116(533), 295–308.
- Zhang Y, Laber E, Tsiatis A, & Davidian M. (2018). Interpretable dynamic treatment regimes. *Journal of the American Statistical Association*, 113(524), 1541–1549. [PubMed: 30774169]




FIGURE 1. Boxplots of the point estimates for ψ under an unadjusted model and when using Monte Carlo sensitivity analysis to adjust for bias due to unmeasured confounding for the one-stage data-generating model.

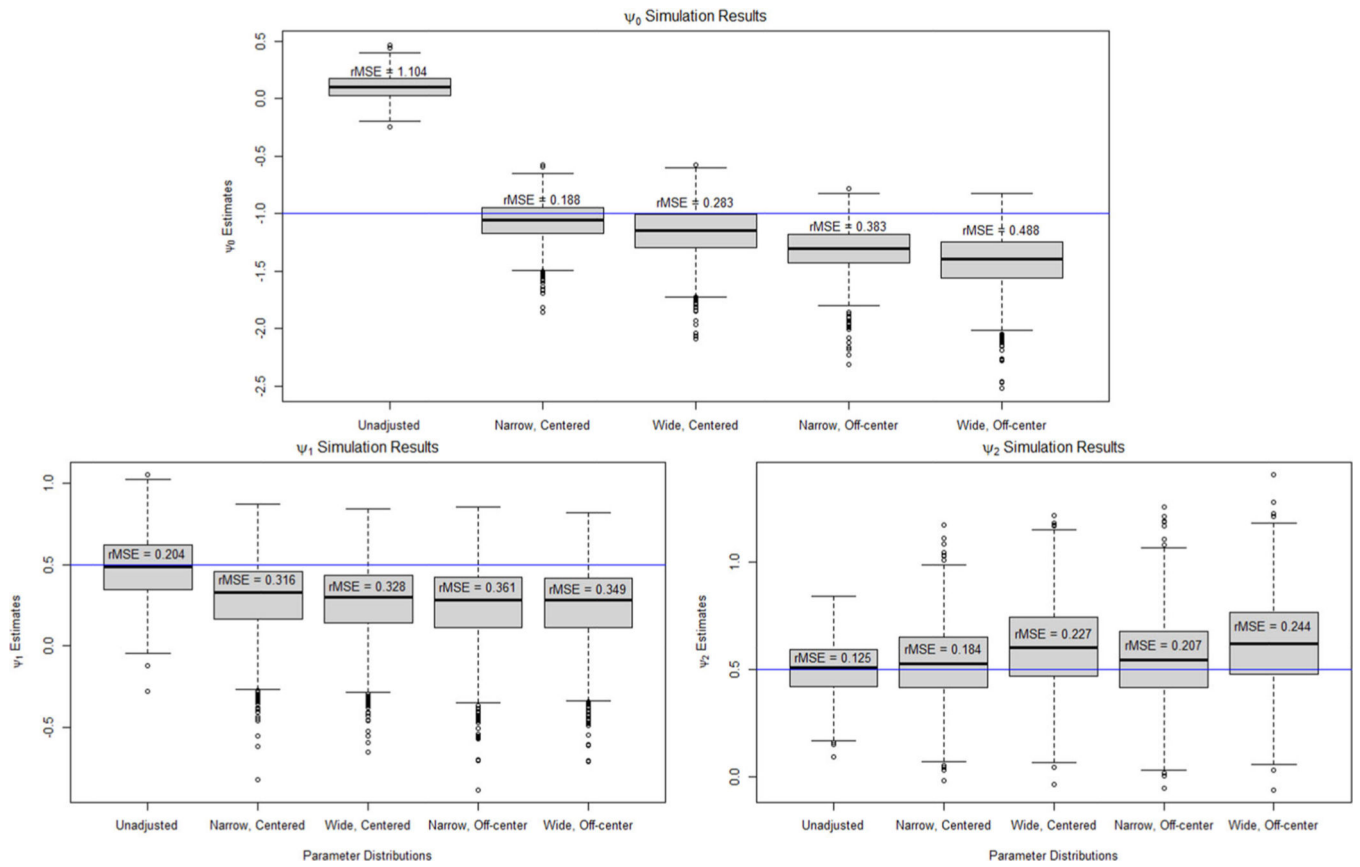


FIGURE 2. Boxplots of the point estimates for ψ under an unadjusted model and when using G-estimation sensitivity analysis to adjust for bias due to unmeasured confounding for the one-stage data-generating model when we assume we know ξ_{11} and ξ_{12} .

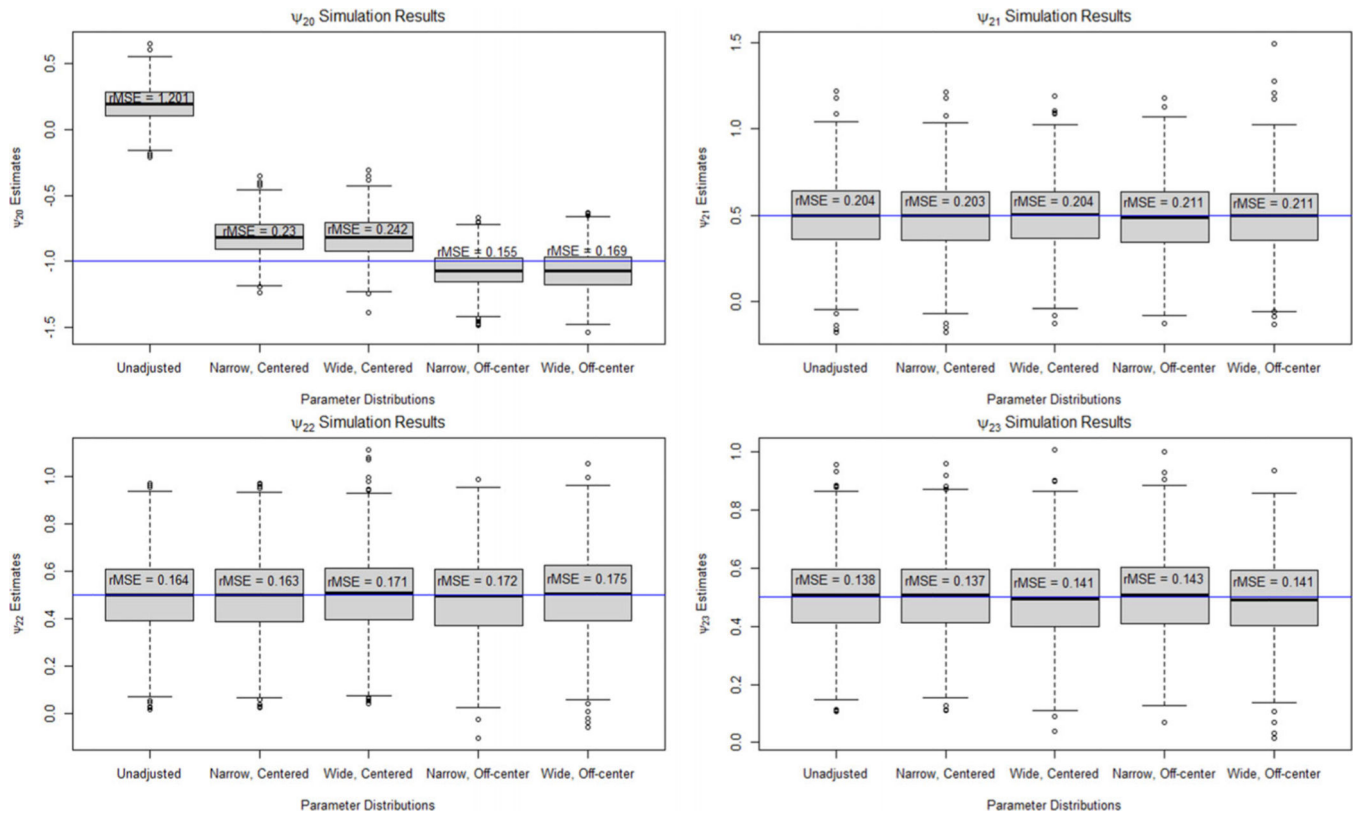


FIGURE 3. Boxplots of the point estimates for ψ_2 under an unadjusted model and when using Monte Carlo sensitivity analysis to adjust for bias due to unmeasured confounding for the two-stage data-generating model.




FIGURE 4. Boxplots of the point estimates for ψ_1 under an unadjusted model and when using Monte Carlo sensitivity analysis to adjust for bias due to unmeasured confounding for the two-stage data-generating model.

TABLE 1

Proportion of patients whose recommended treatment when following each of the estimated regime matches the recommendation of the true optimal regime for the one-stage data-generating model.

Parameter Distr.	Proportion optimal
Unadjusted	0.528
Narrow, Centered	0.956
Wide, Centered	0.953
Narrow, Off-center	0.959
Wide, Off-center	0.958

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

TABLE 2

Coverage (Cvr.) and average width (Wth.) of the 95% confidence intervals for ψ for the unadjusted analysis and sensitivity analysis under each of the posited parameter distributions for the one-stage data-generating model.

Parameter Distr.	Cvr. (Wth.) ψ_0	Cvr. (Wth.) ψ_1	Cvr. (Wth.) ψ_2
Unadjusted	0.000* (0.359)	0.897* (0.426)	0.887* (0.300)
Narrow, Centered	1.000* (2.628)	0.946 (0.520)	0.947 (0.365)
Wide, Centered	1.000* (6.212)	0.954 (0.520)	0.959 (0.365)
Narrow, Off-center	1.000* (2.764)	0.950 (0.519)	0.951 (0.366)
Wide, Off-center	1.000* (6.495)	0.955 (0.519)	0.936* (0.367)

* indicates coverages that are significantly different than 95%.

TABLE 3

Proportion of patients whose recommended treatment when following each of the estimated regime matches the recommendation of the true optimal regime at each stage for the two-stage data-generating model.

Parameter Distr.	Stage 1	Stage 2
Unadjusted	0.814	0.702
Narrow, Centered	0.951	0.937
Wide, Centered	0.949	0.936
Narrow, Off-center	0.951	0.948
Wide, Off-center	0.948	0.947

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

TABLE 4

Coverage (Cvr.) and average width (Wth.) of the 95% confidence intervals for ψ_2 for the unadjusted analysis and sensitivity analysis under each of the posited parameter distributions for the two-stage data-generating model.

Parameter Distr.	Cvr. (Wth.) ψ_{20}	Cvr. (Wth.) ψ_{21}	Cvr. (Wth.) ψ_{22}	Cvr. (Wth.) ψ_{23}
Unadjusted	0.000* (0.339)	0.810* (0.529)	0.810* (0.435)	0.801* (0.359)
Narrow, Centered	1.000* (2.634)	0.948 (0.803)	0.953 (0.658)	0.943 (0.541)
Wide, Centered	1.000* (6.208)	0.948 (0.804)	0.951 (0.654)	0.941 (0.537)
Narrow, Off-center	1.000* (2.775)	0.932* (0.803)	0.949 (0.657)	0.936* (0.542)
Wide, Off-center	1.000* (6.483)	0.934* (0.805)	0.937 (0.659)	0.941 (0.540)

* indicates coverages that are significantly different than 95%.

TABLE 5

Coverage (Cvr.) and average width (Wth.) of the 95% confidence intervals for ψ_1 for the unadjusted analysis and sensitivity analysis under each of the posited parameter distributions for the two-stage data-generating model.

Parameter Distr.	Cvr. (Wth.) ψ_{10}	Cvr. (Wth.) ψ_{11}	Cvr. (Wth.) ψ_{12}
Unadjusted	0.001* (0.456)	0.861* (0.542)	0.868* (0.382)
Narrow, Centered	1.000* (4.236)	1.000* (2.609)	0.989* (2.439)
Wide, Centered	1.000* (7.362)	1.000* (2.828)	1.000* (2.590)
Narrow, Off-center	1.000* (4.607)	1.000* (2.833)	0.995* (2.663)
Wide, Off-center	1.000* (7.857)	1.000* (3.081)	1.000* (2.840)

* indicates coverages that are significantly different than 95%.

TABLE 6

Estimates and 95% confidence intervals for treatment decision rule parameters ψ from models with and without obesity and adjusted estimates from the sensitivity analysis for unmeasured confounding of obesity.

Covariate	Full model	Obesity unmeasured	Adjusted Est. secondary data	Adjusted est. subject expertise
A	-1.56 (-3.11, 0.00)	-1.46 (-3.02, 0.10)	-1.60 (-3.37, 0.12)	-1.48 (-3.13, 0.42)
A × SEX	0.72 (-0.34, 1.79)	0.69 (-0.38, 1.76)	0.73 (-0.58, 2.05)	0.69 (-0.70, 1.98)
A × AGE	0.00 (-0.03, 0.03)	0.00 (-0.03, 0.03)	0.00 (-0.04, 0.04)	0.00 (-0.04, 0.04)
A × PHQ	0.12 (0.02, 0.22)	0.12 (0.02, 0.22)	0.12 (0.00, 0.25)	0.12 (-0.01, 0.24)

TABLE 7

Estimates (Est.), coverage (Cvr.), and average width (Wth.) of 95% confidence intervals for treatment decision rule parameters ψ after adjusting for the unmeasured confounder obesity averaged over 1000 plasmode data sets using data from KPWA.

	Est. ψ_0	Est. ψ_1	Est. ψ_2	Est. ψ_3
True Value	-1.556	0.724	0.001	0.119
Unadjusted	-1.432	0.714	0.001	0.118
Adjusted—Secondary Data	-1.538	0.721	0.001	0.119
Adjusted—Subject Expertise	-1.446	0.715	0.001	0.119
	Cvr. (Wth.) ψ_0	Cvr. (Wth.) ψ_1	Cvr. (Wth.) ψ_2	Cvr. (Wth.) ψ_3
Unadjusted	0.822 [*] (0.665)	0.880 [*] (0.455)	0.870 [*] (0.012)	0.880 [*] (0.042)
Adjusted—Secondary Data	0.975 [*] (0.958)	0.949 (0.586)	0.958 (0.016)	0.959 (0.054)
Adjusted—Subject Expertise	0.940 (0.889)	0.953 (0.588)	0.975 [*] (0.018)	0.981 [*] (0.057)

^{*} indicates coverages that are significantly different than 95%.