
Bayesian Inference and Partial Identification in Multi-Treatment Causal Inference with Unobserved Confounding

Jiajing Zheng
UC Santa Barbara

Alexander D’Amour
Google Research

Alexander Franks
UC Santa Barbara

Abstract

In causal estimation problems, the parameter of interest is often only partially identified, implying that the parameter cannot be recovered exactly, even with infinite data. Here, we study Bayesian inference for partially identified treatment effects in multi-treatment causal inference problems with unobserved confounding. In principle, inferring the partially identified treatment effects is natural under the Bayesian paradigm, but the results can be highly sensitive to parameterization and prior specification, often in surprising ways. It is thus essential to understand which aspects of the conclusions about treatment effects are driven entirely by the prior specification. We use a so-called transparent parameterization to contextualize the effects of more interpretable scientifically motivated prior specifications on the multiple effects. We demonstrate our analysis in an example quantifying the effects of gene expression levels on mouse obesity.

1 INTRODUCTION

In many causal inference settings, the causal quantity of interest cannot be precisely recovered from the distribution of the observed data. In these cases, we say that the causal parameter is not identified. However, even if the parameter is not identified, it is often possible to recover a non-trivial set of valid causal parameters that are equally compatible with the observed data. In this case, we say that parameter is partially identified.

In this paper, we explore the problem of Bayesian

inference for partially identified parameters in multi-treatment causal inference with unobserved confounding. The goal is to infer the effects of multiple simultaneously applied treatments on a scalar outcome, despite the fact that some confounders may be unmeasured. Approaches for Bayesian sensitivity analysis for unobserved confounding in more standard settings have been proposed (e.g. McCandless et al., 2006), but unobserved confounding in the multiple treatment setting is unique in that dependence among the treatments can be leveraged to learn something about unobserved confounders (Wang and Blei, 2018; Wang and Blei, 2019). However, this extra information is still insufficient for identifying the treatment effects without additional strong assumptions (see e.g. D’Amour, 2019; Ogburn et al., 2020). Zheng et al. (2021) establish conditions under which causal effects are at least *partially* identified, and we use this result as a starting point.

In principle, Bayesian posterior distributions over partially identified treatment effects should naturally reflect uncertainty inherent in the model. However, in practice, specifying a credible analysis can be difficult because posterior distributions for partially identified causal parameters are highly sensitive to the choice of prior distribution and model parameterization. On one hand, priors specified in terms of the most scientifically intuitive parameterization mix assumptions with observable and unobservable implications. Thus, it can be difficult to understand how much of the conclusions are driven by the prior rather than the data. On the other hand, so-called transparent parameterizations, which cleanly separates identified and unidentified parameters are often scientifically counterintuitive, making it difficult to bring genuine prior knowledge into a problem (Gustafson et al., 2014; Richardson et al., 2011). This tension can result in a disconnect between many in the causal inference community, who prefer assumptions stated in the transparent form, and scientific practitioners, who typically prefer the directly interpretable parameterization.

Here, we propose to bridge this gap by using a trans-

Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS) 2022, Valencia, Spain. PMLR: Volume 151. Copyright 2022 by the author(s).

parent parameterization recently proposed by Zheng et al. (2021) to study the implications of scientifically-motivated prior specifications for multiple treatments. In addition to flat “convenience” priors, we also consider priors that represent two types of scientific prior assumptions that are often considered to be plausible: a negative control assumption, that some specific interventions are known to have no effect on the outcome, and a sparsity assumption, that a large fraction of interventions have no effect. Assumptions of these types were recently proposed by Miao et al. (2020) to derive identification conditions. Our analysis here yields a deeper understanding of these conditions, particularly what happens when they are relaxed.

Related Work In this paper, we expand on the work by Zheng et al. (2021) who study partial identification in the multi-treatment causal inference problem. While they establish marginal partial identification regions for arbitrary treatment contrasts, they ignore issues of estimation, and do not fully explore how additional assumptions impact the joint partial identification region. Miao et al. (2020) provide a method for multiple treatment inference under a specific sparsity assumption which leads to causal identification. In this work, we provide a more general recipe for understanding how prior restrictions affect inferences.

Contributions Our contributions include 1) a characterization of transparent and scientific parameterizations for multi-treatment causal inference and the often surprising implications of prior choices 2) novel theoretical results about the impact of prior structural assumptions (e.g. null controls or sparsity) on the joint partial identification region and 3) a demonstration of our prior-sensitivity analysis in a mouse gene expression study, showcasing our ability to identify which aspects of the causal conclusions are driven by prior assumptions and which are driven by data.

2 PROBLEM SETTING

2.1 Multi-Treatment Causal Inference and Partial Identification

Let T be a random k -vector of concurrently assigned treatment variables, Y be a scalar random outcome of interest, and t and y be realizations of the respective random variables. Let U be a random m -vector of unobserved confounders. Our estimand of interest is the Population Average Treatment Effect (PATE) for interventions that change the value of T from t_1 to t_2 (i.e., treatment contrasts). We denote this PATE for a treatment contrast in *do*-notation (Pearl, 2009) as:

$$\tau_{t_1, t_2} := E(Y | do(t_1)) - E(Y | do(t_2)). \quad (1)$$

Here, $E(Y | do(t))$ is the expected outcome if we were to intervene and set the level of the treatment vector to $T = t$ for all units. This is distinct from the observed mean for t -treated units, $E[Y | T = t]$, when there are confounders. We denote all probability density/mass functions with f .

Throughout, we make the following assumptions: 1) **latent ignorability**: U block all backdoor paths between T and Y (Pearl, 2009), 2) **positivity**: $f(T = t | U = u) > 0$ for all t and u such that $f(U = u) > 0$. and 3) **SUTVA**: there are no hidden versions of the treatments and there is no interference between units (see Rubin, 1980). Ignorability, a commonly made assumption in observational inference, *does not* hold, since we explicitly assume that U is not directly unobserved. Let $\check{\tau}_{t_1, t_2}$ denote the causal effect for treatment contrast t_1 vs t_2 under the (potentially naive) assumption that there is no unobserved confounding (NUC). Let Bias_{t_1, t_2} denote the associated omitted confounder bias of $\check{\tau}_{t_1, t_2}$.

We posit the following model for conditionally Gaussian outcomes and Gaussian treatments:

$$U = \epsilon_u, \quad \epsilon_u \sim N_m(0, I_m), \quad (2)$$

$$T = BU + \epsilon_{t|u}, \quad \epsilon_{t|u} \sim N_k(0, \sigma_{t|u}^2 I_k), \quad (3)$$

$$Y = g(T) + \gamma'U + \epsilon_{y|t,u}, \quad \epsilon_{y|t,u} \sim N(0, \sigma_{y|t,u}^2), \quad (4)$$

where B is a $k \times m$ matrix characterizing the relationship between multiple treatments and confounders, g is an arbitrary function from $\mathbb{R}^k \rightarrow \mathbb{R}$ and γ is an m -vector characterizing the effect of m confounders on the outcome. We take m to be fixed and known. There may be information in the data about m , e.g. from a scree plot, although in practice it is important to assess sensitivity to m as well.

For this model, although U itself is not recoverable, the conditional confounder density $f(u | T = t)$, which reflects the relationship between confounders and treatments, is identifiable up to a causal equivalence class (Zheng et al., 2021). According to Equations 2–3, $U | T = t \sim N(\mu_{u|t}, \Sigma_{u|t})$, where $\mu_{u|t} = B'(BB' + \sigma_{t|u}^2 I_k)^{-1}t$ and $\Sigma_{u|t} = I - B'(BB' + \sigma_{t|u}^2 I_k)^{-1}B$.

Zheng et al. (2021) show that γ , which characterizes the relationship between confounders and the outcome, is not identified, although its magnitude is bounded by the constraint

$$\gamma' \Sigma_{u|t} \gamma \leq \sigma_{y|t}^2 \quad (5)$$

where $\sigma_{y|t}^2 = \text{Var}(Y | T = t)$ is identifiable from observed data. This leads to a bound on the omitted confounder bias.

Let $\mu_{u|\Delta t}$ be shorthand for the difference in confounder means in each treatment arm, $\mu_{u|t_1} - \mu_{u|t_2}$. Then,

for any given treatment vectors t_1 and t_2 , we have $\text{Bias}_{t_1, t_2} = \gamma' \mu_{u|\Delta t}$ where

$$\text{Bias}_{t_1, t_2}^2 \leq \sigma_{y|t}^2 R_{Y \sim U|T}^2 \|\Sigma_{u|t}^{-1/2} \mu_{u|\Delta t}\|_2^2. \quad (6)$$

and the bound is achieved when γ is collinear with $\Sigma_{u|t}^{-1} \mu_{u|\Delta t}$ (Zheng et al., 2021).

As such, for all t_1 and t_2 , τ_{t_1, t_2} is partially identified in the interval $\check{\tau}_{t_1, t_2} \pm \sqrt{\sigma_{y|t}^2 R_{Y \sim U|T}^2 \|\Sigma_{u|t}^{-1/2} \mu_{u|\Delta t}\|_2^2}$. Establishing this partial identification region is essential for characterizing how further assumptions about treatment effects and/or sensitivity parameters can further inform estimates of the causal effects.

2.2 Model Parameterizations and Prior Specifications

Choosing an appropriate model parameterization is important in Bayesian inference, especially in partially identified models, because of the need to specify well-motivated prior distributions for all model parameters. However, as argued above, one may have many competing priorities in choosing this parameterization (Gelman, 2004).

For example, consider the linear outcome model, where $g(t) = \beta' t$ so that $\tau_{t_1, t_2} = \beta'(t_1 - t_2)$. In the transparent parameterization of the multi-treatment problem, we partition the parameters into two sets (ϕ, λ) so that the distribution of the observed data depends on ϕ but not λ (Gustafson, 2015). With the linear outcome model, $\phi = \{\check{\beta}, \sigma_{y|t}^2, B, \sigma_{t|u}^2\}$ and $\lambda = \{\gamma\}$ is the transparent parameterization because the factor model parameters are identifiable and

$$f(y | T = t) \sim N(\check{\beta}' t, \sigma_{y|t}^2) \quad (7)$$

where $\check{\beta}' t = \beta' t + \gamma' \mu_{u|t}$ and $\sigma_{y|t}^2 = \sigma_{y|t, u}^2 + \gamma' \Sigma_{u|t} \gamma$ are the regression coefficients and residual variance respectively under NUC. Both are identifiable parameters from the observed data, whereas γ is unidentified. The transparent parameterization cleanly separates issues of model fit, which are governed entirely by ϕ , from assumptions about confounding, governed by λ (Franks et al., 2019).

On the other hand, it is more difficult to elicit meaningful prior knowledge about the parameters in the transparent parameterization, which may be nontrivial functions of the most interpretable scientific quantities. In the linear model, for example, $\check{\beta}$, is a convolution of the true effects and confounding bias, and thus hard to reason about directly. The alternative is to consider a scientific parameterization, which, for example, would be the obvious parameterization if the problem were fully identified and all confounders were

observed. It is easier to elicit prior knowledge under a scientific parameterization, since by construction the parameters correspond to the most scientifically meaningful quantities. For multi-treatment causal inference, the natural scientific parameterization would be to specify prior distributions directly on the treatment effects, which, as shown, are not identified in general.

Expressing prior knowledge about τ_{t_1, t_2} directly can be a useful strategy for encoding plausible assumptions which leads to further concentration of the posterior distribution within the partial identification region (Gustafson et al., 2014). Unfortunately, priors on scientific parameters often constrain the model fit in nontrivial ways, so that it can be difficult to ascertain the degree to which the posterior is being influenced by the data, rather than the prior. These constraints can also lead to model misfit. We characterize some of these issues for certain prior choices below, and then argue that some of these problems can be largely alleviated by comparing posteriors to bounds derived under the transparent parameterization.

3 ADDITIONAL PRIOR ASSUMPTIONS

The key partial identification assumption we make throughout this work beyond latent ignorability, overlap, and SUTVA is that model 2-4 holds, so that all unobserved confounding is reflected in the factor model for treatments (see the Discussion for more about how this model can be further relaxed).

In the proposed model, the treatment effects are bounded, but these bounds are often too large to be practically useful. We advocate for exploring the sensitivity of causal effects to a range of additional scientifically motivated assumptions, which can further reduce the size of the partial identification region. We consider three different types of additional assumptions. First, we explore approaches based on reasoning about the fraction of outcome variance explained by confounders, as in Cinelli and Hazlett (2019), Franks et al. (2019), and Veitch and Zaveri (2020). Second, we explore the role of null controls in reducing the size of the partial identification region (Shi et al., 2020). Finally, we explore how to encode assumptions about sparse treatment effects in the Bayesian paradigm. Assumptions about the fraction of outcome variance explained by confounding are untestable (i.e., have no observable implications), whereas assumptions about null controls and sparsity are partially testable (i.e., have a mixture of observable and unobservable implications).

3.1 Partial R-Squared Assumptions

If confounders U were observed and corresponded to meaningful features, then it would be natural to express prior knowledge about these confounders via a prior distribution directly on the regression coefficients γ . However, because U is not observed, its distribution is only identifiable up to a causal equivalence class defined by rotations (see Zheng et al., 2021). As we will show in our simulation study, directly parameterizing the model in terms of γ and using a flat prior leads to an *a posteriori* bias in the assumed strength of confounding.

Instead, we can parameterize γ as

$$\gamma = \sigma_{y|t} \sqrt{R_{Y \sim U|T}^2} \Sigma_{u|t}^{-1/2} d, \quad (8)$$

where the direction $d \in \mathcal{S}^{m-1}$ is a random unit vector on the $(m-1)$ -sphere. $\sigma_{y|t}$ and $\Sigma_{u|t}$ are identifiable, whereas d and $R_{Y \sim U|T}^2$ are not. Here, $R_{Y \sim U|T}^2$ is the partial fraction of outcome variance explained by confounding given treatments and can easily be reasoned about.

There are several strategies for calibrating the value of $R_{Y \sim U|T}^2$ against related observable or hypothetical quantities, e.g. the fraction of variance explained by observed treatments or confounders (Zhang et al., 2019; Cinelli and Hazlett, 2019). We can use these strategies to bound or condition on values of $R_{Y \sim U|T}^2$ which further reduce our ignorance of τ .

Using domain knowledge to inform the prior distribution for the direction d is more difficult. Specifically, since the distribution of U is only identifiable up to rotation, the only seemingly natural prior for d is the rotationally invariant uniform prior over the $(m-1)$ -sphere. However, the following Proposition shows that the uniform prior over the direction has potentially problematic implications about the posterior bias of the naive estimates.

Proposition 1. *Let the direction, d , of γ have a uniform prior distribution on the $(m-1)$ -sphere. Then, conditional on all other parameters,*

$$\text{Bias}_{t_1, t_2} = 2\sigma_{y|t} \sqrt{R_{Y \sim U|T}^2} \|\Sigma_{u|t}^{-1/2} \mu_{u|\Delta t}\|_2 (Z - 1/2)$$

where $Z \sim \text{Beta}((m-1)/2, (m-1)/2)$.

Proof: See Appendix.

In comparison to the maximum bias (Equation 6), this Proposition shows that for models with many confounders, m , the confounding bias concentrates near zero for any causal contrast. Thus, the seemingly-natural uniform prior translates *a posteriori* to a strong assumption about confounding, which often

contradicts real knowledge about the expected magnitude of confounding bias for contrasts of interest. As such, unless additional knowledge can be used to inform the direction of d (see the next two Sections), we recommend fixing d to the vector which maximizes the worst-case bias in each treatment contrast.

3.2 Negative Control Assumptions

In some applications, it is more intuitive to express prior information in terms of direct structural assumptions about the treatment effects. Here, we consider how assumptions about so-called negative control treatments impact the partial identification region for treatment effects τ_{t_1, t_2} . Negative control treatments are a set of treatments which are assumed to have no causal effect on the outcome and which can be used to detect and (partially) remove bias (Shi et al., 2020). If the i th treatment is assumed to be a negative control treatment, then in Equation 4 we have the additional assumption that $g(t) = g(t_{-i})$ where, with some abuse of notation, t_{-i} denotes the vector t with the i th entry removed. In the linear model, a negative control treatment assumption is equivalent to an assumption that a specific entry of β is zero.

Others have explored inference with negative controls in the assumed linear model but focused entirely on identification under a stronger set of assumptions (Gagnon-Bartsch and Speed, 2012). Here, we demonstrate that even when a set of negative controls is insufficient to identify effects, they still reduce the size of the partial identification region by inducing constraints on the unknown vector γ .

Let $\mathcal{C} = \{(t_1^{(1)}, t_2^{(1)}), \dots, (t_1^{(c)}, t_2^{(c)})\}$ be a set of c negative control treatment contrast pairs, such that $\tau_{\mathcal{C}} = [\tau_{t_1^{(1)}, t_2^{(1)}}], \dots, \tau_{t_1^{(c)}, t_2^{(c)}}] = [0, 0, \dots, 0]$ is a c -vector of effects assumed *a priori* to be zero. For these negative control contrasts, the vector of *observed* mean difference in outcomes, $\check{\tau}_{\mathcal{C}}$, must equal the omitted confounder bias. Since the bias is a function of the sensitivity vector, γ , we can establish a constraint on γ via the equation

$$\gamma' \Sigma_{u|t}^{1/2} M_{u|\Delta t_{\mathcal{C}}} = \check{\tau}_{\mathcal{C}} \quad (9)$$

where $M_{u|\Delta t_{\mathcal{C}}} = \Sigma_{u|t}^{-1/2} [\mu_{u|\Delta t_1}, \dots, \mu_{u|\Delta t_c}]$ is a $m \times c$ matrix with columns corresponding to the scaled difference in confounder means for each negative control contrast and we denote $\mu_{u|\Delta t_i} = \mu_{u|t_1^{(i)}} - \mu_{u|t_2^{(i)}}$ for short. We refer to Equation 9 as the “negative control compatibility condition” because it reveals a testable constraint that $\check{\tau}_{\mathcal{C}}$ must be in the row space of $M_{u|\Delta t_{\mathcal{C}}}$ in order for the negative control assumption to be valid. Moreover, in order to make up the difference between the observed effect and the true null

effect of the control treatments, the magnitude of γ must be large enough. That is, constraint 9 implies a lower bound on the fraction of variance explained by confounders. We formalize these ideas in the following proposition.

Proposition 2. *Suppose there are c known negative control treatment contrasts. Then, $\tilde{\tau}_C$ must be in the row space of $M_{u|\Delta t_c}$. Additionally, the partial fraction of variance explained due to confounders given treatments is lower bounded by*

$$R_{Y \sim U|T}^2 \geq R_{\min}^2 := \frac{1}{\sigma_{y|t}^2} \|\tilde{\tau}_C M_{u|\Delta t_c}^\dagger\|_2^2.$$

where $M_{u|\Delta t_c}^\dagger$ denotes a generalized inverse of $M_{u|\Delta t_c}$. *Proof:* See Appendix.

Although the null controls assumption implies a lower bound, R_{\min}^2 , on the magnitude of confounding, negative controls actually reduce the width of the partial identification region, which is no longer centered at $\tilde{\tau}_{t_1, t_2}$. We characterize the partial identification region under the null controls assumption in the following Theorem.

Theorem 1. *For any value of $R_{Y \sim U|T}^2 > R_{\min}^2$ which satisfies negative control compatibility condition, the confounding bias for the treatment effect of contrast t_1 versus t_2 is in the interval*

$$\tilde{\tau}_C M_{u|\Delta t_c}^\dagger \Sigma_{u|t}^{-1/2} \mu_{u|\Delta t} \pm \quad (10)$$

$$\sigma_{y|t} \sqrt{R_{Y \sim U|T}^2 - R_{\min}^2} \left\| P_{M_{u|\Delta t_c}}^\perp \Sigma_{u|t}^{-1/2} \mu_{u|\Delta t} \right\|_2 \quad (11)$$

where $P_{M_{u|\Delta t_c}}^\perp$ is the $m \times m$ projection matrix into the complement of the column space of $M_{u|\Delta t_c}$.

From Theorem 1 the following corollary follows immediately.

Corollary 1. *Under the assumptions established in Theorem 1, negative controls reduce the size of the partial identification ignorance region by a multiplicative factor of*

$$\sqrt{1 - \frac{R_{\min}^2}{R_{Y \sim U|T}^2} \frac{\|P_{M_{u|\Delta t_c}}^\perp \Sigma_{u|t}^{-1/2} \mu_{u|\Delta t}\|_2}{\|\Sigma_{u|t}^{-1/2} \mu_{u|\Delta t}\|_2}} \leq 1 \quad (12)$$

This corollary highlights that there are two ways in which negative controls reduce the width of the worst-case ignorance region: the first term under the radical shows that negative controls constrain the magnitude of the confounding bias (Proposition 2) which proportionally reduces the width of the ignorance region for all contrasts by an equal amount. The second

term is contrast-dependent and indicates that negative controls reduce the ignorance the most for treatment contrasts that have the most similar confounder mean differences. For a treatment contrast t_1 versus t_2 , when $\Sigma_{u|t}^{-1/2} \mu_{u|\Delta t}$ is in the span of the column space of $M_{u|\Delta t_c}$, then the treatment effect is identified; when $\Sigma_{u|t}^{-1/2} \mu_{u|\Delta t}$ is orthogonal to the column space of $M_{u|\Delta t_c}$ then there is no further reduction in the ignorance region for PATE, beyond the constraint on the magnitude. In summary, the best negative controls are those which have large confounding bias and also have confounder distributions similar to those for the primary treatment contrasts of interest. A direct consequence of Theorem 1 for the linear model is that when $M_{u|\Delta t_c}$ is full rank, all treatments are identifiable.

3.3 Assumptions about Sparsity

Several methods have been proposed for imposing sparsity of the coefficients in high-dimensional regression since the introduction of the LASSO (Tibshirani, 1996; Efron et al., 2004). While approximate sparsity is plausible in many applications, it is more plausible when there is no confounding, since confounding variables often make the *observed* effects dense (e.g. Chandrasekaran et al., 2010). By incorporating information about confounders from a latent variable model, we can potentially adjust for this unobserved confounding and account for sparsity of the causal effects. In fact, Miao et al. (2020) show in the multiple treatment setting that when at least half of the treatments have no effect, then the causal effects can be identified. Importantly, unlike the negative controls assumption, there is no need to pre-specify which treatments are likely to have null effects. This is a testable assumption however, since the joint partial identification region may exclude sparse solutions.

Sparse penalization methods have been adapted for Bayesian inference with the development of sparsifying prior distributions. These prior distribution typically have very heavy tails and large mass close to zero (Carvalho et al., 2009; Ishwaran et al., 2005). In Sections 4 and 5, we use the regularized horseshoe prior for the treatment effect parameters (Piironen et al., 2017), which simultaneously weakly regularizes large effects and shrinks smaller ones to zero, and can be viewed as a continuous analog of the spike and slab prior (Ishwaran et al., 2005). The regularized horseshoe includes hyperparameters which correspond to the expected fraction of non-zero effects as well as the prior variance on the effects, and thus is particularly amenable to prior elicitation.

Sparsity is related to negative controls in important

ways. For example, negative controls assumptions can be viewed as degenerate point mass priors at zero on specific elements of τ_{t_1, t_2} and thus it is natural to consider a relaxation of the negative controls assumption by placing a horseshoe prior on treatments which the practitioners believe are most likely, but not certain, to have no causal effect. In Section 5 we explore this strategy on an analysis of gene expression data.

4 SIMULATED DEMONSTRATION

In this Section, we demonstrate how prior distributions of different types can lead to posterior distributions which concentrate in qualitatively different parts of the partial identification region. In principle, the data provide equal evidence for all points in the partial identification region; thus, concentration in this region can be attributed to the influence of the prior. One key finding is that seemingly innocuous prior assumptions can lead to posterior distributions that concentrate in surprising ways.

We explore these issues in a simple simulation with data generated from a linear outcome model. For the linear model, we let $g(t) = \beta^t t$ and for inference consider four different prior distributions for β and γ : 1) the “gamma parameterization” in which we use an improper flat prior directly on γ and β in Equation 4, 2) the “ R^2 parameterization” in which we place an implicit prior on γ via a proper uniform prior on $R_{Y \sim U|T}^2$ and d in Equation 8 (and again an improper flat prior on β), 3) the R^2 parameterization with one assumed null control, and 4) the R^2 parameterization except with the regularized horseshoe prior used for β . We also explore the partial identification implications for non-parametric models like Bayesian Additive Regression Trees (Chipman et al., 2010) with and without null control assumptions. We fit all linear models using the probabilistic programming language STAN (Stan Development Team, 2014) and the BART model using the R package BART (McCulloch et al., 2018).

The data is generated from model 2-4 in which there are $n = 1000$ observations, $k = 10$ treatments, and $m = 2$ confounders. We fix $\beta = 0$ for all treatments, but due to confounding, $\tilde{\beta} = 1$ for the first five treatments and $\tilde{\beta} = -1$ for the last five treatments. We generate the data so that confounding explains 50% of the residual outcome variance, i.e. $R_{Y \sim U|T}^2 = 0.5$. Lastly, we set B so that the first five rows of B are identical and the last five rows of B are orthogonal to the first five. We consider a range of additional settings in the Appendix (see Figures 5 and 6).

For the Bayesian linear regression models, we compare inferences under the four different assumed pri-

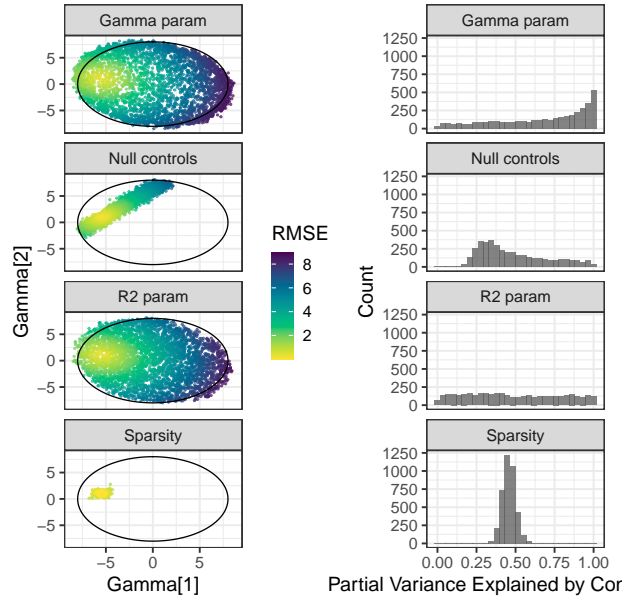


Figure 1: Left) The two dimensions of γ under different prior distributions. Ellipses indicate the theoretical support of γ implied by Equation 5. Right) Distribution of the partial fraction of variance explained by confounders given treatment under the same prior distributions. The true fraction of variance explained is $R_{Y \sim U|T}^2 = 0.5$.

ors. In Figure 1 we plot posterior samples of the two-dimensional sensitivity vector, γ , with points colored according to the RMSE of the inferred causal effects (left). Ellipses indicate the theoretical support of γ implied by Equation 5. On the right panel, we display histograms of the posterior distribution of $R_{Y \sim U|T}^2$. The gamma parameterization is a common default choice when U is observed (i.e. regression with flat priors), but is less reasonable when U is unobserved because the prior induces a posterior that concentrates on the largest values of $R_{Y \sim U|T}^2$ and thus favors models with stronger confounding. We avoid this problem with the R^2 parameterization by separately parameterizing the magnitude and orientation of γ . For the R^2 parameterization, we choose to put a uniform distribution on $R_{Y \sim U|T}^2$, although practitioners can use non-uniform priors or simply condition on a value of $R_{Y \sim U|T}^2$, particularly when domain knowledge is available to facilitate the calibration of this quantity.

When specific information about the magnitudes of the treatment effects are available, we can further constrain inferences. For example, with a single known negative control and $m = 2$ confounders, γ is further constrained to have support which collapses onto a chord of the ellipse (in general, Theorem 1 indicates that the support of γ is the intersection of an

m -dimensional ellipsoid and a hyperplane whose dimension depends on the number of negative controls). The histogram of the posterior distribution of $R_{Y \sim U|T}^2$ also demonstrates how the negative controls induce a lower bound $R_{\min}^2 = 0.33$ on the confounding magnitude $R_{Y \sim U|T}^2$ in this data-generating process (see Proposition 2).

Alternatively, when the magnitudes of the majority of the true treatment effects are known to be sparse, we can constrain inferences with a sparsifying prior, without pre-specifying which effects are null. In this particular example, all the true effects are zero, so the regularizing horseshoe prior leads to a posterior distribution that concentrates on the true value of γ . This is consistent with identification results from Miao et al. (2020), which show that effects are identifiable when more than half of the treatments are negative controls. Note, however, that these assumptions should only be made when well-motivated. In the Appendix, we consider data generating processes for which these assumptions lead to increased bias, despite generating identically distributed observed data (Figure 7). We also include a direct comparison of the joint distribution of different pairs of regression coefficients in Appendix Figure 8.

Finally, we fit the simulated dataset using the nonlinear model, BART. For this model, we focus on estimating

$$\tau_i^x = \tau_{t_i^x, t_i^0} = E[Y | do(T = t_i^x)] - E[Y | do(T = t_i^0)]$$

where t_i^x is the treatment vector where the i th treatment takes on value x and all other treatments are set to 0. In the Appendix Figure 9, we depict the posterior distribution of τ_1^x , τ_2^x and τ_6^x for a range of values, x . On these plots, we include the worst-case bias for different confounding magnitudes $R_{Y \sim U|T}^2$. We compare these plots to the results under the null control assumption that $\tau_1^{-2} = 0$. Since the first two rows of B are identical in the simulated data, fixing the bias of the first contrast via a null control assumption identifies the treatment effects for contrasts of the second treatment. In contrast, the sixth row of B is orthogonal to the first row, and thus the midpoint of the ignorance region remains unchanged. Even though the midpoint doesn’t change, the width of the partial identification region shrinks (see Corollary 1).

5 ANALYSIS OF MICE DATA

In this Section we illustrate a practical workflow for analyzing prior sensitivity in partially identified multi-treatment studies using real data. We reanalyze mice obesity data collected by Wang et al. (2006) and Ghazalpour et al. (2006), integrated by Lin et al. (2015)

and recently analyzed in the multi-treatment context by Miao et al. (2020). The data includes body weight and gene expression levels for 37 genes in each of 227 mice. Miao et al argue that of these 37 genes, only 17 are likely to affect mouse weight, and thus they focus their analysis on these 17 genes only. In contrast, we fit the factor and regression models to all 37 genes, and consider the 20 that were ignored by Miao et al as potential negative controls. In our analysis, we evaluate the robustness of causal effects to a range of additional prior assumptions about the treatment effects and the magnitude of confounding. From the scree plot based on gene expression levels, we determine that an $m = 3$ confounder model is reasonable (Cattell, 1966).

Partial Identification Bounds First, we establish the partial identification region using the transparent parameterization for different strengths of confounding, $R_{Y \sim U|T}^2$. We start by illustrating these bounds for a nonparametric outcome model, by modeling the outcome using Bayesian Additive Regression Tree models (Chipman et al., 2010). Again, we model the observed outcomes conditional on observed treatments using the default prior distributions from the BART package (McCulloch et al., 2018) and infer Bayesian factor model fit using probabilistic programming language STAN (Stan Development Team, 2014).

Here, we let t_i^q be the treatment vector where all genes levels are set to the median except for the i th gene which is set to the q th quantile in the observed population. Then we define a set of estimands for gene i and quantile q as the difference of the intervention mean from the median:

$$\tau_i^q = \tau_{t_i^q, t_i^{0.5}} = E[Y | do(T = t_i^q)] - E[Y | do(T = t_i^{0.5})]$$

With BART, only one gene had 95% posterior credible intervals that did not cover zero for some values of q . In Figure 2 we plot credible intervals for this gene for values of q between 0 and 1 and different degrees of confounding. We can see that posterior intervals cover zero for all $q < 0.7$, even when we assume there is no unobserved confounding. In contrast, we see that for q greater than 0.7, the 95% intervals exclude zero even when confounders explain up to 50% of the residual outcome variance. In short, the evidence suggests that for low levels of Igfbp2, there is no discernible difference from median levels, whereas for high levels of Igfbp2, there is a significant reduction in mouse weight that is robust to confounding at about the $R_{Y \sim U|T}^2 = 0.5$ level.

Sensitivity to Priors We also show how to use the worst-case bounds implied by the transparent parameterization to contextualize the implications of different scientific assumptions. For this analysis, we apply a

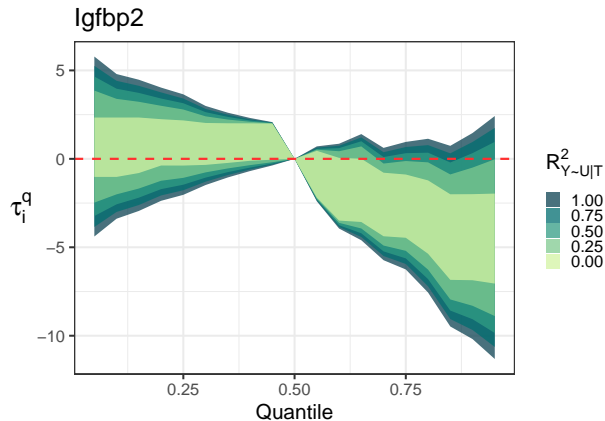


Figure 2: Posterior 95% intervals of τ_i^q for quantiles, q , for different degrees of worst-case confounding parameterized by $R_{Y \sim U|T}^2 \in [0, 1]$. For $q \geq 0.7$, causal effects are significant and robust up to $R_{Y \sim U|T}^2 \approx 0.5$.

standard Bayesian linear regression model to the outcomes, since we have more power under this model to describe variation in the results across genes. We consider several prior specifications on the regression coefficients, β , which correspond to specific treatment effects: flat priors, regularized horseshoe priors and degenerate point mass priors at 0 (i.e. negative controls). For inference with negative controls, we assume the negative controls are the twenty genes that were not considered likely to affect mouse weight by Miao et al. (2020). For the flat priors, we parameterize γ using the R^2 formulation and assume a uniform prior distribution on the sphere, $d \in \mathcal{S}^{m-1}$. We use the R^2 parameterization and assume a uniform prior on $R_{Y \sim U|T}^2 \in [0, k]$ for different upper bounds $k \in 0, 0.5, 1$. All linear models are fit using models implemented in STAN (Stan Development Team, 2014).

Some of the above prior assumptions have testable implications, so we first check whether the data rejects any of these prior specifications. To assess how different models fit the observed data, we use Pareto-Smoothed Importance Sampling estimates of the leave-one-out cross validation loss (Vehtari et al., 2017), and compare differences in expected log predictive density between different prior specifications (See Appendix Table 1). We find that the models specified with sparsifying horseshoe prior have a predictive density within two standard deviations of the models with flat gamma and flat R^2 priors, the latter of which are known to have no observable implications. However, the model fit with negative control constraints showed significantly worse fit. This is consistent with Proposition 2, which indicates that negative control constraints have strong testable implications. In this case, the

negative control assumption implies that the naive effects for the twenty negative controls must lie in the $m = 3$ dimensional row space of $M_{u|\Delta t}$, which they do not. Note that in principle, the data could also reject sparsity assumptions. In this case, there are sparse solutions in the partial identification region, so the horseshoe priors do not seem to reduce the fit of the model.

To relax the rejected negative controls assumption, we fit the linear model with the regularized horseshoe prior applied only to the twenty negative control effects instead of the equivalent implicit degenerate prior at zero for those treatments. The difference in inferences between the strict negative controls assumption and the relaxed horseshoe version is small and can be seen by comparing columns 2 and 3 of Figure 3, where there is agreement in all but 3 of the genes. Importantly, the horseshoe negative controls model fits the observed data better than the strict negative control model and as well as any of the other models.

In Figure 3 we plot the posterior median treatment effect for genes which were significantly different from zero under at least one prior specification. We deem a gene significant if the 95% posterior credible interval excludes zero. *Igfbp2* was found to have a robust and significantly negative effect across all prior specifications, whereas the 20100002N04Rik gene was found to have a significant positive effect on obesity in all but the horseshoe model, for which there were only two significant genes. The first column of Figure 3 corresponds to a standard Bayesian linear regression with flat priors and no unobserved confounding and leads to the largest number of “significant” genes. In Appendix C, we include histograms of the posterior distributions of $R_{Y \sim U|T}^2$ under each prior specification.

In Figure 4 we plot marginal posterior distributions for three different genes across each prior specification and highlight the global bounds of the partial identification regions with red lines. Here, we can see how different prior distributions influence the marginal posterior distributions within the partial identification regions. For genes whose worst-case bounds do not straddle zero, we have confidence that the effect of that genes is robust, even without additional prior assumptions. In other cases, when the posterior concentrates onto a smaller subset of the partial identification region, the significance (or lack thereof) of effects rests on additional unverifiable assumptions. Note that even with “flat priors” under the transparent parameterization for which we might expect the posterior to be uniform in the partial identification region, the posterior concentrates in the middle (see e.g. Proposition 1).

Figure 4 also demonstrates the joint implications of




Figure 3: Analysis of mice obesity data under different prior assumptions about the treatment effects and degree of confounding. Colors reflect the posterior mean and are included only if the 95% posterior credible interval for β excludes zero. All models fit the data equally well, except for the “negative controls” model, for which we see evidence of a violation of the negative controls compatibility condition, Equation 9 (See Appendix Table 1).

the negative controls assumption across genes. Fixing Serpina6 (and others) to zero implies an *a posteriori* negative bias for Socs2 and Igfb2 (relative to the flat priors).

6 CONCLUSION

In this work, we demonstrate the importance of viewing partially identified problems under multiple parameterizations. We establish bounds on causal effects via a transparent parameterization, which then help contextualize how different assumptions under alternative scientific parameterizations constrain the results. Our approach encourages transparency about choices made in an analysis pipeline and can lead to novel insights about the consequences of assumptions




Figure 4: Violin plots for marginal posteriors of treatment effects for three different genes under the various priors. Red lines denote worst-case bounds under the transparent parameterization. The points correspond to posterior samples with the color representing $R^2_{Y \sim U}$.

and robustness of results.

There are still several modeling choices and assumptions that we make in this work that could be further relaxed. First, we illustrate our approach for conditionally Gaussian outcomes and treatments, although the Gaussianity of the treatments and outcomes can be relaxed (see Zheng et al. (2021) for more discussion of partial identification in a more general class of models). There are additional dimensions along which to explore robustness that we did not investigate. For example, we did not explore the robustness of our inferences to, m , the rank of the latent confounder model, which impacts the width of the partial identification regions or scientific parameterizations for more complex nonlinear outcome models.

Finally, in this work, we assume that the outcomes are linear in the confounder and that there are no treatment-confounder interactions. For nonlinear models, we can also consider stronger negative control assumptions about conditional independence between individual treatment variables and the outcome. Exploration of these extensions is left to future work.

Acknowledgements

We would like to thank Victor Veitch for his excellent feedback, as well as the attendees of the ICML 2021

Neglected Assumptions in Causal Inference workshop, the reviewers, and the AC for their comments and suggestions, which improved the work considerably.

References

- Carvalho, C. M., N. G. Polson, and J. G. Scott (2009). Handling sparsity via the horseshoe. In *Artificial Intelligence and Statistics*, pp. 73–80. PMLR.
- Cattell, R. B. (1966). The scree test for the number of factors. *Multivariate behavioral research* 1(2), 245–276.
- Chandrasekaran, V., P. A. Parrilo, and A. S. Willsky (2010). Latent variable graphical model selection via convex optimization. In *2010 48th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pp. 1610–1613. IEEE.
- Chipman, H. A., E. I. George, R. E. McCulloch, et al. (2010). Bart: Bayesian additive regression trees. *The Annals of Applied Statistics* 4(1), 266–298.
- Cinelli, C. and C. Hazlett (2019, 12). Making sense of sensitivity: extending omitted variable bias. *Journal of the Royal Statistical Society Series B (Statistical Methodology)*.
- D’Amour, A. (2019). On multi-cause approaches to causal inference with unobserved confounding: Two cautionary failure cases and a promising alternative. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 3478–3486.
- Efron, B., T. Hastie, I. Johnstone, R. Tibshirani, et al. (2004). Least angle regression. *Annals of statistics* 32(2), 407–499.
- Franks, A., A. D’Amour, and A. Feller (2019). Flexible sensitivity analysis for observational studies without observable implications. *Journal of the American Statistical Association*.
- Gagnon-Bartsch, J. A. and T. P. Speed (2012). Using control genes to correct for unwanted variation in microarray data. *Biostatistics* 13(3), 539–552.
- Gelman, A. (2004). Parameterization and bayesian modeling. *Journal of the American Statistical Association* 99(466), 537–545.
- Ghazalpour, A., S. Doss, B. Zhang, S. Wang, C. Plaisier, R. Castellanos, A. Brozell, E. E. Schadt, T. A. Drake, A. J. Lusic, et al. (2006). Integrating genetic and network analysis to characterize genes related to mouse weight. *PLoS genetics* 2(8), e130.
- Gustafson, P. (2015). *Bayesian inference for partially identified models: Exploring the limits of limited data*, Volume 140. CRC Press.
- Gustafson, P. et al. (2014). Bayesian inference in partially identified models: Is the shape of the posterior distribution useful? *Electronic Journal of Statistics* 8(1), 476–496.
- Ishwaran, H., J. S. Rao, et al. (2005). Spike and slab variable selection: frequentist and bayesian strategies. *Annals of statistics* 33(2), 730–773.
- Lin, W., R. Feng, and H. Li (2015). Regularization methods for high-dimensional instrumental variables regression with an application to genetical genomics. *Journal of the American Statistical Association* 110(509), 270–288.
- McCandless, L. C., P. Gustafson, and A. Levy (2006). Bayesian sensitivity analysis for unmeasured confounding in observational studies. *Statistics in Medicine* 26(11), 2331–2347.
- McCulloch, R., R. Sparapani, R. Gramacy, C. Spanbauer, and M. Pratola (2018). *BART: Bayesian Additive Regression Trees*. R package version 1.9.
- Miao, W., W. Hu, E. L. Ogburn, and X. Zhou (2020). Identifying effects of multiple treatments in the presence of unmeasured confounding.
- Ogburn, E. L., I. Shpitser, and E. J. T. Tchetgen (2020). Counterexamples to” the blessings of multiple causes” by wang and blei. *arXiv preprint arXiv:2001.06555*.
- Pearl, J. (2009). *Causality*. Cambridge university press.
- Piironen, J., A. Vehtari, et al. (2017). Sparsity information and regularization in the horseshoe and other shrinkage priors. *Electronic Journal of Statistics* 11(2), 5018–5051.
- Richardson, T. S., R. J. Evans, and J. M. Robins (2011). Transparent parameterizations of models for potential outcomes. *Bayesian Statistics* 9, 569–610.
- Rubin, D. B. (1980). Comment. *Journal of the American Statistical Association* 75(371), 591–593.
- Shi, X., W. Miao, and E. T. Tchetgen (2020). A selective review of negative control methods in epidemiology. *arXiv preprint arXiv:2009.05641*.
- Stan Development Team (2014). Stan: A C++ library for probability and sampling, version 2.2.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)* 58(1), 267–288.
- Vehtari, A., J. Gabry, M. Magnusson, Y. Yao, P.-C. Bürkner, T. Paananen, and A. Gelman (2020). loo: Efficient leave-one-out cross-validation and waic for bayesian models. R package version 2.4.1.
- Vehtari, A., A. Gelman, and J. Gabry (2017). Practical bayesian model evaluation using leave-one-out cross-validation and waic. *Statistics and computing* 27(5), 1413–1432.

- Veitch, V. and A. Zaveri (2020). Sense and sensitivity analysis: Simple post-hoc analysis of bias due to unobserved confounding. *arXiv preprint arXiv:2003.01747*.
- Wang, S., N. Yehya, E. E. Schadt, H. Wang, T. A. Drake, and A. J. Lusa (2006). Genetic and genomic analysis of a fat mass trait with complex inheritance reveals marked sex specificity. *PLoS genetics* 2(2), e15.
- Wang, Y. and D. M. Blei (2018, May). The Blessings of Multiple Causes. *arXiv e-prints*, arXiv:1805.06826.
- Wang, Y. and D. M. Blei (2019). Multiple causes: A causal graphical view. *arXiv preprint arXiv:1905.12793*.
- Zhang, L., Y. Wang, A. Ostropolets, J. J. Mulgrave, D. M. Blei, and G. Hripcsak (2019). The medical deconfounder: Assessing treatment effects with electronic health records. *arXiv preprint arXiv:1904.02098*.
- Zheng, J., A. D'Amour, and A. Franks (2021). Copula-based sensitivity analysis for multi-treatment causal inference with unobserved confounding. *arXiv preprint arXiv:2102.09412*.

Supplementary Material: Bayesian Inference and Partial Identification in Multi-Treatment Causal Inference with Unobserved Confounding

A Proofs of Theoretical Results

Proof of Proposition 1.

$$\begin{aligned}
\text{Bias}_{t_1, t_2} &= \gamma' \mu_{u|\Delta t} \\
&= \sigma_{y|t} \sqrt{R_{Y \sim U|T}^2} d' \Sigma_{u|t}^{-1/2} \mu_{u|\Delta t} \\
&= \left(\sigma_{y|t}^2 R_{Y \sim U|T}^2 \|\Sigma_{u|t}^{-1/2} \mu_{u|\Delta t}\|_2 \right) d'u
\end{aligned}$$

where $u = \frac{\Sigma_{u|t}^{-1/2} \mu_{u|\Delta t}}{\|\Sigma_{u|t}^{-1/2} \mu_{u|\Delta t}\|_2}$. Since $d = (d_1, \dots, d_m)$ uniformly distributed on the m -sphere, without loss of generality, we can assume that $u = (1, 0, \dots, 0)$ is an m -dimensional unit vector where the first coordinate is 1 and the rest are 0. Then, $d_1 = d'u$ and it suffices to show that $z \sim \text{Beta}((m-1)/2, (m-1)/2)$ where $d_1 = 2(z-1/2)$. We can derive the density by of z by first converting to spherical coordinates, where $d_1 = \cos(\theta_1)$, $d_i = \cos(\theta_i) \prod_{j=2}^{i-1} \sin(\theta_j)$ for $1 < i < m$ and $d_m = \prod_{j=1}^{m-1} \sin(\theta_j)$. The uniform density m-sphere in spherical coordinates is proportional to the spherical volume element which can be computed from the jacobian as $\prod_{i=1}^{M-2} \sin^{M-i-1}(\theta_i)$ for $\theta_1, \dots, \theta_{M-1} \in [0, \pi]^{M-1}$. We derive the density of d_1 as a transformation of the density of θ_1 . Since all angles are independent, $p(\theta_1) \propto \sin(\theta_1)^{m-2}$. Then,

$$\begin{aligned}
p_d(d_1) &\propto p_{\theta_1}(\arccos(d_1)) \left| \frac{d\theta_1}{dd_1} \right| \\
&= \sin(\arccos(d_1))^{m-2} (1 - d_1^2)^{-1/2} \\
&= (1 - d_1^2)^{(m-1)/2} \\
&= (1 - d_1)^{(m-1)/2} (1 + d_1)^{(m-1)/2}
\end{aligned}$$

Changing variables again from d_1 to z , it is immediate that $p_Z(z) \propto (1 - z)^{(m-1)/2} z^{(m-1)/2}$ which is proportional to a $\text{Beta}((m-1)/2, (m-1)/2)$ density.

Proof of Proposition 2, Theorem 1 and Corollary.

Proof. Assume there are c negative control treatment contrasts, satisfying

$$\gamma' \Sigma_{u|t}^{1/2} M_{u|\Delta t_c} = \check{\tau}_c \tag{13}$$

The solution for above equation exists if and only if $\check{\tau}_c M_{u|\Delta t_c}^\dagger M_{u|\Delta t_c} = \check{\tau}_c$ holds, which ensures that the negative control assumptions are compatible. Under this condition, all solutions to equation 9 can be expressed as

$$\gamma' = \tau'_{naive} \Delta t_c M_{u|\Delta t_c}^\dagger \Sigma_{u|t}^{-1/2} + w'(I - M_{u|\Delta t_c} M_{u|\Delta t_c}^\dagger) \Sigma_{u|t}^{-1/2}, \tag{14}$$

Since $\gamma' \Sigma_{u|t} \gamma = \sigma_{y|t}^2 R_{Y \sim U|T}^2$, w can be any $m \times 1$ vector satisfying

$$\begin{aligned} & \| w'(I - M_{u|\Delta t_c} M_{u|\Delta t_c}^\dagger) \|_2^2 = \\ & \sigma_{y|t}^2 R_{Y \sim U|T}^2 - \| \tilde{\tau}_C M_{u|\Delta t_c}^\dagger \|_2^2 \end{aligned}$$

Further,

$$\| w'(I - M_{u|\Delta t_c} M_{u|\Delta t_c}^\dagger) \|_2^2 \geq 0$$

must hold, we know that $R_{Y \sim U|T}^2$ must be at least

$$\frac{1}{\sigma_{y|t}^2} \| \tilde{\tau}_C M_{u|\Delta t_c}^\dagger \|_2^2 .$$

which proves Proposition 2.

For treatment contrast Δt , the omitted variable bias of $\text{PATE}_{\Delta t}$ is $\text{Bias}_{\Delta t} = \gamma' \mu_{u|\Delta t}$ and so

$$\begin{aligned} & \tilde{\tau}_C M_{u|\Delta t_c}^\dagger \Sigma_{u|t}^{-1/2} \mu_{u|\Delta t} \pm \\ & \sigma_{y|t} \sqrt{R_{Y \sim U|T}^2 - R_{\min}^2} \| (I - M_{u|\Delta t_c} M_{u|\Delta t_c}^\dagger) \Sigma_{u|t}^{-1/2} \mu_{u|\Delta t} \|_2, \end{aligned} \tag{15}$$

where the bounds are achieved when $(I - M_{u|\Delta t_c} M_{u|\Delta t_c}^\dagger) w$ has the largest cosine similarity with $\Sigma_{u|t}^{-1/2} \mu_{u|\Delta t}$.

By dividing the right hand side of 15 by the right hand side of 6, we see that the width of ignorance region is shrunk by a multiplicative factor of

$$\sqrt{1 - R_{\min}^2 / R_{Y \sim U|T}^2} \frac{\| (I - M_{u|\Delta t_c} M_{u|\Delta t_c}^\dagger) \Sigma_{u|t}^{-1/2} \mu_{u|\Delta t} \|_2}{\| \Sigma_{u|t}^{-1/2} \mu_{u|\Delta t} \|_2}. \tag{16}$$

B Additional Figures from the Simulation Section

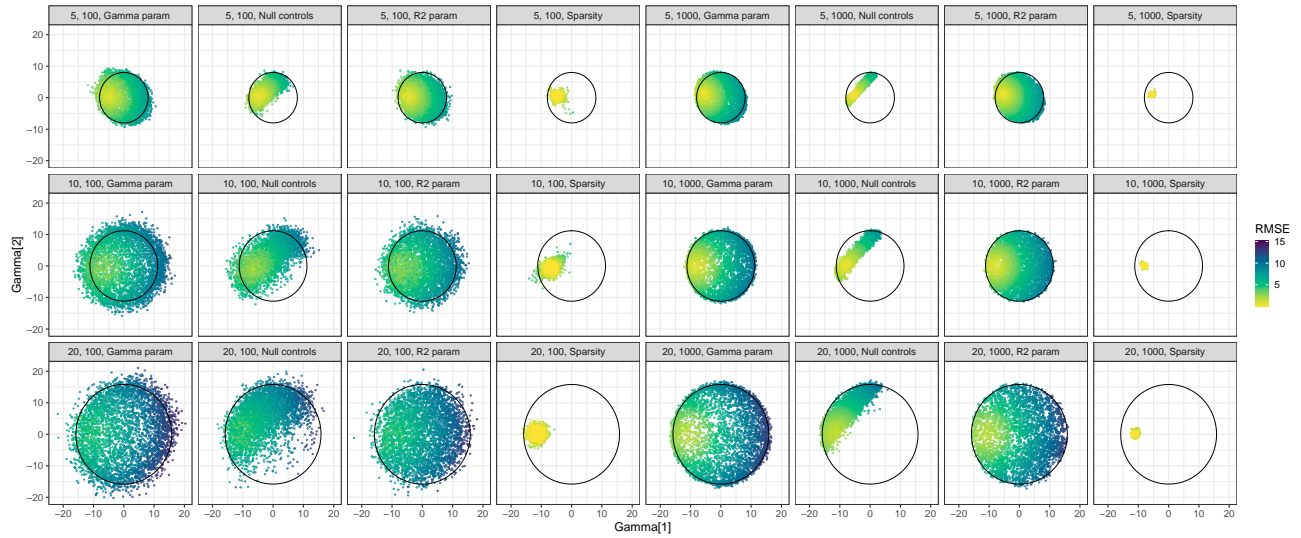


Figure 5: Additional simulation results for $k \in \{5, 10, 20\}$ and $n \in \{100, 1000\}$

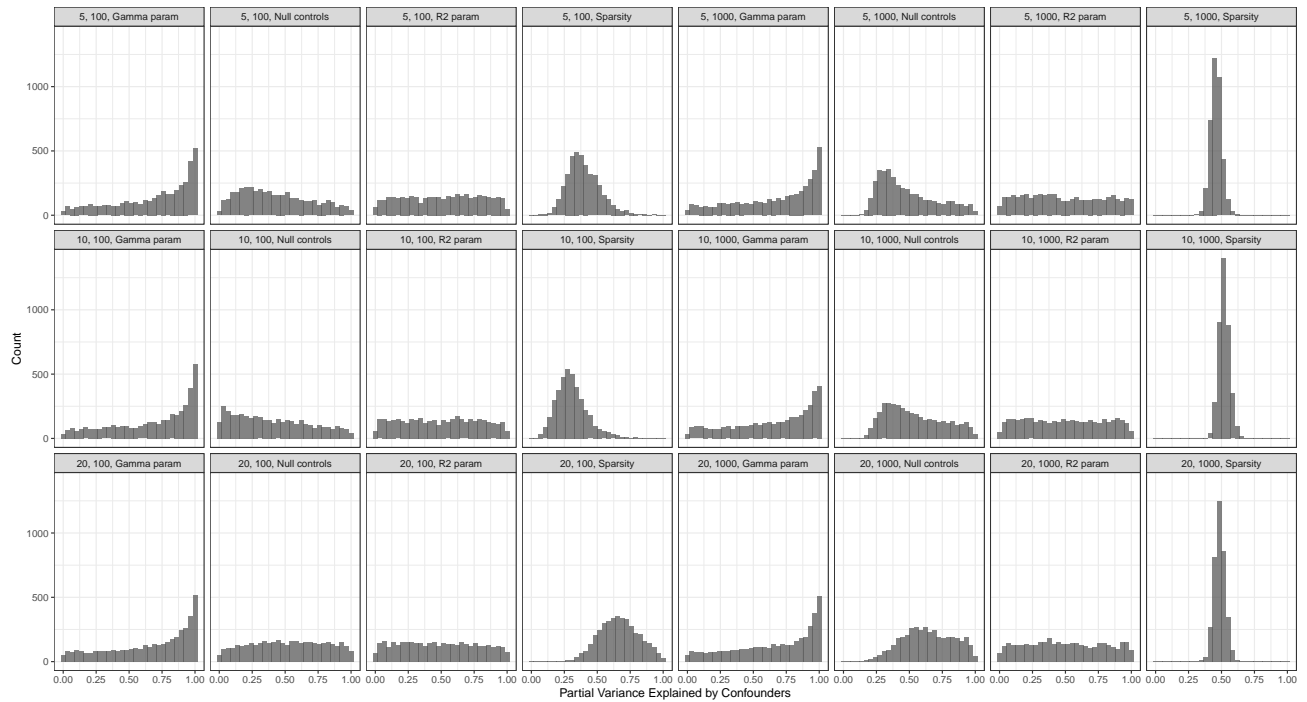


Figure 6: Additional simulation results for $k \in \{5, 10, 20\}$ and $n \in \{100, 1000\}$

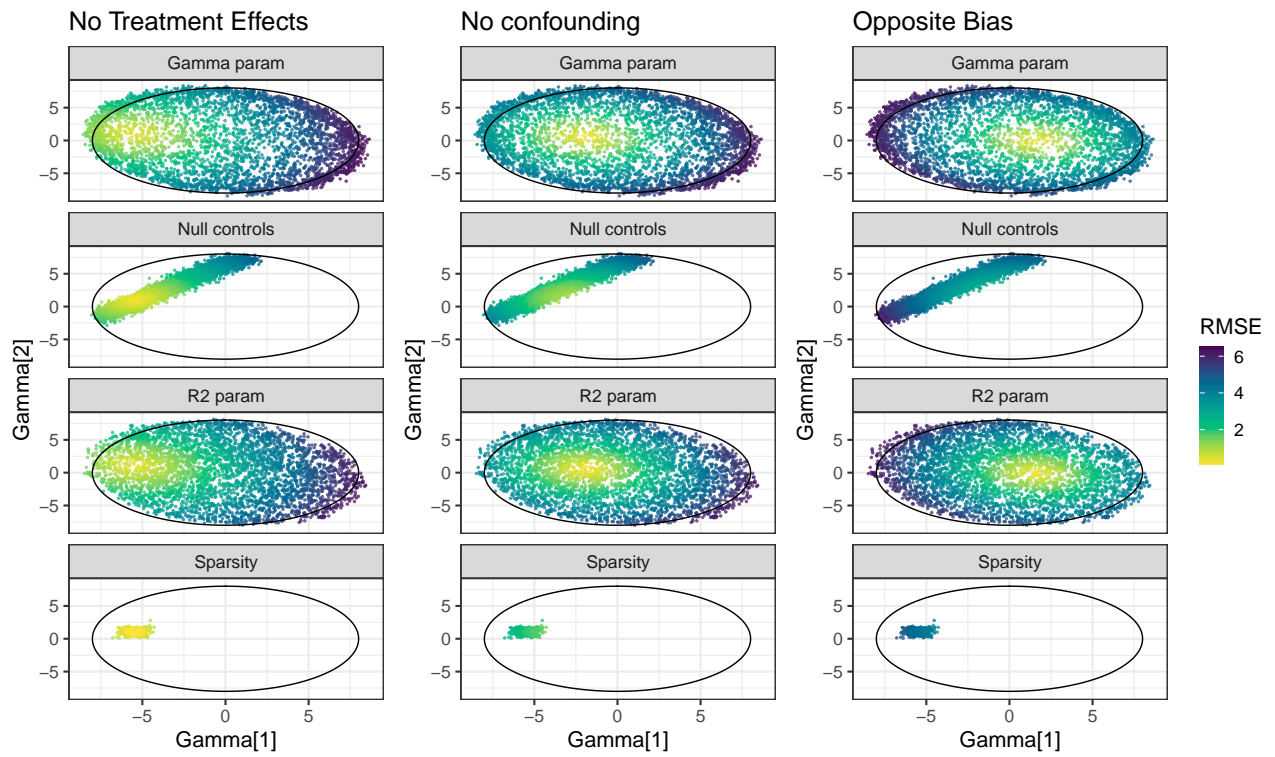


Figure 7: Data generated from the linear model assuming $n = 1000$, $k = 5$ and $m = 2$. We consider three different true effects: one in which there is no true treatment effects (left), one in which there is no confounding ($\check{\beta} = \tau_{t_1, t_2}$ middle), and one in which the bias is opposite to that in the "no treatment effects setting (right).

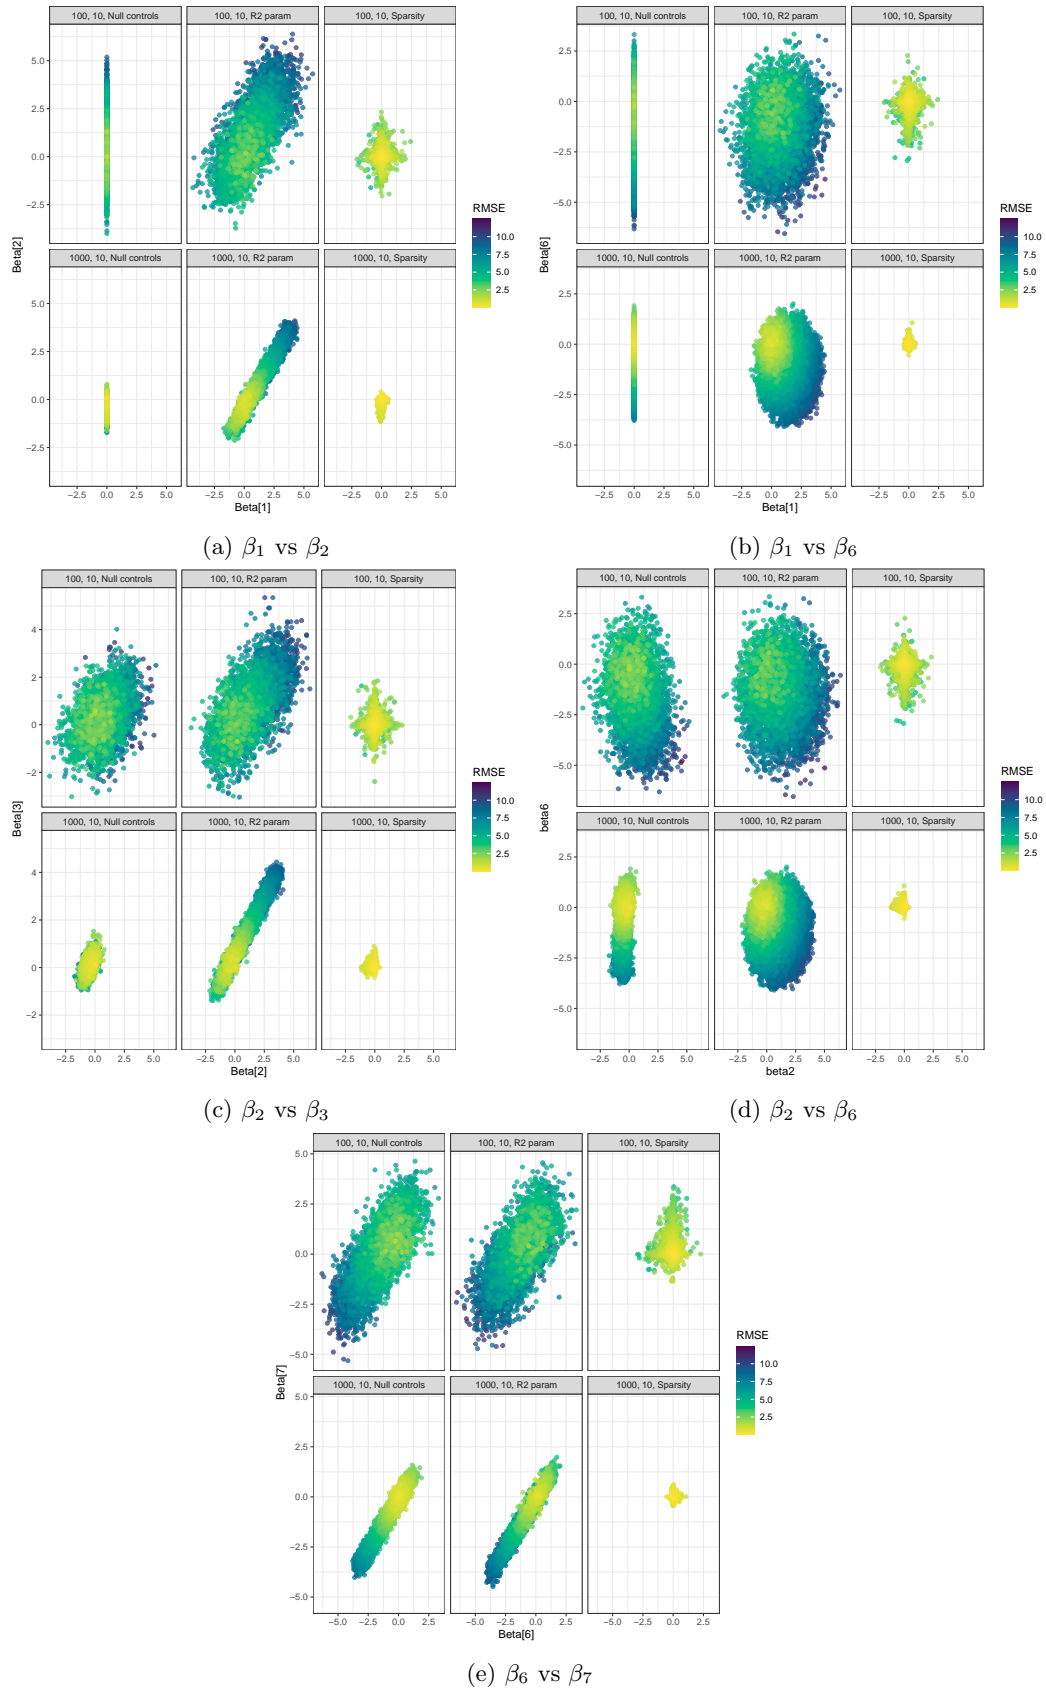


Figure 8: Joint posterior distributions for pairs of causal effects on different treatments under different models.

B.1 BART applied to Simulated Data

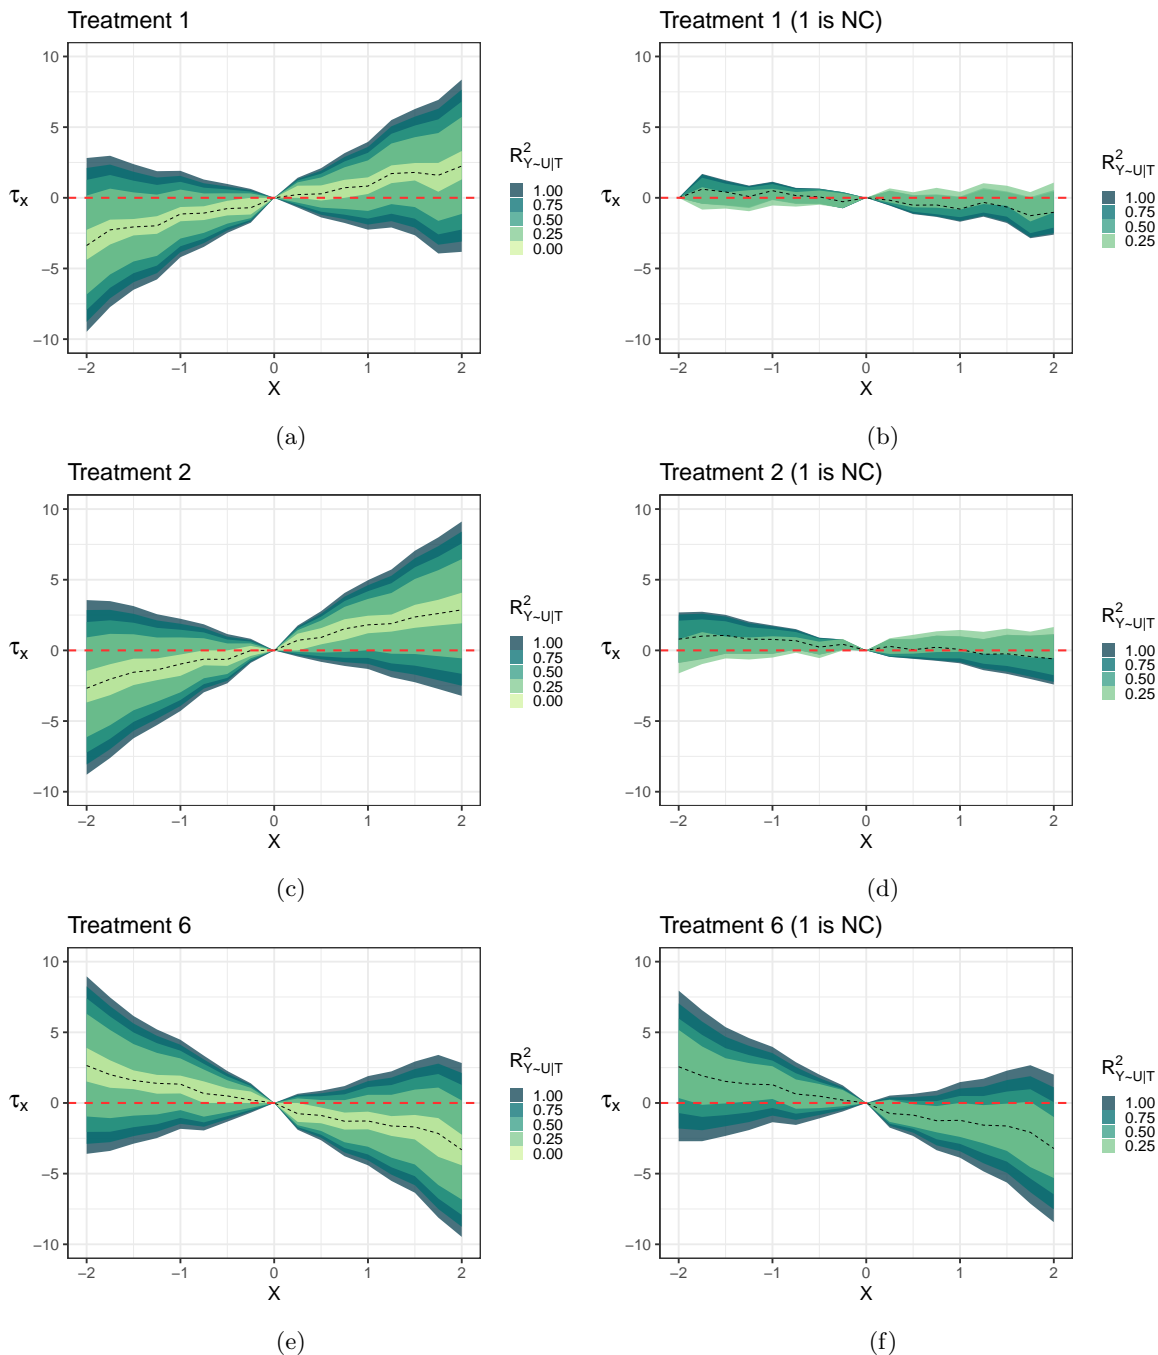


Figure 9: BART analysis of the linear data from Section 4. Dashed black lines represent the midpoint of the ignorance region at each level x of treatment t_i for $i \in \{1, 2, 6\}$. The data on the left depicts the worst-case biases for different values of $R^2_{Y \sim U|T}$ in the transparent parameterization. In the right column, we depict the posterior distribution with worst-case biases under the additional assumption that $\tau_1^{-2} = 0$, i.e. there is no causal difference in the outcome between treatment $(-2, 0, \dots, 0)$ and $(0, 0, \dots, 0)$. The ignorance regions collapse to approximately zero for τ_1^x and τ_2^x for all x , since the first two rows of B are identical. The midpoint of the ignorance regions for the contrast on treatment 6 remain unchanged because the 6th row of B is orthogonal to the first. The size of the ignorance region still shrinks for any value of $R^2_{Y \sim U|T}$ since R^2_{min} % of the outcome variance is due to variables which are nonconfounding for this contrast (See Corollary 1).

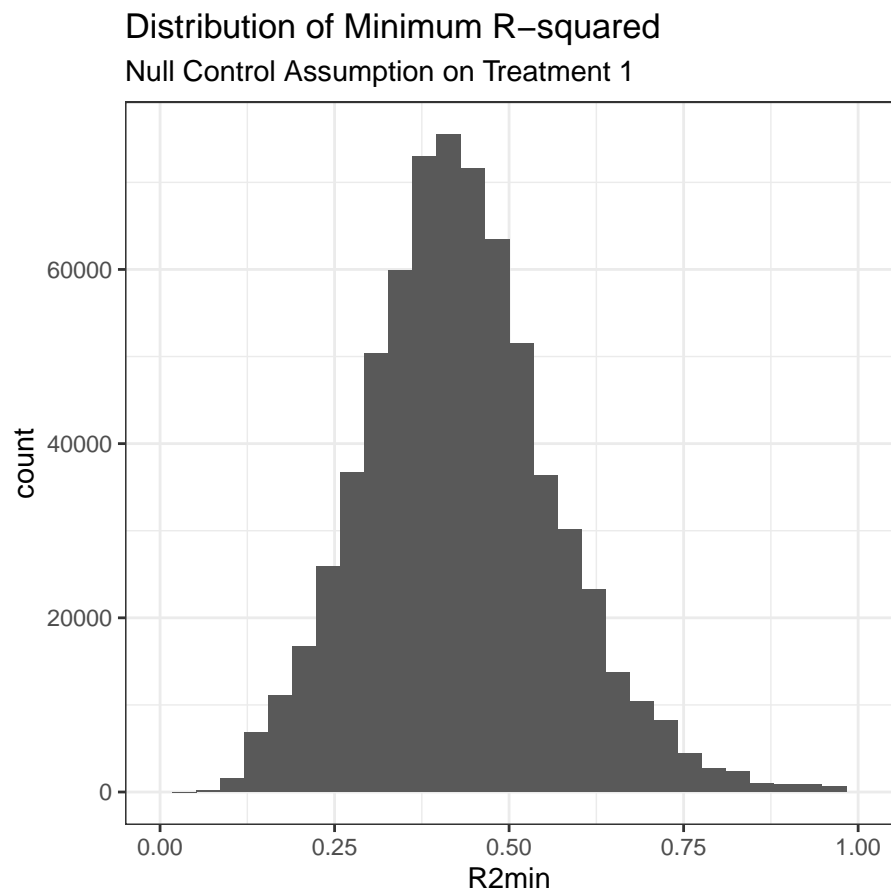


Figure 10: Posterior distribution of R_{\min}^2 for the BART analysis of the simulated data in Section 4, under the null control assumption that $\tau_1^{-2} = 0$.

C Additional Results for the Mice Analysis

	ELPD Difference	SE
Horseshoe	0.00	0.00
Horseshoe (Negative controls only)	-0.14	4.69
Flat prior	-0.35	4.64
Flat priors ($R^2 = 0$)	-1.54	4.91
Horseshoe ($R^2 = 0$)	-1.74	3.35
Negative controls	-16.63	7.26

Table 1: Difference expected log posterior density (ELPD) under various prior distributions and the associated error. Results were computed using `loo` package (Vehtari et al., 2020). All priors except for the negative controls model are within two standard deviation of the estimated expected log posterior density for the top performing model.

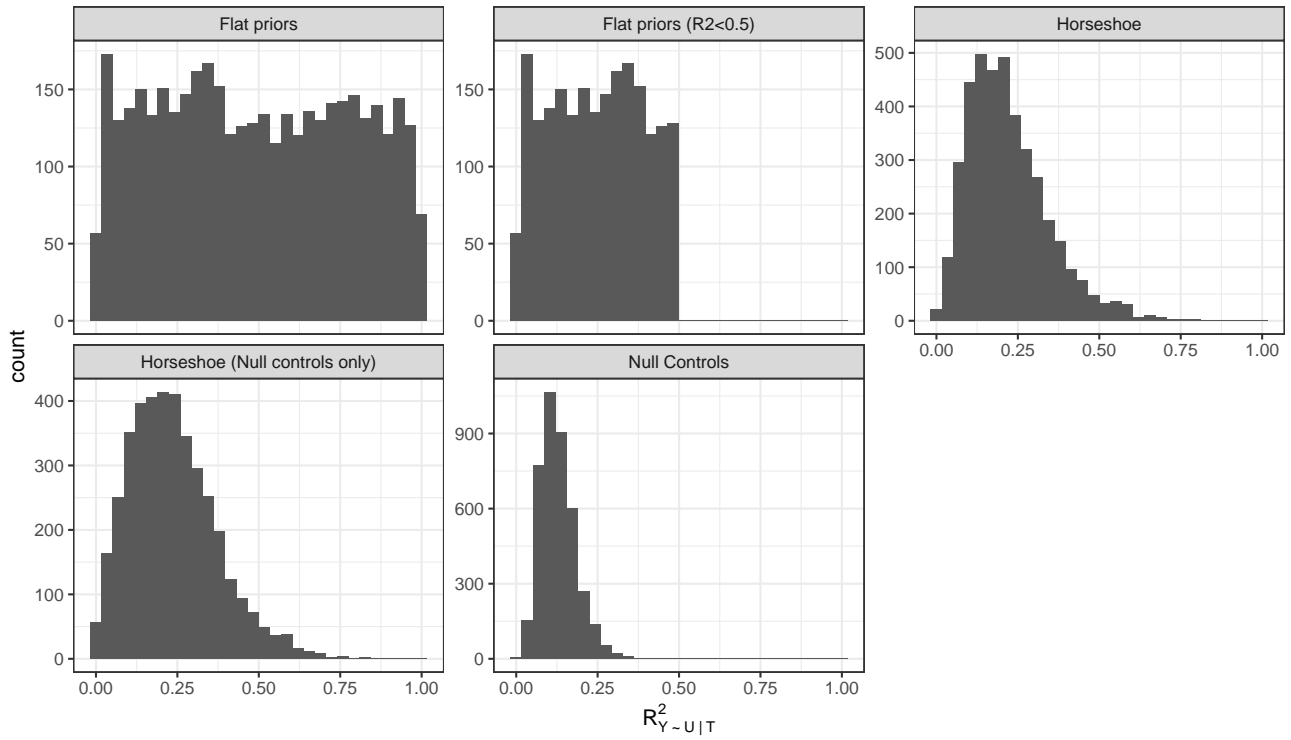


Figure 11: Histograms of $R^2_{Y \sim U|T}$ under the prior specifications chosen in Section 5.