

High-Dimensional Confounding Adjustment Using Continuous Spike and Slab Priors

Joseph Antonelli^{*}, Giovanni Parmigiani[†], and Francesca Dominici[‡]

Abstract. In observational studies, estimation of a causal effect of a treatment on an outcome relies on proper adjustment for confounding. If the number of the potential confounders (p) is larger than the number of observations (n), then direct control for all potential confounders is infeasible. Existing approaches for dimension reduction and penalization are generally aimed at predicting the outcome, and are less suited for estimation of causal effects. Under standard penalization approaches (e.g. Lasso), if a variable X_j is strongly associated with the treatment T but weakly with the outcome Y , the coefficient β_j will be shrunk towards zero thus leading to confounding bias. Under the assumption of a linear model for the outcome and sparsity, we propose continuous spike and slab priors on the regression coefficients β_j corresponding to the potential confounders X_j . Specifically, we introduce a prior distribution that does not heavily shrink to zero the coefficients (β_j s) of the X_j s that are strongly associated with T but weakly associated with Y . We compare our proposed approach to several state of the art methods proposed in the literature. Our proposed approach has the following features: 1) it reduces confounding bias in high dimensional settings; 2) it shrinks towards zero coefficients of instrumental variables; and 3) it achieves good coverages even in small sample sizes. We apply our approach to the National Health and Nutrition Examination Survey (NHANES) data to estimate the causal effects of persistent pesticide exposure on triglyceride levels.

Keywords: high-dimensional data, causal inference, bayesian variable selection, shrinkage priors.

1 Introduction

In observational studies, we are often interested in estimating the causal effect of a treatment T on an outcome Y , which requires proper adjustment of a set of potential confounders \mathbf{X} . In the context of high-dimensional data, where the number of potential measured confounders p could be even larger than the sample size n , standard methods for confounding adjustment such as regression or propensity scores (Rosenbaum and Rubin, 1983) will fail.

In the context of prediction, a variety of methods exist for imposing sparsity in regression models with a high-dimensional set of covariates. Arguably the most popular,

^{*}Department of Statistics, University of Florida, 102 Griffin-Floyd Hall, P.O. Box 118545, Gainesville, FL, 32611, USA, jantonelli@ufl.edu

[†]Department of Biostatistics and Computational Biology, CLS 11007, Dana-Farber Cancer Institute, 450 Brookline Ave, Boston, MA, 02215, USA, gp@jimmy.harvard.edu

[‡]Department of Biostatistics, Harvard T.H. Chan School of Public Health, 677 Huntington Avenue, Boston, MA, 02115, USA, fdominic@hsph.harvard.edu

the lasso (Tibshirani, 1996) places a penalty on the absolute value of the coefficients from a regression model, thus shrinking many of them to be exactly zero, leading to a more parsimonious model. A variety of extensions to the lasso have been proposed such as the SCAD, elastic net, and adaptive lasso penalties, to name a few (Fan and Li, 2001; Zou and Hastie, 2005; Zou, 2006). One challenge encountered with all of these approaches is the difficulty to provide a meaningful assessment of uncertainty around estimates of the regression coefficients. While progress has been made recently on this topic (Lockhart et al., 2014; Taylor and Tibshirani, 2016), it remains difficult to obtain valid confidence intervals for parameters under complex, high-dimensional models.

Bayesian models can alleviate these issues by providing valid inference from posterior samples. Much of the recent work has centered around shrinkage priors, which can be represented as scale mixtures of Gaussian distributions and allow for straightforward posterior sampling. Park and Casella (2008) introduced the Bayesian lasso: a scale mixture of Gaussians with an exponential mixing distribution that induces wider tails than a standard normal prior. More recently, global-local shrinkage priors have been advocated that have a global shrinking parameter that applies to all parameters, as well as local shrinking parameters which are unique to the individual coefficients. Carvalho et al. (2010) introduced the horseshoe prior, which is a scaled mixture of Gaussians with a half-Cauchy mixing distribution that has been shown empirically to have good performance in high-dimensional settings. Bhattacharya et al. (2015) introduced a new class of distributions that are also scaled mixtures of Gaussians with an additional Dirichlet mixing component and proved that its posterior concentrates at the optimal rate. Ročková and George (2016) differ somewhat in that they adopt the spike and slab formulation of George and McCulloch (1993), however both the spike and slab priors are Laplace distributions. All of these approaches are aimed at obtaining ideal amounts of shrinkage in high-dimensional settings where large coefficients should be shrunk a small amount, while others are shrunk heavily towards zero.

These and many other approaches have been based on the same principles of aiming to reduce shrinkage for important covariates in the context of prediction of Y . Several authors have pointed out that frequentist and Bayesian procedures for variable selection or shrinkage that focus on predicting Y perform poorly when the inferential goal is estimation of the effect of T on Y (Crainiceanu et al., 2008; Wang et al., 2012; Belloni et al., 2014, 2017; Hahn et al., 2017). A variety of data driven methods have been developed to select confounders in causal inference (van der Laan and Gruber, 2010; De Luna et al., 2011; Vansteelandt et al., 2012; Wang et al., 2012; Zigler and Dominici, 2014). Many of these approaches rely on the specification of a treatment model $E(T|\mathbf{X})$, and an outcome model $E(Y|T, \mathbf{X})$. Wang et al. (2012) introduced a Bayesian model averaging approach for estimating the effect of T on Y averaged across models that include different sets of potential confounders. They assume a priori that if a covariate X_j is associated with the treatment T then this covariate should have high probability to be included into the outcome model, even if this covariate is weakly associated with the outcome. Many ideas have been built on this prior specification to address the issue of confounder selection and model uncertainty (Talbot et al., 2015; Wang et al., 2015; Cefalu et al., 2016; Antonelli et al., 2017). All of the aforementioned approaches have been shown to work well in identifying confounders or adjusting for confounding;

however, none of these approaches are well-suited to a high-dimensional vector of confounders. Recently, there has been increased attention to estimating treatment effects when $p \geq n$. Wilson and Reich (2014) introduced a decision theoretic approach to confounder selection for $p \geq n$. They showed that their approach has connections to the adaptive lasso, but with weights aimed at reducing shrinkage of confounders, rather than predictors. Belloni et al. (2014) applied standard lasso models on both the treatment model $E(T|\mathbf{X})$, and the outcome model $E(Y|T, \mathbf{X})$, separately. Then, they identify as confounders the union of the variables that were not shrunk to zero in the two models, and estimate the causal effect using this reduced set of covariates. Farrell (2015) first applies lasso models to treatment and outcome models to select confounders. Second, they calculate a double robust estimator using the resulting unpenalized treatment and outcome models. Antonelli et al. (2018) implemented a similar doubly robust estimation approach using standard lasso outcome and treatment models but in context of matching on both the propensity and prognostic scores. Ertefaie et al. (2015) proposed an alternative approach for selecting confounders in high-dimensional settings by penalizing a joint likelihood for both the treatment and the outcome model to ultimately lead to the selection of important confounders. Shortreed and Ertefaie (2017) used similar ideas by fitting an adaptive lasso to a propensity score model, and show that it leads to the inclusion of only covariates necessary for confounding adjustment or outcome model prediction. Hahn et al. (2016) utilized horseshoe priors on a re-parameterized likelihood that aims to reduce shrinkage for important confounders. Athey et al. (2016) combined high-dimensional regression with the balancing weights of Zubizarreta (2015) to obtain valid inference of treatment effects even when the true data generating models are not sparse. All of these approaches have the advantage of being able to handle settings with $p \geq n$, however, as we will demonstrate in simulations, existing approaches relying on asymptotic theory can provide coverage below the nominal level in finite samples. Also, many of these approaches will tend to include instrumental variables, are not applicable to studying the effects of continuous treatments (see Section 5), or both.

In Section 2 we propose spike and slab priors for confounding adjustment in the context of homogeneous and heterogeneous treatment effects. In Section 3, we detail the Bayesian computations, including the selection of tuning parameters to achieve a good compromise between sparsity and addressing confounding bias. In Section 4, we present results from several simulation studies that include homogeneous and heterogeneous treatment effects, strong and weak confounders, instrumental variables, and sparse and non sparse settings. In Section 5, we present the data analysis which considers continuous treatments. In Section 6 we conclude with a summary of the strengths and weaknesses of the proposed approach and future research directions.

2 Spike and slab priors for confounding adjustment

Throughout, we will assume that we observe $\mathbf{D}_i = (Y_i, T_i, \mathbf{X}_i)$ for i, \dots, n , where n is the sample size of the observed data, Y_i is the outcome, T_i is the treatment, and \mathbf{X}_i is a p -dimensional vector of pre-treatment covariates for subject i . We will assume for simplicity that Y_i is continuous, though we will not make assumptions regarding T_i , as it can be binary, continuous, or categorical. The extension to binary outcomes is

straightforward using latent variable techniques introduced in Albert and Chib (1993). In general we will be working under the high-dimensional scenario of $p \geq n$, where we let $p \rightarrow \infty$. Our estimand of interest is the average treatment effect (ATE), defined as $\Delta(t_1, t_2) = E(Y(t_1) - Y(t_2))$, where $Y_i(t)$ is the potential outcome subject i would receive under treatment t . We will assume that the probability of receiving any value of treatment is greater than 0 for any combination of the covariates, commonly referred to as positivity. We make the stable unit treatment value assumption (SUTVA) (Little and Rubin, 2000), which states that the treatment received by one observation or unit does not affect the outcomes of other units and the potential outcomes are well-defined. We will further assume strong ignorability conditional on the observed covariates, and that the covariates necessary for ignorability are an unknown subset of \mathbf{X} . Strong ignorability implies that potential outcomes are independent of T conditional on \mathbf{X} .

2.1 Model formulation

In this section we assume a homogeneous treatment effect, i.e. that the treatment effect is the same across all values of the covariates \mathbf{X} . We will relax this assumption in Section 2.5. We introduce the following hierarchical formulation:

$$\begin{aligned} Y_i | T_i, \mathbf{X}_i, \beta_0, \beta_t, \boldsymbol{\beta}, \sigma^2 &\sim \text{Normal}(\beta_0 + \beta_t T_i + \mathbf{X}_i \boldsymbol{\beta}, \sigma^2), \\ P(\boldsymbol{\beta} | \boldsymbol{\gamma}) &= \prod_{j=1}^p \gamma_j \psi_1(\beta_j) + (1 - \gamma_j) \psi_0(\beta_j), \\ P(\boldsymbol{\gamma} | \boldsymbol{\theta}) &= \prod_{j=1}^p \theta^{w_j \gamma_j} (1 - \theta^{w_j})^{1 - \gamma_j}, \\ P(\boldsymbol{\theta} | a, b) &\sim \text{Beta}(a, b), \\ P(\sigma^2 | c, d) &\sim \text{InvGamma}(c, d), \\ P(\beta_0), P(\beta_t) &\sim \text{Normal}(0, K). \end{aligned} \tag{1}$$

Under these assumptions, $\Delta(t_1, t_2) = (t_1 - t_2)\beta_t$. It is straightforward to allow the treatment effect to be nonlinear by replacing $\beta_t T_i$ with $f(T_i)$, which can be approximated using basis functions. If $\gamma_j = 1$ then $\beta_j \sim \psi_1(\beta_j)$ – the slab component of the prior. If $\gamma_j = 0$ then $\beta_j \sim \psi_0(\beta_j)$ – the spike component of the prior. Therefore $\gamma_j = 1$ indicates that X_j is potentially an important confounder. We set $\psi_1(\cdot)$ and $\psi_0(\cdot)$ to be Laplace distributions with densities $\psi_1(\beta_j) = \frac{\lambda_1}{2\sigma} e^{-\frac{\lambda_1 |\beta_j|}{\sigma}}$ and $\psi_0(\beta_j) = \frac{\lambda_0}{2\sigma} e^{-\frac{\lambda_0 |\beta_j|}{\sigma}}$, respectively. More specifically, when $\gamma_j = 1$, the prior standard deviation of β_j is $\sigma\sqrt{2}/\lambda_1$ and when $\gamma_j = 0$ the prior standard deviation is $\sigma\sqrt{2}/\lambda_0$. Scaling the prior variance with σ^2 is common in Bayesian hierarchical models (Park and Casella, 2008) and allows for more stable estimation and better interpretation of λ_1 and λ_0 . Finally, we have θ and w_j , which control the prior probability that $\gamma_j = 1$. The global parameter θ dictates the probability that $\gamma_j = 1$ when $w_j = 1$ and can be thought of as the overall sparsity level in the data. The weights w_j are tuning parameters that we will use to prioritize variables to have $\gamma_j = 1$ if they are also associated with the treatment. We will discuss the selection of w_j in more detail in Sections 2.3–2.4.

2.2 Hyper prior selection

Our prior formulation has a number of hyper parameters, and it is important for these to be set to reasonable values to obtain good inference for the treatment effect of interest. The first hyper parameters are λ_0 and λ_1 , which control the variance of the spike and slab priors. Following Ročková and George (2016), we will fix λ_1 to a small value, say 0.1, so that the prior variance for coefficients in the slab component of the prior is high enough to be reasonably uninformative, and important parameters will not be shrunk heavily towards 0. We assess the sensitivity to this choice in the supplementary materials (Antonelli et al., 2018) and find that results are robust to the choice of λ_1 . Results can be quite sensitive to the choice of λ_0 , therefore we will estimate it using empirical Bayes to let the data determine how much to shrink coefficients that are placed in the spike component of the prior.

The parameters a and b dictate the prior for θ meaning they control the amount of sparsity induced a priori. We will adopt standard practice in the high-dimensional Bayesian literature and set a to a constant, and set $b \propto p$ (Zhou et al., 2015; Ročková and George, 2016). This prior more aggressively shrinks coefficients to the spike component of the prior as p grows. This feature is desirable in high-dimensional models where we must more aggressively shrink parameters as the covariate space grows to avoid the curse of dimensionality (Scott et al., 2010). Throughout the paper we will use $a = 1$ and $b = 0.1p$, though we assessed the sensitivity to these choices in the supplementary materials and found the results are robust to the selection of a and b . We will assume conjugate and uninformative priors for σ^2 , β_0 and β_t . Finally, we must choose the tuning parameter, w_j , which we will use to prioritize potential confounders in the prior formulation. We will discuss the selection of w_j in Section 2.4.

2.3 Probability of inclusion into the slab

To better understand how to select w_j , it helps to understand the probability that a parameter for a given covariate is included in the slab component of the prior. There are two crucial quantities that we can study to gain intuition into whether an important covariate is effectively included in the model. The first is the conditional probability that a parameter is included in the slab component of the prior, which can be defined as follows:

$$p_{\theta}^*(\beta_j) = P(\gamma_j = 1 | \beta_j, \theta) = \frac{\theta^{w_j} \psi_1(\beta_j)}{\theta^{w_j} \psi_1(\beta_j) + (1 - \theta^{w_j}) \psi_0(\beta_j)}.$$

This is also the expression from which we update γ_j in a Gibbs sampler, and therefore gives insight into the probability that a parameter is included in the slab component of the prior. The second quantity involves the posterior mode of our model. As seen in Ročková and George (2016) the posterior mode will be sparse in the sense that many of the parameters will be set exactly to zero. An important quantity in the estimation

of the posterior mode is defined as

$$\Delta_j \equiv \inf_{t>0} \left\{ nt/2 + \frac{\lambda_1 t - \log \left(\frac{p_\theta^*(0)}{p_\theta^*(t)} \right)}{t} \right\}.$$

This expression is important because the posterior mode estimate of $\hat{\beta}_j$ is nonzero if $\mathbf{X}_j'(\mathbf{Y} - \sum_{l \neq j} \mathbf{X}_l \hat{\beta}_l - \mathbf{T} \hat{\beta}_t - \hat{\beta}_0) > \Delta_j$. Now that we have defined these two quantities, we can look at them with respect to w_j to gain some intuition as to how w_j impacts the variable selection and subsequent shrinkage of important parameters. Figure 1 shows these two quantities as a function of w_j when $\lambda_1 = 0.1$, $\lambda_0 = 30$, and $\theta = 0.05$. The left panel shows that covariates strongly associated with the outcome ($\beta_j = 0.4$) always enter the slab regardless of w_j , while variables with no association ($\beta_j = 0$) never enter the slab unless w_j is very close to 0. Covariates with a mild association with the outcome ($\beta_j = 0.2$) change drastically depending on w_j , as small values of w_j lead to inclusion probabilities near 1 and large values of w_j lead to posterior inclusion probabilities near 0. The right panel of Figure 1 shows that the threshold for a parameter having a nonzero posterior mode greatly decreases when w_j is small. Not seen in the figure is that values of w_j greater than 0.1 are essentially the same as $w_j = 1$ in terms of the probability of being nonzero in the posterior mode.

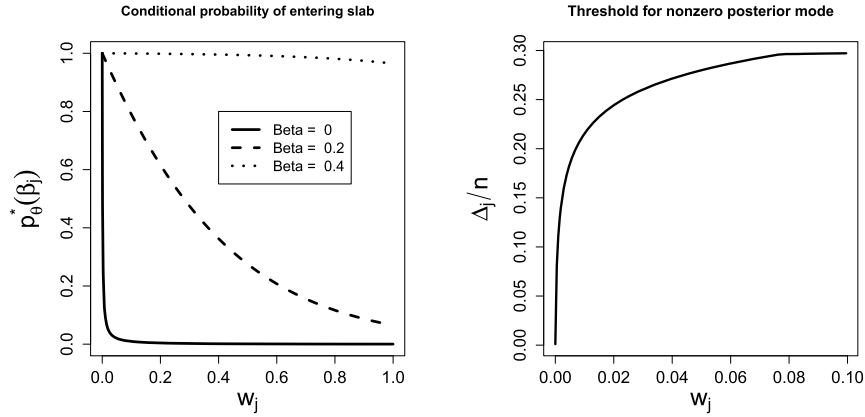


Figure 1: The left panel shows $p_\theta^*(\beta_j)$ for a variety of values of β_j as a function of w_j . The right panel shows Δ_j/n as a function of w_j . Here we fixed $\lambda_1 = 0.1$, $\lambda_0 = 30$, and $\theta = 0.05$.

2.4 Selection of w_j

In the previous section we saw how the weight w_j can impact variable selection as it is decreased towards 0. Our goal is improved estimation of β_t , the treatment effect, and therefore we want to prioritize variables that are also associated with the treatment,

T_i . If a variable X_j is associated with T , then omitting X_j in the outcome model could lead to confounding bias. Consequently, it is desirable to increase the prior probability that β_j is in the slab component of the prior. With this guiding principle in mind, we do the following: 1) we use lasso to fit the exposure model $E(T|\mathbf{X})$ (Tibshirani, 1996); 2) for each X_j that has a non zero regression coefficient from the lasso estimation of the exposure model, we set $w_j = \delta$ where $0 < \delta < 1$. Please note that if $\delta < 1$, then $\theta^\delta > \theta$ which leads to a higher prior probability for β_j to be included into the slab component of the prior. Smaller values of δ lead to more protection against omitting an important confounder. However, values of δ too small might lead to inclusion of instrumental variables which decrease efficiency and can amplify bias in the presence of unmeasured confounding (Pearl, 2011).

Now we provide guidance on how to select a reasonable value of δ . Figure 1 can be used to guide our choice of δ . We assume $w_j = \delta \forall j$, and we want to set δ to be as small as possible to protect us against shrinking the coefficients for important confounders, but we also want to ensure that coefficients for instrumental variables or noise variables are heavily shrunk towards 0. We can see in Figure 1 that we can set δ to be as small as possible to increase $p_\theta^*(\beta_j)$ for variables with moderate associations with the outcome, while still keeping $p_\theta^*(0)$ low. One possibility is to select the minimum value of δ such that $p_\theta^*(0)$ is less than some threshold, such as 0.1. This would imply that the probability of including a parameter for an instrument or noise variable into the slab component of the prior would be 0.1. Intuitively, this threshold represents the point at which we can get the most protection against residual confounding bias while alleviating the impact of instrumental variables.

Additionally, we have found that when the treatment assignment is not sparse, assigning weights of δ to all covariates identified by a treatment model can lead to poor performance in the subsequent outcome model. One approach to these types of issues is to cap the number of variables that are prioritized by the treatment model to k covariates. We will explore a scenario where this is the case in the simulation study of Section 4 and assess the extent that this problem is corrected when we limit the number of variables prioritized.

2.5 Heterogeneous treatment effects

In this section we describe our approach under the more general case of treatment effect heterogeneity and in the context where the treatment variable is binary or categorical. Addressing treatment effect heterogeneity in the presence of continuous treatments is a more difficult problem, which is beyond the scope of this paper. In the case of a binary treatment, we now specify the same model as in Section 2 but separately for $t = 1$ and $t = 0$:

$$Y_i | T_i = t, \mathbf{X}_i, \beta_0^{(t)}, \boldsymbol{\beta}^{(t)}, \sigma^{2(t)} \sim \text{Normal}(\beta_0^{(t)} + \mathbf{X}_i \boldsymbol{\beta}^{(t)}, \sigma^{2(t)}).$$

It is important to note that for each treatment level t , we fit a separate model only using subjects i with $T_i = t$. To estimate the treatment effect in this setting, we can

exploit the fact that the average treatment effect can be estimated as

$$\Delta(t_1, t_2) = \int_{\mathbf{X}} E(Y|T = t_1, \mathbf{X}) - E(Y|T = t_2, \mathbf{X}) dF_n(\mathbf{X}),$$

where $F_n(\mathbf{X})$ is the empirical distribution of the covariates. If we let s denote the s^{th} posterior draw obtained in our Markov chain Monte Carlo (MCMC) algorithm then the posterior mean of the treatment effect can be defined as

$$E(\Delta(t_1, t_2)|\mathbf{D}) = \frac{1}{Sn} \sum_{s=1}^S \sum_{i=1}^n \beta_{0(s)}^{(t_1)} - \beta_{0(s)}^{(t_2)} + \mathbf{X}_i(\beta_{(s)}^{(t_1)} - \beta_{(s)}^{(t_2)}).$$

This provides us with a valid estimate of the treatment effect, however, it does not provide us with a valid credible interval. This estimate is marginalizing over the covariates, and our posterior distribution does not take into account this additional uncertainty, so we will utilize the Bayesian bootstrap (Rubin et al., 1981) to account for it. Specifically, we can define $u_0 = 0$, $u_n = 1$, and u_1 through u_{n-1} to be the order statistics from $n - 1$ draws from a uniform distribution. Then we can define weights, $\xi_i = u_i - u_{i-1}$. We will do this M separate times leading to weights ξ_{mi} for $m = 1, \dots, M$ and $i = 1, \dots, n$. Finally, for each of the S posterior samples and M weight vectors we can calculate

$$\Delta_{(s,m)}(t_1, t_2) = \frac{1}{n} \sum_{i=1}^n \xi_{mi} \left(\beta_{0(s)}^{(t_1)} - \beta_{0(s)}^{(t_2)} + \mathbf{X}_i(\beta_{(s)}^{(t_1)} - \beta_{(s)}^{(t_2)}) \right) \quad (2)$$

and we can use the quantiles of these values to create credible intervals. In brief, we have built separate regression models for each of the treatment levels and then taken the difference in the mean predicted values from these models for each observation in the data. To account for the additional uncertainty from marginalizing over covariates, we randomly re-weighted the data using the Bayesian bootstrap. If we were interested in estimating treatment effects within particular subgroups such as the treatment effect on the treated, then the sum in equation (2) would be over only those subjects of interest. Note that while we have separated the estimation of the models for each treatment group, it is possible to posit a hierarchical model to borrow information between treatment groups. This could exploit the fact that we expect the parameters to be similar across treatment groups, and would amount to shrinking the heterogeneous model closer to the homogeneous model. This sort of shrinkage has been shown to work well in related contexts (Hahn et al., 2017), and merits future research.

3 Bayesian computation

Posterior distributions of all the unknown parameters can be easily obtained via standard Gibbs-sampling as each parameter is conditionally conjugate, with the exception of θ , which is easy to sample from since it is univariate. A key component driving the

ease of sampling is that the Laplace distribution has the following representation as a scale mixture of Gaussians with an exponential mixing weight

$$\frac{\lambda}{2} e^{-\lambda|\beta|} = \int_0^\infty \frac{1}{\sqrt{2\pi\tau^2}} e^{-\beta^2/2\tau^2} \frac{\lambda^2}{2} e^{-\lambda^2\tau^2/2} d\tau.$$

This makes the prior distribution of β multivariate normal with a covariance matrix equal to $\text{diag}(\tau_1^2, \dots, \tau_p^2)$. Full details of this mixture as well as posterior implementation can be found in the supplementary materials. The most important parameter of the procedure is λ_0 , which dictates how strongly parameters are shrunk towards zero when they are included in the spike part of the model, i.e. $\gamma_j = 0$. Bayesian inference allows for viable alternatives over cross validation, which is commonly used in the penalized likelihood literature. We will examine estimation of λ_0 using an empirical Bayes procedure, though it is possible to utilize a fully Bayesian specification that places a prior on λ_0 as discussed in Park and Casella (2008).

3.1 Selection of λ_0

In many complex settings, such as the current one, empirical Bayes estimators of tuning parameters can not be done analytically. To alleviate this issue Casella (2001) proposed a Monte Carlo based approach to finding empirical Bayes estimates of hyperparameter values. The general idea is very similar to the expectation-maximization (EM) algorithm for estimating missing or unknown parameters, however, expectations in the E-step are calculated using draws from a Gibbs Sampler. In our example, we set $\lambda_0^{(0)} = \lambda_0^*$, a starting value of the algorithm. Then for iteration k , set

$$\lambda_0^{(k)} = \sqrt{\frac{2 \left(p - E_{\lambda_0^{(k-1)}} \left(\sum_j \gamma_j \right) \right)}{\sum_j E_{\lambda_0^{(k-1)}} \left(\tau_j^2 1(\gamma_j = 0) \right)}},$$

where the expectations are approximated with averages from the previous iteration's Gibbs Sampler. Due to Monte Carlo error, this algorithm will not exactly converge, but rather will bounce around the maximum likelihood estimate. The more posterior samples used during each iteration, the less this will occur. Once this has run long enough and the maximum likelihood estimate of λ_0 is found, inference can proceed by running the same Gibbs sampler with the selected λ_0 . A derivation of this quantity can be found in the supplementary materials.

3.2 Posterior mode estimation

The most natural implementation of the above formulation is within the Bayesian paradigm, where we can obtain samples of γ directly. This is advantageous as we can examine $p(\gamma|\mathbf{D})$, which provides an assessment of model uncertainty and can be used to identify the best-fitting models. In some situations, however, MCMC can become burdensome if p is very large. An alternative approach is to formulate model estimation as

a penalized likelihood problem, in which we estimate the posterior mode of the model. While in this paradigm we lose some of the aforementioned features of Bayesian inference, estimation can be done in a fraction of the time. Furthermore, the posterior mode of our model will be sparse, i.e. many of the regression coefficients will be estimated to be exactly zero allowing us to quickly perform confounder selection in high-dimensions. The details on the penalized likelihood implementation are in the supplementary materials, where we also show that the posterior mode of Δ from model (1) is consistent at a rate equal to $O(\sqrt{\frac{\log p}{n}})$.

4 Simulation study

In this section we compared our proposed approach with several state of the art alternatives for confounding adjustment in the context of $p \geq n$. We consider three data generating mechanisms: 1) homogeneous treatment effect and sparsity; 2) heterogeneous treatment effect and sparsity; and 3) homogeneous treatment effect and non sparsity. Our goal is always to estimate the average treatment effect. Before we detail the data generating mechanisms to be examined, we describe the approaches being compared.

1. Proposed approach using MCMC, where we estimate λ_0 using the empirical Bayes approach described in Section 3.1 and choose δ as described in Section 2.4 (we will refer to this as EM-SSL)
2. EM-SSL for heterogeneous treatment effect
3. Outcome lasso that includes treatment and covariates, but only places an l_1 penalty on the covariates
4. Re-fit an unpenalized regression model using the covariates identified by the outcome lasso approach above (Post selection lasso)
5. Double post selection approach of Belloni et al. (2014)
6. Doubly robust lasso approach of Farrell (2015)
7. Approximate residual de-biasing approach of Athey et al. (2016)

The purpose of this simulation study is to assess the performance of our proposed approach compared to competitors when the true outcome model is linear. In the context of non linear relationships between the covariates and T or Y none of the methods compared here would perform well.

For some of these estimators, an extension to heterogeneous treatment effects exist, while others implicitly account for treatment effect heterogeneity. With the exception of our estimator, we will always use the version of the estimator that matches the data generating mechanism, e.g. the homogeneous version of the estimators will be used for the homogeneous simulation studies. In all simulations the covariates are drawn from a multivariate normal distribution with marginal variances set to 1 and correlation of

0.6 between all covariates. For each scenario, we compare average percent bias, mean squared error (MSE), 95% interval coverage, and the ratio of the average estimated standard errors and the true standard errors. For our approach, the interval coverage is calculated as the percentage of the time our posterior credible interval covers the true parameter, while all other approaches use frequentist confidence intervals. Finally, we present additional simulation results for differing sample sizes and differing confounding strengths in the supplementary materials, though we found the results to be very similar to those seen here.

4.1 Homogeneous treatment effects

We now examine our approach in a high-dimensional setting where $p = 500$ and $n = 200$, in which there exist strong confounders, weak confounders, and instruments. We simulate the treatment and outcome from the following models:

$$Y_i = T_i + \mathbf{X}_i\boldsymbol{\beta} + \epsilon_i,$$

$$\text{logit}(p(T_i = 1)) = \mathbf{X}_i\boldsymbol{\psi},$$

where $\epsilon_i \sim \mathcal{N}(0, 1)$. The first 8 elements of $\boldsymbol{\beta}$ are $(1, -1, 0.3, -0.3, 0, 0, 1, -1)$, while the remaining elements are drawn from a normal distribution with a standard deviation of 0.1. The first 6 elements of $\boldsymbol{\psi}$ are $(1, -1, 1, -1, 1, -1)$ and the remaining values are set to zero. This leads to a treatment prevalence of 50%. In this setting, covariates 1 and 2 are strong confounders, covariates 3 and 4 are so called “weak” confounders that are weakly associated with the outcome and strongly associated with treatment, covariates 5 and 6 are instruments, and covariates 7 and 8 are strong predictors of the outcome. The remaining covariates have no association with the treatment and a small to moderate association with the outcome. This situation is not strictly sparse due to the small signals in $\boldsymbol{\beta}$, however, it is approximately sparse in the sense that only a small number of covariates are needed to obtain unbiased estimates of the treatment effect.

type	% Bias	MSE	95% interval coverage	$E(\widehat{SE}(\hat{\beta}_t))/SD(\hat{\beta}_t)$
Outcome lasso	49.1	0.34		
Post selection lasso	27.8	0.18		
Double post selection	16.8	0.15	0.81	0.78
Approximate residual de-biasing	43.2	0.28	0.73	1.03
Doubly robust lasso	18.7	0.26	0.72	0.64
EM-SSL	12.9	0.10	0.93	1.01
EM-SSL Heterogeneous	22.5	0.15	0.88	1.04

Table 1: Results for estimating the average treatment effect under the simulation scenario of Section 4.1.

Table 1 shows the results of the proposed simulation across 1000 simulated datasets, and we see that the proposed approach performs the best with respect to all metrics. EM-SSL achieves the minimum bias of 12.9% and the minimum MSE of 0.10. The heterogeneous version of the EM-SSL procedure performs slightly worse in terms of bias and efficiency, which is to be expected given that it splits the sample into the treated

and controls and estimates models separately in the two groups. The next best performing estimator in terms of bias and MSE was the double post selection procedure that had a bias of 16.8% and an MSE of 0.15. The doubly robust lasso had a small bias of 18.7%, though it was quite variable due to the instability of weights in high-dimensions. In terms of interval estimation we do the best in terms of 95% interval coverages (93%) whereas all the other estimators have coverages well below the nominal level (81%, 73%, and 72%). Looking at the ratio of the average estimated to true standard errors, our approach does well (1.01), while most procedures were substantially smaller than 1. The approximate residual de-biasing procedure does well at estimating the standard errors, but is too biased to achieve good interval coverages. Our EM-SSL Heterogeneous procedure also does well at estimating the standard errors, but has an interval coverage of 88% due to the larger amount of bias relative to the homogeneous version.

Figure 2 shows the posterior inclusion probabilities for the homogeneous EM-SSL model for each of the different types of covariates in the model. We see that the variables strongly associated with both the treatment and outcome (X_1 and X_2) have the highest value of $P(\gamma_j = 1|\mathbf{D})$. Variables X_3 and X_4 have weak associations with the outcome, but strong relationships with the treatment and they enter into the slab the next highest percentage of the time. Due to our weights, strong instrumental variables (X_5 and X_6) are in the slab approximately 20% of the time. Strong predictors of the outcome enter the slab slightly more often than instruments, while the remaining variables almost never enter the slab. These posterior inclusion probabilities highlight why there exists bias in our estimates of the treatment effect. The important confounders are included in the spike component of the prior during some MCMC scans leading to more shrinkage of important components of β and biased estimates of the ATE. It is important to note that even when a coefficient is included in the spike, it is not eliminated from the model completely, but rather is more aggressively shrunk to zero. This small amount of bias seems to come with improved efficiency, however, as our estimator has the smallest MSE overall.




Figure 2: Posterior inclusion probabilities from the homogeneous model for simulations in Sections 4.1 and 4.2.

4.2 Heterogeneous treatment effects

We now simulate data with $n = 400$ and $p = 800$ and a heterogeneous treatment effect. The data generating models are of the following form:

$$\begin{aligned} Y_i &= T_i + 0.6(T_i \times X_{1i}) + 0.4(T_i \times X_{3i}) + X_{1i} - X_{2i} + \\ &\quad 0.3X_{3i} - 0.3X_{4i} + X_{7i} - X_{8i} + \epsilon_i, \\ \text{logit}(p(T_i = 1)) &= -1.5 + X_{1i} - X_{2i} + X_{3i} - X_{4i} + X_{5i} - X_{6i}, \end{aligned}$$

where $\epsilon_i \sim \mathcal{N}(0, 1)$. This simulation is similar to the previous section with three changes: 1) covariates 9 through 500 have no association with either treatment or outcome; 2) there is an interaction between the treatment and covariates 1 and 3; and 3) the prevalence of the treatment has been dropped from approximately 50% to 25%. We have increased the sample size from 200 to 400 since we lowered the prevalence of the treatment, and all methods explored need a sufficient sample size in both the treated and control groups to estimate heterogeneous treatment effects. Table 2 shows the results averaged across 1000 simulations. Results are similar to the homogeneous treatment effects setting. The double post selection approach again fares relatively well across all metrics, however, is outperformed by the EM-SSL approach. The approximate residual de-biasing approach has fairly substantial amounts of bias in this setting, which also leads to poor interval coverages. The doubly robust lasso again has a higher MSE due to the instability of inverse propensity weights in high-dimensional settings. Our EM-SSL procedure does not achieve the nominal interval coverage in this setting, though this is due to the bias incurred by assuming a homogeneous treatment effect. Our EM-SSL heterogeneous procedure achieves interval coverages of 95% and an average estimated standard error that is close to the truth (0.99). Figure 2 shows the posterior inclusion probabilities for the homogeneous EM-SSL model, and we see that all the important confounders and predictors are included nearly 100% of the time into the slab component of the prior. This shows that the increased sample size in this simulation, compared with the previous simulation, leads to the improved variable selection. Further, it shows that the bias we see in the EM-SSL estimator is not caused by shrinkage of important parameters, but rather because it assumes homogeneity when the treatment effect is truly heterogeneous.

type	% Bias	MSE	95% interval coverage	$E(\widehat{SE}(\widehat{\beta}_t))/SD(\widehat{\beta}_t)$
Outcome LASSO	52.0	0.30		
Post LASSO	25.2	0.09		
Double post selection	9.0	0.10	0.85	0.67
Approximate residual de-biasing	35.0	0.16	0.45	0.93
Doubly robust LASSO	18.0	0.14	0.73	0.69
EM-SSL	17.0	0.05	0.76	0.95
EM-SSL Heterogeneous	5.0	0.04	0.95	0.99

Table 2: Results for estimating the average treatment effect under the simulation scenario of Section 4.2.

4.3 Dense treatment model

Here we simulate data as described in Athey et al. (2016) where the treatment model is purposely chosen to be dense. More specifically, first we define 20 clusters, $\{\mathbf{c}_1, \dots, \mathbf{c}_{20}\}$ where $\mathbf{c}_k \sim \mathcal{N}(0, I_{p \times p})$. Second, we draw \mathbf{C}_i uniformly at random from one of the 20. Third, we draw the covariates from a multivariate normal distribution centered at \mathbf{C}_i with the identity matrix as the covariance. Fourth, we set $T_i = 1$ with probability 0.1 for the first 10 clusters, and $T_i = 1$ with probability 0.9 for the remaining clusters. Finally, we generate data from the outcome model defined as $Y_i = 10T_i + \mathbf{X}\beta + \epsilon_i$, where $\beta \propto (1, \frac{1}{\sqrt{2}}, \dots, \frac{1}{\sqrt{p}})$ and is normalized such that $\|\beta\|_2^2 = 18$. Here we will again set $n = 200$ and $p = 500$. Intuitively, this is a simulation scenario in which the outcome model is approximately sparse, though the treatment model is dense as all of the covariates are associated with the treatment.

Results are summarized in Table 3. Because the data generating mechanism does not assume sparsity, our original EM-SSL procedure performs poorly relative to the post selection lasso approach. We obtain an MSE of 1.07, while also doing very poorly at estimating the standard errors of our approach as the ratio of the average estimated to true standard errors is 0.59. However, under a non sparse setting, if we impose a restriction that only the top $k = 10$ variables most associated with the treatment (identified by the magnitude of their coefficients in the treatment lasso model) are prioritized with $w_j = \delta$, then our approach (EM-SSL Restricted) performs the best in terms of MSE (0.59) and interval coverage (93%). It is also important to note that while we did not show the restricted results in the other simulation scenarios that had sparse treatment models, the restricted approach performed almost identically to the original EM-SSL approach. While there is no principled way of selecting k , we have found that other values of k , such as $k = 20$, perform similarly well.

type	% Bias	MSE	95% interval coverage	$E(\widehat{SE}(\hat{\beta}_t))/SD(\hat{\beta}_t)$
Outcome LASSO	0.0	0.88		
Post LASSO	0.0	0.76		
Double post selection	0.0	1.06	0.82	0.69
Approximate residual de-biasing	0.0	1.22	0.81	0.69
Doubly robust lasso	0.0	1.85	0.49	0.40
EM-SSL	0.0	1.07	0.74	0.59
EM-SSL Heterogeneous	0.0	1.64	0.88	0.79
EM-SSL Restricted	0.0	0.59	0.93	0.93
EM-SSL Restricted Heterogeneous	0.0	1.11	0.93	0.97

Table 3: Results for estimating the average treatment effect under the simulation scenario of Section 4.3.

4.4 Choosing between homogeneous and heterogeneous models

An important question is how to decide between the homogeneous and heterogeneous versions of our model in practice. One potential solution to this is to use the Watanabe-Akaike information criterion (WAIC) (Watanabe, 2010; Gelman et al., 2014), which is a Bayesian analog to traditional model selection tools. We applied WAIC to each of the three simulation scenarios described above to evaluate its effectiveness in choosing

the right model. The correct model in Sections 4.1 and 4.3 is the homogeneous model, and the correct model in Section 4.2 is the heterogeneous one. In the three simulation scenarios, the WAIC chose the correct model 93%, 82%, and 99% of the time, respectively. This shows that it is possible to automate the decision between the homogeneous and heterogeneous versions of the model, leading to reductions in bias and MSE. These results should be taken with caution, however, as credible intervals from a model chosen using WAIC will not account for the additional uncertainty incurred from the model selection process. This didn't seem to impact our results greatly as our credible interval coverages were 93%, 92%, and 94% for the three simulations explored when using the model chosen by WAIC.

5 Analysis of NHANES data

Recent work (Wild, 2005; Patel et al., 2010; Louis et al., 2012; Patel et al., 2012; Patel and Ioannidis, 2014) has centered on studying the effects of a vast set of exposures on disease. These analyses, termed environmental wide association studies (EWAS), examine environmental factors and aim to improve understanding of the long term effects of different exposures and toxins that humans are invariably exposed to on a daily basis. The National Health and Nutrition Examination Survey (NHANES), is a cross-sectional data source made publicly available by the Centers for Disease Control and Prevention (CDC). The data has also been aggregated and made available by Patel et al. (2016). The NHANES data is a nationally representative study, in which participants were questioned regarding their health status, with a subset of these patients providing extensive clinical and laboratory tests to provide information on a variety of environmental attributes such as chemical toxicants, pollutants, allergens, bacterial/viral organisms, and nutrients (Patel et al., 2010).

Our analysis will center on data from the 1999–2000, 2001–2002, 2003–2004, and 2005–2006 surveys. We build on the analysis described in Patel et al. (2012), by applying our proposed methodology to estimating the effects of volatile compounds (VCs) on triglyceride levels in humans. VCs were measured in $n = 177$ subjects, and there were $p = 127$ covariates. The list of potential confounders consists of other volatile compounds, their interactions, other persistent pesticides measured in the respective subsample, body measurements, demographic, and socioeconomic variables. To evaluate the proposed approach in a setting with $p > n$, we ran an additional analysis, which looked at the effect of VCs on triglycerides in subjects over 40 years old. This led to a sample size of $n = 77$ subjects.

In previous work (Patel et al., 2012) these exposures were evaluated individually without controlling for the remaining pesticides, and only a small subset of pre-selected covariates such as age, body mass index (BMI), and gender were controlled for. Patel and Ioannidis (2014) wrote that persistent pesticides tend to be highly correlated with one another and that many of the associations found by previous exposome studies (Patel et al., 2012) could simply be to confounding bias that was unadjusted for due to the small set of confounders used and the fact that other pesticides were not adjusted for. This highlights the need for an analysis that adjusts for all potential confounders, but in our analyses we have $p = 127$ covariates with small sample sizes. Therefore, due

to the large model space, some confounder selection or shrinkage is required to obtain efficient estimates of exposure effects.

5.1 VC analyses

We now analyze the NHANES data described above. In particular, we will examine the effect of each of 10 volatile compounds on triglyceride levels while controlling for an extensive set of potential confounders. We analyze the effect of each volatile compound using three approaches: 1) an unadjusted model that regresses the outcome on the treatment without confounder adjustment, 2) the homogeneous EM-SSL procedure described above, and 3) the double post selection approach described in Belloni et al. (2014). The approximate residual de-biasing and the doubly robust lasso approaches are only applicable to categorical treatments and are therefore left out. We restrict attention to the homogeneous treatment effect application of our approach, because addressing heterogeneity in the setting of continuous treatments would require additional work that is beyond the scope of this paper.

Figure 3 shows the point estimates and 95% confidence intervals (credible interval for EM-SSL approach) from the analysis across the 10 volatile compounds for the three approaches under consideration and each of the two data sets being analyzed. The results are qualitatively very similar across the three approaches for each of the 10 exposures we looked at, in both the full data set and the data set restricting to older subjects. One exception is VC7 in the analysis of older subjects where the EM-SSL estimate has a credible interval that does not contain zero, while the confidence intervals for the other two approaches do. In general, however, the results are very similar in magnitude and direction across the approaches, with the only major difference coming in the widths of the corresponding confidence (and credible) intervals.

5.2 Comparison of standard errors across approaches

While there are not drastic differences in point estimates, there are large differences in the widths of the confidence intervals of the two approaches that aim to adjust for confounding. Of interest is the ratio of the standard errors for the EM-SSL procedure and the double post selection approach. These can be seen in Figure 4. We see that in the full data set the EM-SSL procedure is more efficient overall than the double post selection approach. The majority of the analyses (8/10) had smaller confidence intervals under the EM-SSL procedure with an average confidence interval ratio of 0.9 as indicated by the dashed line in Figure 4. This means that the EM-SSL procedure on average has a 10% smaller standard error than the double post selection approach and occasionally has a standard error 30% smaller. The results are even more striking when we subset the data to the $n = 77$ subjects who are over 40 years old. In this case all analyses were more efficient using the EM-SSL procedure, with an average standard error ratio of 0.78, highlighting the ability of our estimator in high-dimensional scenarios. That our estimator is more efficient than the double post selection estimator is not surprising. The goal of the double post selection estimator was to obtain valid inference in high-dimensional scenarios, not to provide the most efficient estimate of the treatment effect.

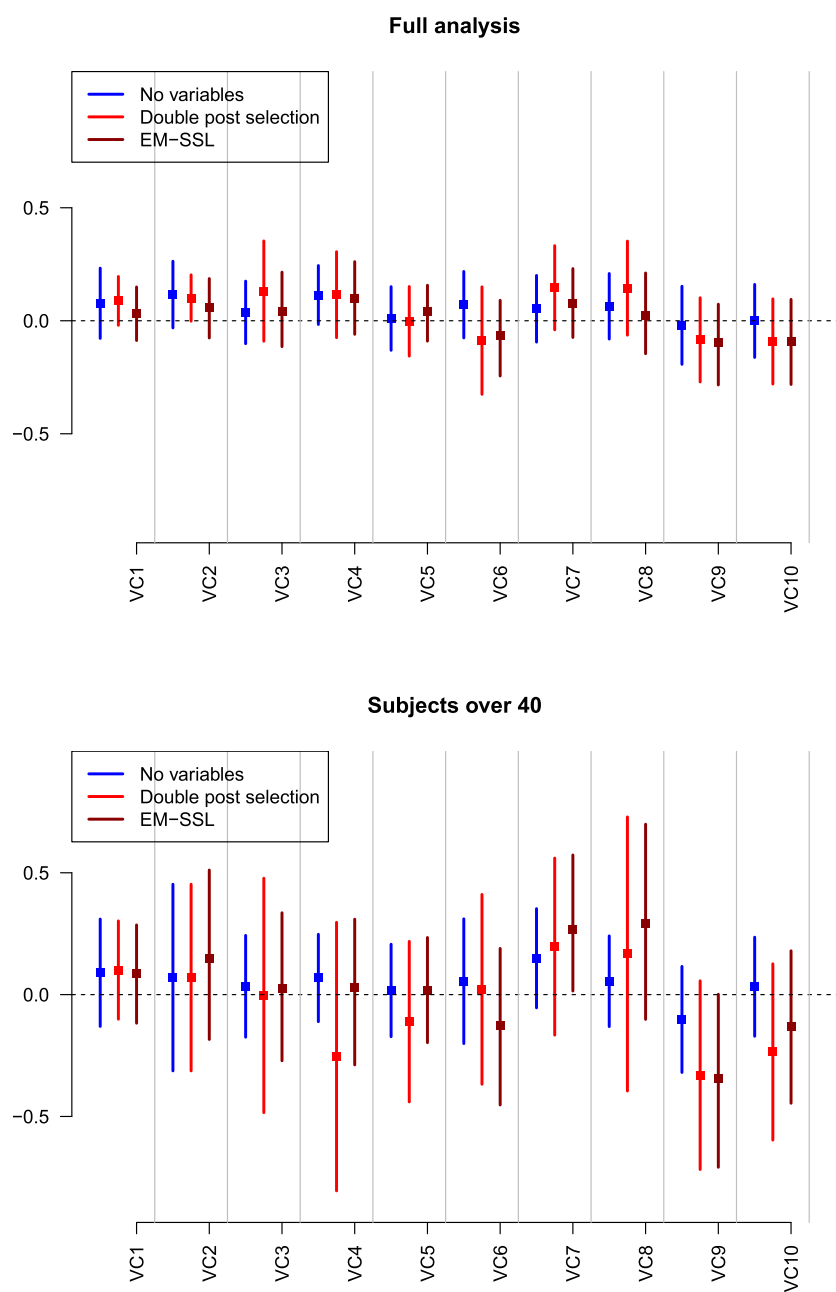


Figure 3: Results from analysis of volatile compounds on triglycerides. The upper panel shows the results using the full, $n = 177$, sample. The lower panel shows the results for just the $n = 77$ subjects who are over 40 years old.

Nonetheless, this analysis highlights an important difference between the approaches in their finite sample performance and how they address instrumental variables.




Figure 4: The left panel shows a histogram of the ratios of standard errors for the EM-SSL approach and the double post selection approach for the analysis of volatile compounds in the full data. The right panel shows the corresponding histogram for the analysis of subjects over the age of 40. The dashed vertical line is the mean of the ratios that make up the histogram.

6 Discussion

In this paper we have introduced a novel approach for estimating treatment effects in high-dimensional settings. We introduced a generalization of the spike and slab formulation to allow the prior probability that a parameter for a given covariate is included in the slab component of the prior to depend on the association between each potential confounder and the treatment. We highlighted how this could drastically reduce the shrinkage of important confounders, while still shrinking to zero the coefficients of instrumental and noise variables. Through simulation we showed that our proposed approach has better performance than state of the art approaches under data generating mechanisms that are more or less sparse and also in the context of heterogeneous treatment effects. By tackling the problem within the Bayesian paradigm we achieve good interval coverage rates even in small samples unlike existing approaches in the literature. Importantly, we applied the proposed approach to an exposome study and found that our approach gave smaller confidence intervals than existing approaches for confounding adjustment in high dimensional settings.

Our prior is purposely constructed to improve small sample performance of ATE estimation. It shares some commonalities with doubly robust approaches that aim to

model both the treatment and outcome to reduce bias. A crucial difference, however, is that we try to eliminate variables only associated with the treatment while still prioritizing the inclusion of potential confounders in the outcome model, to minimize both bias and variance in small and finite samples. This differs from existing approaches (Belloni et al., 2014), which aims at eliminating confounding bias by including all variables associated with either the treatment or outcome. Further, a nice feature of Bayesian approaches is that they account for all the sources of uncertainty in the estimation of the ATE, thus performing better in finite samples than asymptotic approaches. This is because asymptotic approaches often assume that higher-order terms from asymptotic expansions are asymptotically negligible. Such assumptions do not hold in finite samples. Bayesian approaches do not rely on asymptotic expansions nor on the assumption of negligibility detailed above. Statistical uncertainty associated with all the model parameters is indeed accounted for in the credible intervals for the treatment effect, leading to improved finite sample coverage as seen in Section 4. As the sample size increases, the differences between our Bayesian approach and approaches based on asymptotic approximations will diminish.

Our proposed approach has limitations. First, we make the strong assumption of a linear outcome model which allows us to: 1) handle ultra high-dimensional covariate spaces; and 2) borrow information from the treatment model when estimating the causal effects. While similar assumptions of linearity are also used in existing approaches to high-dimensional confounding adjustment (Belloni et al., 2014; Farrell, 2015; Athey et al., 2016; Antonelli et al., 2018; Shortreed and Ertefaie, 2017), a topic of future research would be to extend these ideas to nonlinear settings to overcome challenges inherent to model misspecification. Second, we also make the assumption of sparsity of both the treatment and outcome models. While we showed that we can potentially overcome a lack of sparsity in the treatment model by restricting that only a small percentage of the total number of covariates be prioritized in the outcome model, our approach still relies on sparsity of the outcome model to obtain good results in terms of estimation and interval coverages. Third, although we consider scenarios of large p , MCMC can become computationally intensive in ultra-high dimensions where the number of covariates is in the tens of thousands. In this setting, alternative approaches such as double post selection, the doubly robust lasso, and approximate residual debiasing could be used. Lastly, a common criticism of confounder selection is that there is a nonzero probability of excluding a confounder, which leads to bias in the estimation of the causal effect. While excluding a confounder is certainly an issue, we have constructed our prior to avoid this problem as much as possible, by only excluding confounders when this will have very low probability of contributing to bias. In small samples a small amount of bias can be acceptable in high-dimensional scenarios to improve efficiency.

Our prior construction involves building a lasso model as a pre-processing step to build an informative prior. From a Bayesian regression modeling perspective, using the covariates in a regression model to inform the prior distribution of the regression parameters is widely accepted whenever the covariates can be treated as fixed. In standard regression models with covariates \mathbf{X} , Zellner's g-prior is frequently used, which sets the prior on the variance proportional to $(\mathbf{X}^T \mathbf{X})^{-1}$. In our approach, the outcome model is treating both T and \mathbf{X} as fixed. Therefore constructing a prior from the association

between T and \mathbf{X} is consistent with standard practice, as long as we don't use the outcome Y to inform the prior. From a causal inference perspective, a prevalent perspective is to separate the design and analysis stages of causal inference (Rubin et al., 2008). The design phase consists of any steps that occur before analyzing the outcome and can include propensity score modeling, building of matched sets, design of the study, etc. In this perspective, uncertainty in the design phase is typically not accounted for in the analysis phase. Our model follows this perspective, as our prior construction relies on only the propensity score model. Alternative perspectives are now emerging Liao and Zigler (2018), and extensions of our approach would be worthy of consideration for future work.

There are a number of extensions of the proposed ideas that merit further research. In the current manuscript we restricted attention to the case where $w_j = \delta$ for all variables associated with the treatment. We can relax this assumption to let w_j vary for each covariate j associated with the treatment, potentially improving on the current approach, though future research would be required to find an optimal strategy. The ideas in this paper could also be used to improve finite sample estimation for doubly robust estimators. Doubly robust estimators typically combine an outcome regression and a treatment model, and our ideas could be used to improve the outcome model in this estimator. This, coupled with improved estimation of the treatment model using similar ideas as done in Shortreed and Ertefaie (2017), could lead to improved doubly robust estimators. Finally, the idea to borrow information from the treatment model to guide the amount of shrinkage in the outcome model can be extended to other high-dimensional priors beyond the spike and slab one seen here. A number of priors are used in high-dimensional Bayesian modeling, and this idea can potentially be extended to many of them.

Supplementary Material

Supplementary materials for “High-dimensional confounding adjustment using continuous spike and slab priors” (DOI: [10.1214/18-BA1131SUPP](https://doi.org/10.1214/18-BA1131SUPP); .pdf). Here we give additional details and derivations for estimation of the empirical Bayes variance and posterior calculation. We further illustrate the estimation of the posterior mode of our model and give additional simulation results. An R package implementing the approach for both binary and continuous outcomes is available at github.com/jantonelli1111/HDconfounding.

References

- Albert, J. H. and Chib, S. (1993). “Bayesian analysis of binary and polychotomous response data.” *Journal of the American statistical Association*, 88(422): 669–679. [MR1224394](#). [808](#)
- Antonelli, J., Cefalu, M., Palmer, N., and Agniel, D. (2018). “Doubly robust matching estimators for high dimensional confounding adjustment.” *Biometrics*. [807](#), [823](#)

- Antonelli, J., Parmigiani, G., and Dominici, F. (2018). “Supplementary materials for “High-dimensional confounding adjustment using continuous spike and slab priors”.” *Bayesian Analysis*. doi: <https://doi.org/10.1214/18-BA1131SUPP>. 809
- Antonelli, J., Zigler, C., and Dominici, F. (2017). “Guided Bayesian imputation to adjust for confounding when combining heterogeneous data sources in comparative effectiveness research.” *Biostatistics*, 18(3): 553–568. MR3799594. doi: <https://doi.org/10.1093/biostatistics/kxx003>. 806
- Athey, S., Imbens, G. W., and Wager, S. (2016). “Approximate residual balancing: De-biased inference of average treatment effects in high dimensions.” *arXiv preprint arXiv:1604.07125*. 807, 814, 818, 823
- Belloni, A., Chernozhukov, V., Fernández-Val, I., and Hansen, C. (2017). “Program Evaluation and Causal Inference With High-Dimensional Data.” *Econometrica*, 85(1): 233–298. MR3611771. doi: <https://doi.org/10.3982/ECTA12723>. 806
- Belloni, A., Chernozhukov, V., and Hansen, C. (2014). “Inference on treatment effects after selection among high-dimensional controls.” *The Review of Economic Studies*, 81(2): 608–650. MR3207983. doi: <https://doi.org/10.1093/restud/rdt044>. 806, 807, 814, 820, 823
- Bhattacharya, A., Pati, D., Pillai, N. S., and Dunson, D. B. (2015). “Dirichlet–Laplace priors for optimal shrinkage.” *Journal of the American Statistical Association*, 110(512): 1479–1490. MR3449048. doi: <https://doi.org/10.1080/01621459.2014.960967>. 806
- Carvalho, C. M., Polson, N. G., and Scott, J. G. (2010). “The horseshoe estimator for sparse signals.” *Biometrika*, asq017. MR2650751. doi: <https://doi.org/10.1093/biomet/asq017>. 806
- Casella, G. (2001). “Empirical Bayes Gibbs sampling.” *Biostatistics*, 2(4): 485–500. 813
- Cefalu, M., Dominici, F., Arvold, N., and Parmigiani, G. (2016). “Model averaged double robust estimation.” *Biometrics*. MR3665958. doi: <https://doi.org/10.1111/biom.12622>. 806
- Crainiceanu, C. M., Dominici, F., and Parmigiani, G. (2008). “Adjustment uncertainty in effect estimation.” *Biometrika*, 95(3): 635–651. MR2443180. doi: <https://doi.org/10.1093/biomet/asn015>. 806
- De Luna, X., Waernbaum, I., and Richardson, T. S. (2011). “Covariate selection for the nonparametric estimation of an average treatment effect.” *Biometrika*, asr041. MR2860329. doi: <https://doi.org/10.1093/biomet/asr041>. 806
- Ertefaie, A., Asgharian, M., and Stephens, D. (2015). “Variable selection in causal inference using a simultaneous penalization method.” *arXiv preprint arXiv:1511.08501*. 807
- Fan, J. and Li, R. (2001). “Variable selection via nonconcave penalized likelihood and

- its oracle properties.” *Journal of the American statistical Association*, 96(456): 1348–1360. MR1946581. doi: <https://doi.org/10.1198/016214501753382273>. 806
- Farrell, M. H. (2015). “Robust inference on average treatment effects with possibly more covariates than observations.” *Journal of Econometrics*, 189(1): 1–23. MR3397349. doi: <https://doi.org/10.1016/j.jeconom.2015.06.017>. 807, 814, 823
- Gelman, A., Carlin, J. B., Stern, H. S., and Rubin, D. B. (2014). *Bayesian data analysis*, volume 2. Chapman & Hall/CRC Boca Raton, FL, USA. MR2027492. 818
- George, E. I. and McCulloch, R. E. (1993). “Variable selection via Gibbs sampling.” *Journal of the American Statistical Association*, 88(423): 881–889. 806
- Hahn, P. R., Carvalho, C., and Puelz, D. (2016). “Bayesian Regularized Regression for Treatment Effect Estimation from Observational Data.” *Available at SSRN*. 807
- Hahn, P. R., Murray, J. S., and Carvalho, C. (2017). “Bayesian regression tree models for causal inference: regularization, confounding, and heterogeneous effects.” *arXiv preprint arXiv:1706.09523*. 806, 812
- Liao, S. and Zigler, C. (2018). “Uncertainty in the Design Stage of Two-Stage Bayesian Propensity Score Analysis.” *arXiv preprint arXiv:1809.05038*. 824
- Little, R. J. and Rubin, D. B. (2000). “Causal effects in clinical and epidemiological studies via potential outcomes: concepts and analytical approaches.” *Annual review of public health*, 21(1): 121–145. 808
- Lockhart, R., Taylor, J., Tibshirani, R. J., and Tibshirani, R. (2014). “A significance test for the lasso.” *Annals of statistics*, 42(2): 413. MR3210970. doi: <https://doi.org/10.1214/13-AOS1175>. 806
- Louis, B., Germaine, M., and Sundaram, R. (2012). “Exposome: time for transformative research.” *Statistics in medicine*, 31(22): 2569–2575. MR2972269. doi: <https://doi.org/10.1002/sim.5496>. 819
- Park, T. and Casella, G. (2008). “The bayesian lasso.” *Journal of the American Statistical Association*, 103(482): 681–686. MR2524001. doi: <https://doi.org/10.1198/016214508000000337>. 806, 808, 813
- Patel, C. J., Bhattacharya, J., and Butte, A. J. (2010). “An environment-wide association study (EWAS) on type 2 diabetes mellitus.” *PloS one*, 5(5): e10746. 819
- Patel, C. J., Cullen, M. R., Ioannidis, J. P., and Butte, A. J. (2012). “Systematic evaluation of environmental factors: persistent pollutants and nutrients correlated with serum lipid levels.” *International journal of epidemiology*, 41(3): 828–843. 819
- Patel, C. J. and Ioannidis, J. P. (2014). “Studying the elusive environment in large scale.” *Jama*, 311(21): 2173–2174. 819
- Patel, C. J., Pho, N., McDuffie, M., Easton-Marks, J., Kothari, C., Kohane, I. S., and Avillach, P. (2016). “A database of human exposomes and phenomes from the US National Health and Nutrition Examination Survey.” *Scientific data*, 3. 819

- Pearl, J. (2011). “Invited commentary: understanding bias amplification.” *American journal of epidemiology*, 174(11): 1223–1227. 811
- Ročková, V. and George, E. I. (2016). “The spike-and-slab lasso.” *Journal of the American Statistical Association*, (just-accepted). MR3803476. doi: <https://doi.org/10.1080/01621459.2016.1260469>. 806, 809
- Rosenbaum, P. R. and Rubin, D. B. (1983). “The central role of the propensity score in observational studies for causal effects.” *Biometrika*, 70(1): 41–55. MR0742974. doi: <https://doi.org/10.1093/biomet/70.1.41>. 805
- Rubin, D. B. et al. (1981). “The bayesian bootstrap.” *The annals of statistics*, 9(1): 130–134. MR0600538. 812
- Rubin, D. B. et al. (2008). “For objective causal inference, design trumps analysis.” *The Annals of Applied Statistics*, 2(3): 808–840. 824
- Scott, J. G., Berger, J. O., et al. (2010). “Bayes and empirical-Bayes multiplicity adjustment in the variable-selection problem.” *The Annals of Statistics*, 38(5): 2587–2619. MR2722450. doi: <https://doi.org/10.1214/10-AOS792>. 809
- Shortreed, S. M. and Ertefaie, A. (2017). “Outcome-adaptive lasso: Variable selection for causal inference.” *Biometrics*. MR3744525. doi: <https://doi.org/10.1111/biom.12679>. 807, 823, 824
- Talbot, D., Lefebvre, G., and Atherton, J. (2015). “The Bayesian causal effect estimation algorithm.” *Journal of Causal Inference*, 3(2): 207–236. 806
- Taylor, J. and Tibshirani, R. (2016). “Post-selection inference for l1-penalized likelihood models.” *arXiv preprint arXiv:1602.07358*. MR3767165. doi: <https://doi.org/10.1002/cjs.11313>. 806
- Tibshirani, R. (1996). “Regression shrinkage and selection via the lasso.” *Journal of the Royal Statistical Society. Series B (Methodological)*, 267–288. MR1379242. 806, 811
- van der Laan, M. J. and Gruber, S. (2010). “Collaborative double robust targeted maximum likelihood estimation.” *The international journal of biostatistics*, 6(1). MR2653848. doi: <https://doi.org/10.2202/1557-4679.1181>. 806
- Vansteelandt, S., Bekaert, M., and Claeskens, G. (2012). “On model selection and model misspecification in causal inference.” *Statistical methods in medical research*, 21(1): 7–30. MR2867536. doi: <https://doi.org/10.1177/0962280210387717>. 806
- Wang, C., Dominici, F., Parmigiani, G., and Zigler, C. M. (2015). “Accounting for uncertainty in confounder and effect modifier selection when estimating average causal effects in generalized linear models.” *Biometrics*, 71(3): 654–665. MR3402601. doi: <https://doi.org/10.1111/biom.12315>. 806
- Wang, C., Parmigiani, G., and Dominici, F. (2012). “Bayesian effect estimation accounting for adjustment uncertainty.” *Biometrics*, 68(3): 661–671. MR3055168. doi: <https://doi.org/10.1111/j.1541-0420.2011.01731.x>. 806

- Watanabe, S. (2010). “Asymptotic equivalence of Bayes cross validation and widely applicable information criterion in singular learning theory.” *Journal of Machine Learning Research*, 11(Dec): 3571–3594. [MR2756194](#). 818
- Wild, C. P. (2005). “Complementing the genome with an “exposome”: the outstanding challenge of environmental exposure measurement in molecular epidemiology.” *Cancer Epidemiology Biomarkers & Prevention*, 14(8): 1847–1850. 819
- Wilson, A. and Reich, B. J. (2014). “Confounder selection via penalized credible regions.” *Biometrics*, 70(4): 852–861. [MR3295746](#). doi: <https://doi.org/10.1111/biom.12203>. 807
- Zhou, J., Bhattacharya, A., Herring, A. H., and Dunson, D. B. (2015). “Bayesian factorizations of big sparse tensors.” *Journal of the American Statistical Association*, 110(512): 1562–1576. URL <http://www.tandfonline.com/doi/abs/10.1080/01621459.2014.983233#.VNQ2p1WUd5k>. [MR3449055](#). doi: <https://doi.org/10.1080/01621459.2014.983233>. 809
- Zigler, C. M. and Dominici, F. (2014). “Uncertainty in propensity score estimation: Bayesian methods for variable selection and model-averaged causal effects.” *Journal of the American Statistical Association*, 109(505): 95–107. [MR3180549](#). doi: <https://doi.org/10.1080/01621459.2013.869498>. 806
- Zou, H. (2006). “The adaptive lasso and its oracle properties.” *Journal of the American statistical association*, 101(476): 1418–1429. [MR2279469](#). doi: <https://doi.org/10.1198/016214506000000735>. 806
- Zou, H. and Hastie, T. (2005). “Regularization and variable selection via the elastic net.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2): 301–320. [MR2137327](#). doi: <https://doi.org/10.1111/j.1467-9868.2005.00503.x>. 806
- Zubizarreta, J. R. (2015). “Stable weights that balance covariates for estimation with incomplete outcome data.” *Journal of the American Statistical Association*, 110(511): 910–922. [MR3420672](#). doi: <https://doi.org/10.1080/01621459.2015.1023805>. 807

Acknowledgments

The authors are grateful for Chirag Patel and his advice regarding the NHANES data analysis. Funding for this work was provided by National Institutes of Health (ES000002, ES024332, ES007142, ES026217, ES028033, P01CA134294, R01GM111339, R35CA197449, P50MD010428, DP2MD012722), The U.S. Environmental Protection Agency (83615601, 83587201-0), and The Health Effects Institute (4953-RFA14-3/16-4).