

ARTICLE

Bayesian Sensitivity Analysis for Unmeasured Confounding in Causal Panel Data Models

Licheng Liu¹  and Teppei Yamamoto²

¹Political Science, University of Michigan, USA; ²Waseda University, Japan

Corresponding author: Licheng Liu; Email: lichengl@umich.edu

(Received 14 April 2025; revised 1 September 2025; accepted 4 September 2025)

Abstract

Despite the recent methodological advancements in causal panel data analysis, concerns remain about unobserved unit-specific time-varying confounders that cannot be addressed by unit or time fixed effects or their interactions. We develop a Bayesian sensitivity analysis (BSA) method to address the concern. Our proposed method is built upon a general framework combining Rubin's Bayesian framework for model-based causal inference (Rubin [1978], *The Annals of Statistics* 6(1), 34–58) with parametric BSA (McCandless, Gustafson, and Levy [2007], *Statistics in Medicine* 26(11), 2331–2347). We assess the sensitivity of the causal effect estimate from a linear factor model to the possible existence of unobserved unit-specific time-varying confounding, using the coefficients of the treatment variable and observed confounders in the model for the unobserved confounding as sensitivity parameters. We utilize priors on these coefficients to constrain the hypothetical severity of unobserved confounding. Our proposed approach allows researchers to benchmark the assumed strength of confounding on observed confounders more systematically than conventional frequentist sensitivity analysis techniques. Moreover, to cope with convergence issues typically encountered in nonidentified Bayesian models, we develop an efficient Markov chain Monte Carlo algorithm exploiting transparent parameterization (Gustafson [2005], *Statistical Science* 20(2), 111–140). We illustrate our proposed method in a Monte Carlo simulation study as well as an empirical example on the effect of war on inheritance tax rates.

Keywords: Bayesian sensitivity analysis; causal inference; linear factor model; hierarchical shrinkage priors; transparent parameterization

Edited by: Daniel J. Hopkins and Brandon M. Stewart

1. Introduction

Researchers are often concerned about unobserved confounding when analyzing time-series cross-section data. To address such confounding, scholars routinely use linear outcome models with additive two-way fixed effects (TWFEs) to take advantage of the panel data structure. However, the underlying assumptions for identifying causal effects with a linear model with additive fixed effects, often heuristically understood as “parallel trends” (PTs) of the outcomes in the absence of treatment, are known to be rather restrictive (e.g., Manski and Pepper 2018).

Recent model-based approaches relax the assumption of PT by replacing the TWFE with an interactive factor term (e.g., Abadie, Diamond, and Hainmueller 2010, 2015). The factor term, also known as interactive fixed effects (IFEs; Bai 2009), allows the common shocks to have heterogeneous influence on different units and incorporates the TWFE as a special case. In addition, the counterfactual-prediction-based estimators, like the synthetic control (SC) estimator, allow the treatment effect to be heterogeneous. Such causal panel data models (Athey *et al.* 2021) are gaining popularity in empirical

social sciences research. Still, there may be an unobserved confounder U that varies both across units and time periods in a way that cannot be fully captured by either the TWFE or the IFE.

To address this problem, we develop a fully Bayesian sensitivity analysis (BSA) method. Our method is based on a general framework for causal BSA which combines Rubin's Bayesian framework for model-based causal inference (Rubin 1978) with parametric BSA (McCandless, Gustafson, and Levy 2007). That is, we adopt a strategy similar to McCandless *et al.* (2007) by assuming that there is an unobserved continuous confounder U and investigate how its existence affect the treatment effect estimate under a given prior belief about the magnitude of confounding by U .

The proposed methodology fits within the literature that assesses the robustness of treatment effect estimation in the presence of potential violations of PT. Rambachan and Roth (2023) propose a sensitivity analysis framework that treats observed pre-trends as informative about possible violations of PT. Ye *et al.* (2024) assume that the outcome trajectories of treated units are bounded by those of the set of control units. Both approaches yield partial identification of causal effects. The proposed BSA framework assesses robustness by explicitly modeling the presence of an unobserved time-varying confounder. Built on a counterfactual prediction approach that accommodates heterogeneous treatment effects, it can be viewed as an extension of the omitted variable bias (OVB) framework in Cinelli and Hazlett (2020) to panel data settings. Recent developments in more interpretable sensitivity analysis (e.g., Cinelli and Hazlett 2020; VanderWeele and Ding 2017) may be adapted to causal panel settings, although doing so often involves additional structural assumptions, such as treatment effect homogeneity.

We make two important contributions in this article. First, conceptually, our proposed method addresses a key challenge that is common to most existing non-BSA techniques for unobserved confounding: the specification of “plausible” sensitivity parameter values. While there have been many successful attempts to address this in a frequentist setting—for instance, researchers can assess how large an influence unobserved confounding would need to have to explain away the estimated treatment effects (e.g., Altonji, Elder, and Taber 2005; Imbens 2003)—existing techniques typically leave such specification up to the discretion of the researcher, providing little systematic guidance as to how to set those parameter values. BSA, on the other hand, provides a principled solution to this problem via a formally justified procedure to incorporate the researcher's prior belief about those values. That is, in BSA, a researcher would begin with a prior distribution on the sensitivity parameter values and the results would be summarized using the posterior distribution of the sensitivity parameter, e.g., by averaging the causal effect estimates over the posterior on the sensitivity parameter.

In many BSA methods, however, this advantage remains only theoretical, as they also provide little guidance on how to specify such priors. In contrast, our proposed approach adopts the approach of Gustafson *et al.* (2010) and borrows information from the observed covariates to make an informed choice about the sensitivity parameter prior. Specifically, we assign priors to a subset of sensitivity parameters with restrictions to guarantee that the dispersion of the unobserved confounder is comparable to that of observed confounders. Then, we assign exchangeable priors to the coefficients of both unobserved and observed confounders. We view this as a Bayesian analog of “benchmarking” sensitivity parameters by observed pre-treatment confounders, as commonly done in frequentist approaches to make results of sensitivity analysis more interpretable (Cinelli and Hazlett 2020; Imbens 2003).

Second, computationally, BSA often runs into convergence issues because the model with sensitivity parameters is by design nonidentified. To cope with this problem, we develop an efficient Markov chain Monte Carlo (MCMC) algorithm by utilizing the transparent parameterization proposed by Gustafson (2005). Moreover, the shrinkage prior assigned to the coefficient of unobserved confounder induces substantial indirect learning of the coefficient given data, making results based on BSA significantly different from those based on Monte Carlo sensitivity analysis (MCSA; Greenland 2003), a procedure that does not involve the Bayesian updating of sensitivity parameters.

The rest of this article is organized as follows: in Section 2, we introduce our general framework for causal BSA, as well as the details of our proposed methodology; in Section 3, we illustrate details of computational issues; in Section 4, we conduct Monte Carlo studies; and in Section 5, we apply the proposed method to a real-world example. The last section concludes.

2. Methodology

2.1. The General Framework

Researchers often rely on unconfoundedness assumptions to infer causality from observational data, where they do not directly control the treatment assignment mechanism. A canonical example of such an assumption, often called *conditional ignorability*, posits that the treatment assignment is independent of the potential outcomes given a set of observed pre-treatment confounders, i.e.,

$$D_i \perp\!\!\!\perp \{Y_i(0), Y_i(1)\} \mid X_i, \quad (1)$$

where $D_i \in \{0, 1\}$ is the treatment assigned to unit i , $Y_i(d)$ is unit i 's potential outcome under treatment $d \in \{0, 1\}$, and X_i is a vector of observed pre-treatment covariates. In many real-world applications, however, this assumption may not hold. A typical example is observational studies with non-randomized treatment assignment, where researchers cannot rule out the possibility of unobserved confounders correlated with both the treatment and potential outcomes. Another example, which we focus on in this article, is longitudinal data where researchers collect repeated observations for each unit. In such data, unobserved unit-level heterogeneity in the potential outcomes and time-varying common shocks to the outcomes may also affect the adoption of treatment.

A common approach to addressing the violation of an unconfoundedness assumption is sensitivity analysis. Specifically, instead of the relationship (1), it is assumed that there is an unobserved confounder U_i such that

$$D_i \perp\!\!\!\perp \{Y_i(0), Y_i(1)\} \mid X_i, U_i, \quad (2)$$

which is also known as “latent ignorability” (Rubin 1976). Various sensitivity analysis methods have been developed on the basis of assumptions similar to (2) (e.g., Imbens 2003; Oster 2019; VanderWeele and Ding 2017). In this article, we propose a fully Bayesian approach, which directly builds upon the Bayesian causal inference framework of Rubin (1978) (see also Imbens and Rubin 1997; Rubin *et al.* 2008). Specifically, we consider causal inference to be the problem of imputing missing potential outcomes (Y_{mis}) given observed data (D, X and Y_{obs}) as well as the maintained assumptions. Under latent ignorability, we have

$$\begin{aligned} & Pr(Y_{mis} \mid D, X, Y_{obs}) \\ &= \frac{Pr(D, X, Y(0), Y(1))}{Pr(D, X, Y_{obs})} \\ &\propto Pr(D, X, Y(0), Y(1)) \\ &= \int Pr(D, X, U, Y(0), Y(1)) dU \\ &= \int Pr(D \mid X, U, Y(0), Y(1)) Pr(X, U, Y(0), Y(1)) dU \\ &= \int Pr(D \mid X, U) Pr(X, U, Y(0), Y(1)) dU, \end{aligned} \quad (3)$$

where the last equality holds because of (2). Note that $Pr(D \mid X, U)$ is the treatment assignment mechanism and $Pr(X, U, Y(0), Y(1))$ is what Rubin *et al.* (2008) call the “science.” Thus, the posterior predictive distribution of the missing outcome $Pr(Y_{mis} \mid D, X, Y_{obs})$ is proportional to the treatment assignment mechanism multiplied by the underlying science, marginalizing over the unobserved confounder. Many existing sensitivity analysis techniques implicitly utilize this general relationship with additional parametric assumptions for $Pr(D \mid X, U)$ and $Pr(X, U, Y(0), Y(1))$ (e.g., Imbens 2003).

Here, we pursue an alternative factorization which establishes an explicit connection between Rubin's Bayesian causal inference framework and the Bayesian, non-counterfactual framework for sensitivity analysis (McCandless *et al.* 2007). That is, we have

$$\begin{aligned} & Pr(Y_{mis} \mid D, X, Y_{obs}) \\ &\propto \int Pr(D \mid X, U) Pr(X, U, Y(0), Y(1)) dU \end{aligned}$$

$$\begin{aligned}
&= \int \frac{Pr(D, X, U)}{Pr(X, U)} Pr(X, U, Y(0), Y(1)) dU \\
&= \int Pr(U|D, X) Pr(D, X) \frac{Pr(X, U, Y(0), Y(1))}{Pr(X, U)} dU \\
&\propto \int Pr(U|D, X) Pr(Y(0), Y(1)|X, U) dU,
\end{aligned} \tag{4}$$

suggesting an alternative procedure for inference based on the imputation of the latent variable U as a function of D and X , as well as a model for the potential outcomes. Indeed, the standard BSA framework utilizes this imputation-based strategy for inference (Gustafson *et al.* 2010; McCandless *et al.* 2007), first imputing U using a parametric model for $Pr(U|D, X)$, and then including the imputed values of U in an outcome model $f(Y_{obs}|D, U, X)$. That is, a model for the observed data (Y_{obs}, D, X) can be obtained by integrating U out:

$$\begin{aligned}
Pr(Y_{obs}, D, X) &\propto Pr(Y_{obs}|D, X) \\
&= \int Pr(Y_{obs}, U|D, X) dU \\
&\propto \int Pr(U|D, X) Pr(Y_{obs}|D, X, U) dU.
\end{aligned} \tag{5}$$

Comparing (5) to the posterior predictive distribution (4) reveals that the standard BSA framework can be used for sensitivity analysis with respect to causal quantities involving missing potential outcomes. That is, a latent ignorability assumption, such as (2), implies sufficient constraints on the joint potential outcome distribution conditional on X and U so that a causal quantity of interest is identified by the distribution of the observed outcome Y_{obs} conditional on D , X , and U .

2.2. Setup and Notation

Now, we set up our notational framework. Suppose that we observe a panel data set that consists of N units and spans over T periods, with $i = 1, 2, \dots, N$ and $t = 1, 2, \dots, T$ indexing unit and time, respectively. We are interested in the effect of a binary treatment indicator, $D_{it} \in \{0, 1\}$, on some outcome variable $Y_{it} \in \mathbb{R}$. For notational convenience, we assume that the panel data set is balanced, i.e., there are no missing observations. We adopt the potential outcomes framework (Rubin 1974) to define the causal quantities of interest. We denote by $\mathbf{D}_i = (D_{i1}, D_{i2}, \dots, D_{iT})$ the full treatment assignment vector for unit i and $\mathbf{D}_i^t = (D_{i1}, D_{i2}, \dots, D_{it})$ the treatment assignment vector up to time t for unit i . $Y_{it}(\mathbf{d}_i)$ and $Y_{it}(\mathbf{d}_i')$ are two different potential outcomes for unit i in period t given treatment assignment vectors $\mathbf{d}_i = (d_{i1}, d_{i2}, \dots, d_{iT})$ and $\mathbf{d}_i' = (d'_{i1}, d'_{i2}, \dots, d'_{iT})$. We make the following assumptions.

Assumption 1 (No anticipation of future treatments). *Treatment assignment in the future does not affect current potential outcomes, i.e.,*

$$Y_{it}(\mathbf{d}_i) = Y_{it}(\mathbf{d}_i') \quad \text{if} \quad \mathbf{d}_i^t = \mathbf{d}_i'^t \tag{6}$$

for any $i \in \{1, \dots, N\}$ and $t \in \{1, \dots, T\}$.

Assumption 2 (No carryover effects). *Past treatment assignments do not affect current potential outcomes, i.e.,*

$$Y_{it}(\mathbf{d}_i^t) = Y_{it}(\mathbf{d}_i'^t) \quad \text{if} \quad d_{it} = d'_{it} \tag{7}$$

for any $i \in \{1, \dots, N\}$ and $t \in \{1, \dots, T\}$.¹

¹This assumption is also called *no dynamics* in Arkhangelsky and Imbens (2022).

Under these assumptions, we can re-write the potential outcome as $Y_{it}(\mathbf{d}_i) = Y_{it}(d_{it})$, and there are two potential outcomes $Y_{it}(1)$ and $Y_{it}(0)$ for each unit in each period. The individual treatment effect is defined as the difference in potential outcomes:

$$\delta_{it} = Y_{it}(1) - Y_{it}(0). \quad (8)$$

We focus on the identification and estimation of the average treatment effect on the treated (ATT), defined as $ATT = \frac{\sum_{i,t} D_{it} \delta_{it}}{\sum_{i,t} D_{it}}$. The results also apply to other convex combinations of individual effects, like ATT in a specific period relative to the occurrence of the treatment.

We consider three types of unobserved confounders: unit-specific time-invariant confounders γ_i , time-specific unit-invariant confounders f_t , and confounders that vary both across units and time U_{it} . Below, we will develop a method for sensitivity analysis based on the following latent ignorability assumption with respect to these unobserved confounders.

Assumption 3 (Panel latent ignorability).

$$\{Y_{it}(0), Y_{it}(1)\} \perp\!\!\!\perp D_{it} \mid X_{it}, \gamma_i, f_t, U_{it} \quad (9)$$

for $i \in \{1, \dots, N\}$ and $t \in \{1, \dots, T\}$, where X_{it} is a vector of observed confounders.

2.3. Causal Panel Data Models

The panel latent ignorability assumption implies that we need to control for observed confounders X_{it} as well as the unobserved confounders γ_i, f_t , and U_{it} to identify the effect of D_{it} on Y_{it} . Allowing U_{it} to be time-varying enables the model to capture both unobserved time-varying confounders and the time-varying influence of time-invariant confounders, making it a more flexible specification. Suppose the observed confounder X_{it} is a vector of length P , i.e., $X_{it} = (X_{it,1}, X_{it,2}, \dots, X_{it,P})$, where each element $X_{it,p}$ is normalized to have zero mean and unit variance. Existing literature has proposed both design-based approaches like doubly-robust estimator (Arkhangelsky and Imbens 2022) and model-based approach to counterfactual outcome prediction. The latter is widely adopted in empirical studies, where scholars assume a parametric model for the potential outcomes. A general functional form for the observed outcome model is written as:

$$Y_{it} = g(D_{it}, X_{it}, \gamma_i, f_t, U_{it}, \varepsilon_{it}), \quad (10)$$

where ε_{it} is an error term with zero mean. In empirical studies, scholars add more constraints on the functional form to identify causal parameters despite the presence of unobserved confounding. For example, Pang, Liu, and Xu (2022) assume that the observed terms in $g(\cdot)$ are additive and that unobserved confounding can be represented by functions of γ_i and f_t , such as the TWFE and the latent factor term or IFE in the SC methods. With a statistical model for potential outcomes, scholars can predict the counterfactual outcomes under control for the treated observations and then average the predicted differences between the observed outcomes and predicted counterfactuals as an estimate for the ATT. Such counterfactual prediction-based approaches are called “causal panel data models” in Athey *et al.* (2021).

Here, we focus on a linear model incorporating γ_i and f_t by way of a two-way interaction term or a latent factor term:

$$Y_{it} = \delta_{it} D_{it} + \beta_0 + X'_{it} \beta + \gamma'_i f_t + \varepsilon_{it}, \quad (11)$$

$$\gamma_i \mid \Omega_0 \sim \mathcal{N}(0, \Omega_0), \quad f_t \sim \mathcal{N}(0, I_r), \quad \varepsilon_{it} \sim \mathcal{N}(0, \sigma^2),$$

where β_0 is the global intercept and $\beta = (\beta_1, \beta_2, \dots, \beta_P)$ is the slope coefficients for X_{it} . In this setup, γ_i can be regarded as an $(r \times 1)$ vector of factor loadings and f_t an $(r \times 1)$ vector of common factors. The number of factors r is unknown by the researcher *a priori*. The functional form above implies that $Y_{it}(0) = \beta_0 + X'_{it} \beta + \gamma'_i f_t + \varepsilon_{it}$, i.e., we assume a latent factor model (LFM) for counterfactual outcomes

under control. From a Bayesian perspective, Model (11) is a static variant of the state space model (Klinenberg 2022; Pang *et al.* 2022). Following the literature of Bayesian factor analysis, we assume that the factor loadings are drawn from an i.i.d. multivariate normal distribution $\mathcal{N}(0, \Omega_0)$, where $\Omega_0 = \text{Diag}\{\omega_1^2, \dots, \omega_r^2\}$ is an $(r \times r)$ diagonal matrix, and that common factors are drawn from an i.i.d. multivariate normal distribution $\mathcal{N}(0, I_r)$, where I_r is the identity matrix. The error term ε_{it} is drawn from i.i.d. normal distribution with mean 0 and variance σ^2 . Model (11) entails the assumption that the panel latent ignorability (9) holds without conditioning on U_{it} , i.e., unobserved confounding that varies both over units and time can be fully captured by the r -dimensional factor term $\gamma'_{it}f_t$.

It is worth noting that there is an extensive literature on estimating Model (11) using both frequentist and Bayesian approaches.² While the Bayesian approach offers greater flexibility in estimation, it relies on additional distributional assumptions for the factor term and the error term. In the frequentist approach, the factor term $\gamma'_{it}f_t$, also known as “interactive fixed effects” (Bai 2009), can be used to address endogeneity for the identification of coefficients β , when some covariates in X_{it} are endogenous. Additionally, instead of assuming that error terms are drawn from an independent standard normal distribution, they can be serially correlated within a unit.

In this article, we focus on BSA for treatment effects and thus assume independence between the factor term and the observed covariates, with i.i.d. error terms. However, Model (11) can be extended to incorporate endogenous covariates by directly modeling the relationship between X_{it} and $\gamma'_{it}f_t$. For the error terms, while we assume that ε_{it} are i.i.d., the factor term, which includes unit random effects as a special case, provides a way to account for within-unit correlation (Cameron and Miller 2015). More complex structures for the error term, such as stochastic volatility (Belmonte, Koop, and Korobilis 2014), can also be incorporated.

Determining the number of factors is a frequently discussed issue in the literature on estimating LFM-like (11). The prior variances for factor loadings, $\{\omega_1^2, \dots, \omega_r^2\}$, which control the effective number of factors, play a key role in the Bayesian analysis of factor models. For example, if $\omega_1^2 = 0$, then the factor loading that corresponds to the first factor equals 0 for each unit, and thus it has no effect on the potential outcomes and is essentially excluded from the model. However, a well-known challenge in specifying prior for variance is that the conjugate inverse gamma distribution performs poorly for the purpose of selection of factor numbers. Since the support of the inverse gamma distribution includes only positive real numbers, even a conventional vague inverse gamma prior exerts a substantial influence on the posterior of ω^2 when its true value is close to 0 (Frühwirth-Schnatter and Wagner 2010). There are two main approaches to this problem, both based on Bayesian shrinkage: spike-and-slab prior and continuous shrinkage prior. In this article, we adopt the continuous shrinkage approach. We follow Frühwirth-Schnatter and Wagner (2010) and re-parameterize the factor loadings as follows:

$$\begin{aligned} y_{it} &= \beta_0 + \delta_{it}D_{it} + X'_{it}\beta + (\omega_\gamma \cdot \tilde{\gamma}_{it})'f_t + \varepsilon_{it}, \\ \tilde{\gamma}_{it} &\sim \mathcal{N}(0, I_r), \quad f_t \sim \mathcal{N}(0, I_r), \quad \varepsilon_{it} \sim \mathcal{N}(0, \sigma^2), \end{aligned} \quad (12)$$

where the factor loadings γ_{it} in model (11) are re-parameterized as the product of ω_γ and $\tilde{\gamma}_{it}$. In model (12), $\tilde{\gamma}_{it}$ is a vector of “standardized” factor loadings with each entry drawn from i.i.d. standard normal distribution and ω_γ is now interpreted as a vector of coefficients controlling the influence of each factor. If the k^{th} element in ω_γ equals 0, then the k^{th} factor is excluded from the model. The re-parameterization allows the use of a prior distribution for ω_γ that contains 0 in the interior of its support, enabling model selection by shrinkage.³

²Bayesian methods for causal inference have been increasingly developed and applied in political science research (Gill 2012). For example, Hollenbach, Montgomery, and Crespo-Tenorio (2019) propose a model-based approach, and Keele and Quinn (2017) develop a sensitivity analysis method for binary outcome data. Bayesian approaches have also been applied to quantify the uncertainty of Rosenbaum bounds in sensitivity analysis (Mebane Jr and Poast 2013).

³The factor terms are unidentified because of the label-switching problem, such that $\omega_\gamma \cdot \tilde{\gamma}_{it} = (-\omega_\gamma) \cdot (-\tilde{\gamma}_{it})$. Following the literature, we introduce a step of random permutation in our sampler to address this problem.

2.4. Bayesian Sensitivity Analysis

Although the factor term in model (11) is highly flexible, it may not fully capture the unobserved confounding that varies across units and time. To address this common concern, we develop a sensitivity analysis method for treatment effect estimation under Assumption 3, where the existence of general confounding varying across units and time is permitted. As shown above in terms of a general setting (4) and (5), Assumption 3 enables imputation-based sensitivity analysis for causal effects by jointly modeling Y_{it} conditional on $(D_{it}, X_{it}, \gamma_i, f_i, U_{it})$ and U_{it} conditional on $(D_{it}, X_{it}, \gamma_i, f_i)$. Specifically, we expand Model (12) to incorporate U_{it} in the model:⁴

$$\begin{aligned} Y_{it} &= \beta_0 + \delta_{it}D_{it} + X'_{it}\beta + \beta_u U_{it} + (\omega_\gamma \cdot \tilde{\gamma}_i)'f_i + \varepsilon_{it}, \\ U_{it} &= \lambda_0 + \lambda_d D_{it} + X'_{it}\lambda_x + \lambda_f (\omega_\gamma \cdot \tilde{\gamma}_i)'f_i + \nu_{it}, \\ \tilde{\gamma}_i &\sim \mathcal{N}(0, I_r), \quad f_i \sim \mathcal{N}(0, I_r), \\ \varepsilon_{it} &\sim \mathcal{N}(0, \sigma^2), \quad \nu_{it} \sim \mathcal{N}(0, c_1^2), \end{aligned} \quad (13)$$

where the coefficient β_u measures the effect of U_{it} on the outcome of interest. We follow common practice (e.g., Gustafson *et al.* 2010) to model U_{it} as a linear function of the treatment indicator D_{it} , observed confounders X_{it} , as well as the factor terms $(\omega_\gamma \cdot \tilde{\gamma}_i)'f_i$, with their associated coefficients denoted by $\lambda \equiv [\lambda_0, \lambda_d, \lambda'_x, \lambda'_f]'$. The variance of the normal error term ν_{it} equals c_1^2 , a fixed constant specified by the researcher. We call Model (13) “Bayesian sensitivity analysis for the linear factor model” (BSA-LFM).

λ and β_u represent our sensitivity parameters and deserve several additional remarks. First, as discussed in the next section, we specify the prior for each element in λ to be symmetric and centered at zero, indicating that the grand mean of the unobserved confounder is 0 *a priori*. Second, existing methods for sensitivity analysis often assume that the unobserved confounder is only determined by the treatment indicator to simplify the analysis (e.g., Lin, Psaty, and Kronmal 1998). However, the assumption that U_{it} is independent of the observed covariates and the factor term conditional on treatment is often unrealistic and has been challenged by VanderWeele (2008).⁵ In addition, the common correlations of the outcome and the unobserved confounder imply additional learning (Gustafson *et al.* 2010). That is, while β_u and λ_d are the only two sensitivity parameters that affect the estimation of the treatment effect,⁶ allowing non-zero correlation of U_{it} with observed confounders as well as the factor terms can help narrow the inference. If at least one element in λ_x or λ_f is not equal to 0, we can achieve substantive updating in β_u given the data. Furthermore, allowing for correlation between observed covariates and unobserved confounders implies that some covariates may be endogenous (Diegert, Masten, and Poirier 2022), which is a weaker condition than treating all covariates as exogenous.

Finally, it is important to note that we must assume some correlation between λ_0 and λ_d . Since Model (13) allows for arbitrary heterogeneous treatment effects, like most counterfactual imputation methods, only observations under control ($D_{it} = 0$) are used for model fitting. For these observations, Model (13) reduces to:

⁴ An alternative specification of U_{it} assumes that it is orthogonal to the observed covariates and the latent factor term conditional on D_{it} ; that is, $U_{it} = \lambda_0 + \lambda_d D_{it} + \nu_{it}$. As a reviewer noted, this specification reduces the number of sensitivity parameters, thereby simplifying the model. This approach aligns with the MCSA (Greenland 2003), which draws inference based on the prior distribution of sensitivity parameters. However, as Gustafson and McCandless (2018) shows, compared to MCSA, the proposed BSA, which links U_{it} to the observed covariates and latent factors, enables inference based on the posterior distribution of sensitivity parameters (e.g., β_u) and may yield more precise uncertainty estimates.

⁵ As noted by VanderWeele (2008), if neither the observed covariates nor the factor term affects treatment assignment, it is reasonable to assume $\lambda_0 = \lambda_x = \lambda_f = 0$. This conventional practice is included as a special case within our proposed model.

⁶ As one reviewer helpfully pointed out, re-parameterizing the model of U to express D as a function of U makes it clearer that λ_d , rather than the other λ terms or β_u , is the essential sensitivity parameter. Suppose D is continuous. Then, the model of U can be equivalently written as: $D_{it} = -\frac{\lambda_0}{\lambda_d} + \frac{1}{\lambda_d} U_{it} - \left\{ \frac{\lambda_x}{\lambda_d} X_{it} + \frac{\lambda_f}{\lambda_d} (\omega_w \cdot \tilde{\gamma}_i)'f_i + v'_{it} \right\}$, where $v'_{it} = -\frac{1}{\lambda_d} \nu_{it}$. This re-parameterization demonstrates that $1/\lambda_d$ is the key sensitivity parameter.

$$\begin{aligned} Y_{it|D_{it}=0} &= \beta_0 + X'_{it}\beta + \beta_u U_{it} + (\omega_\gamma \cdot \tilde{\gamma}_i)' f_t + \varepsilon_{it}, \\ U_{it|D_{it}=0} &= \lambda_0 + X'_{it}\lambda_x + \lambda_f (\omega_\gamma \cdot \tilde{\gamma}_i)' f_t + \nu_{it}. \end{aligned} \quad (14)$$

Thus, we can only infer λ_0 directly from the data and indirectly estimate λ_d through its assumed correlation with λ_0 . Details of the prior specifications are provided below.

2.5. Prior Specification

Prior specification plays a key role in BSA. Because data contain no direct information about unobserved confounders, posteriors of sensitivity parameters (λ and β_u) are dominated by priors assigned to them. In fact, priors for the sensitivity parameters encode the researcher's *a priori* belief about the direction and magnitude of unobserved confounding. Such beliefs are typically incorporated in frequentist sensitivity analysis only partially and informally by fixing the sensitivity parameters to particular values. In contrast, our Bayesian approach provides a more formal, principled way to represent researcher's belief about confounding in the form of a probability distribution.

From a practical standpoint, the choice of a prior distribution for sensitivity parameters entails a bias–variance trade-off. An overly diffuse prior would result in a credible interval that has low power to detect a non-zero treatment effect, while a prior with high precision risks bias that never disappears asymptotically. Here, we adopt the approach proposed by McCandless, Gustafson, and Levy (2008), which makes use of observed confounders to benchmark our prior on the effects of unobserved confounders. That is, we assume that the confounding effects of U_{it} and $X_{it,j}$ ($j \in \{1, \dots, p\}$) are comparable *a priori* after standardization, in the sense that β_u and β have exchangeable prior distributions. Formally, we have

$$\begin{aligned} \mathbb{E}(U_{it}) &= 0, \quad \text{Var}(U_{it}) = 1, \\ f_{\beta, \theta}(\beta_0, \beta_1, \dots, \beta_p, \beta_u) &= \int \prod_{k=0}^p f_\beta(\beta_k | \theta) f_\beta(\beta_u | \theta) f_\theta(\theta) d\theta, \end{aligned} \quad (15)$$

where the second line in (15) is by de Finetti's theorem, and the first line adds several constraints on the prior distributions of sensitivity parameters λ as well as the variance of ν_{it} . We assign conjugate normal priors to λ with constraints on prior variances⁷ and exchangeable priors to β_u and β . More specifically, our prior on λ is specified as follows:

- λ_0 and λ_d :

$$\begin{aligned} (\lambda_0, \lambda_d)' &\sim \mathcal{N}(0, A), \\ A &= \begin{pmatrix} c_2^2 & -c_3^2/2 \\ -c_3^2/2 & c_3^2 \end{pmatrix}, \end{aligned} \quad (16)$$

where the hyperparameters are constrained such that $0 < c_2 \leq 1$, $c_3 > 0$, and $2c_2 > c_3$ to guarantee positive definiteness of the covariance matrix.

- λ_x :

$$\lambda_x \sim \mathcal{N}\left(0, \frac{c_4^2}{p} I_p\right). \quad (17)$$

- λ_f :

$$\lambda_f \sim \mathcal{N}\left(0, \frac{c_5^2(k_1 - 1)}{2rk_2}\right), \quad (18)$$

⁷As one reviewer pointed out, another common choice is to assign uniform priors to the sensitivity parameters, reflecting vague prior knowledge. In this case, it is necessary to specify lower and upper bounds to satisfy the constraint on the variance of U_{it} .

where k_1 and k_2 are hyperparameters for ω_γ defined below. Taken together, our prior specification for λ implies the following proposition.⁸

Proposition 1 (Prior variance of unobserved confounder). *Priors (16)–(18) imply that the prior variance of U_{it} is independent of treatment status D_{it} and given by,*

$$\text{Var}(U_{it}|D_{it}) = \text{Var}(U_{it}) = c_1^2 + c_2^2 + c_4^2 + c_5^2. \quad (19)$$

A formal proof of Proposition 1 can be found in Appendix 1 of the Supplementary Material. Since we standardize the observed and unobserved confounders to have unit variance, Proposition 1 implies the following constraint on the hyperparameters:

$$c_1^2 + c_2^2 + c_4^2 + c_5^2 = 1. \quad (20)$$

Now, we turn to parameters in the outcome model and prior variances of factor loadings. There are several sensible options for exchangeable priors assigned to β and β_u . Of particular note, while an independent, weakly informative normal prior is perhaps the most intuitive choice, Gustafson *et al.* (2010) use the i.i.d. t distribution with a certain degree of freedom, noting that the latter often yields substantially narrower inferences than a normal independence prior with the same marginal variability. This prior has a hierarchical form: the priors for $\beta_0, \beta_1, \dots, \beta_p, \beta_u$ are i.i.d. $N(0, \tau^2 S)$ given S , and the hyperparameter S is drawn from inverted-gamma distribution $S \sim \text{Inv-Gamma}(d/2, d/2)$.

Here, we propose a more flexible specification, allowing each β_k and β_u to have heterogeneous prior variances drawn from a multi-level model. Specifically, we use the Laplacian prior (Park and Casella 2008), which is a continuous shrinkage prior that have heavier tails than the t distribution and more densities near 0 to “let the data speak for itself.”

- β_j for $j \in \{0, 1, \dots, p\}$:⁹

$$\begin{aligned} \beta_j | \tau_{\beta_j}^2 &\sim \mathcal{N}(0, \tau_{\beta_j}^2), \\ \tau_{\beta_j}^2 | \xi_\beta &\sim \text{Exp}\left(\frac{\xi_\beta^2}{2}\right), \\ \xi_\beta^2 &\sim \text{Gamma}(a_1, a_2). \end{aligned} \quad (21)$$

- β_u :

$$\begin{aligned} \beta_u | \tau_{\beta_u}^2 &\sim \mathcal{N}(0, \tau_{\beta_u}^2), \\ \tau_{\beta_u}^2 | \xi_\beta &\sim \text{Exp}\left(\frac{\xi_\beta^2}{2}\right), \\ \xi_\beta^2 &\sim \text{Gamma}(a_1, a_2). \end{aligned} \quad (22)$$

We also assign the Laplacian prior to ω_γ for the purpose of factor number selection.

- ω_{γ_j} for $j \in \{1, \dots, r\}$:

$$\begin{aligned} \omega_{\gamma_j} | \tau_{\gamma_j}^2 &\sim \mathcal{N}(0, \tau_{\gamma_j}^2), \\ \tau_{\gamma_j}^2 | \xi_\gamma &\sim \text{Exp}\left(\frac{\xi_\gamma^2}{2}\right), \\ \xi_\gamma^2 &\sim \text{Gamma}(k_1, k_2). \end{aligned} \quad (23)$$

⁸If the researcher chooses $k_1 \leq 1$, another prior specification for λ_f is $\lambda_f \sim \mathcal{N}\left(0, \frac{c_5^2(k_1+1)}{2rk_2}\right)$. See the Supplementary Material for more details.

⁹The parameterization of gamma distribution has mean = $\frac{a_1}{a_2}$.

Table 1. Suggested default specifications for the hyperparameters.

Hyperparameter	Function	Default values
$(c_1^2, c_2^2, c_3^2, c_4^2, c_5^2)$	Controlling the dispersion of U_{it}	Subject to restrictions
(a_1, a_2)	Controlling the effects of X_{it} and U_{it}	(0.001, 0.001)
(k_1, k_2)	Determining the number of factors	(1.001, 0.001)
(e_1, e_2)	Controlling the error variance	(0.001, 0.001)

Finally, we use the conventional inverted-gamma prior for the error variance σ^2 .

• σ_ε^2 :

$$\sigma^2 \sim \text{Inv} - \text{Gamma}(e_1, e_2). \tag{24}$$

Taken together, our model for sensitivity analysis requires 11 hyperparameter values to be set by the researcher, subject to the constraints as specified above. Appropriately setting hyperparameter values is often a challenging task in implementing Bayesian methods in practice, particularly for complex models. Fortunately, sensible default values that approximate a lack of prior knowledge are available for most of the hyperparameters in our model, including those related to the error variance, coefficients in the outcome model, and factor selection. This substantially simplifies the task of specifying priors in applied research.

Table 1 summarizes such default values for these hyperparameters, along with their substantive interpretations. For the prior variances of the sensitivity parameters λ s, we impose constraints to keep their scales small while allowing β_u to be potentially large. In practice, selecting these values may require some prior knowledge. For instance, if researchers believe that U_{it} is informative given X_{it} or the factor terms, they may choose relatively large values for c_4^2 or c_5^2 . Otherwise, they may favor a larger random component, reflected in a larger c_1^2 .

3. Computation

Suppose the true data-generating process (DGP) is model (13) but we fit the LFM in model (12) without the unobserved confounder U_{it} . Then, the bias for the ATT estimate is represented by the product $\beta_u \lambda_d$. Frequentist sensitivity analysis usually sets a range of values for β_u and λ_d and investigates how large their values must be to explain away the estimated treatment effect (Imbens 2003; Rosenbaum and Rosenbaum 2002). By contrast, BSA assigns priors to all parameters and makes inference on the posterior distribution. Model (12) implies a joint distribution of the observed outcome Y_{it} and the unobserved confounder U_{it} , and thus it seems appealing to apply the MCMC algorithm to the joint posterior distributions of parameters as well as U_{it} . However, since none of the elements in the parameter vector $\theta = \{\beta_0, \beta, \beta_u, A, F, \lambda_0, \lambda_d, \lambda_x, \lambda_f, \sigma^2\}$ is known, the model is nonidentified. For example, define $\tilde{\theta} = \{\beta_0, \beta, \beta_u/2, A, F, 2\lambda_0, 2\lambda_d, 2\lambda_x, 2\lambda_f, \sigma^2\}$, an alternative parameter vector. Then, $f(Y|\theta) = f(Y|\tilde{\theta})$.

The non-identifiability of such a model results in practical issues of simulation-based model fitting, like the MCMC. For example, Gelfand and Sahu (1999) argue that with flat priors, the trajectories of the Markov chain for some parameters tend to drift to extreme values, causing poor mixing of Markov chains and make them hard to converge. In this article, we address the issue of computation by adopting the “transparent parameterization” of the original nonidentified model (13) to improve the efficiency of the MCMC sampler, as recommended by Gustafson (2005).

3.1. Transparent Parameterization of the Nonidentified Model

We begin by briefly summarizing the general idea of transparent parameterization proposed by Gustafson (2005). Suppose the original parameter vector in the nonidentified model is θ .

we reparameterize θ to new parameter vector $\phi \equiv \{\phi_I, \phi_N\}$ such that $f(\text{data}|\phi) = f(\text{data}|\phi_I)$. Priors assigned to ϕ can be factorized as $f(\phi) = f(\phi_N|\phi_I)f(\phi_I)$, and the corresponding posteriors are:

$$\begin{aligned} f(\phi_I|\text{data}) &\propto f(\text{data}|\phi_I)f(\phi_I), \\ f(\phi_N|\text{data}) &= \int f(\phi_N|\phi_I)f(\phi_I|\text{data})d\phi_I. \end{aligned} \quad (25)$$

The posteriors (25) imply qualitatively different asymptotic behaviors for ϕ_I and ϕ_N . That is, *direct Bayesian learning* occurs for ϕ_I , in that under regularity conditions, $f(\phi_I|\text{data})$ converges to the point probability mass at the true value of ϕ_I as sample size grows, and *indirect Bayesian learning* for ϕ_N , which converges to the conditional distribution $f(\phi_N|\phi_I)$ at the true value of ϕ_I . This, in turn, implies that ϕ_I and ϕ_N will also exhibit distinct behaviors in an MCMC sampler.

Gustafson *et al.* (2010) suggest that transparent parameterization could be exploited for the purpose of building a computationally efficient posterior sampler, making jumps in spaces for those parameters informed by the data while moving fast through spaces for those parameters with little information from the data. We adopt a similar strategy for our proposed BSA-LFM. Specifically, we substitute U_{it} in the first line for potential outcome with the second line in model (13) that describe the DGP of U_{it} , essentially integrating out the unobserved confounder:

$$Y_{it} = (\delta_{it} + \beta_u \lambda_d) D_{it} + (\beta_0 + \beta_u \lambda_0) + X'_{it}(\beta + \beta_u \lambda_x) + (1 + \beta_u \lambda_f)(\omega_\gamma \cdot \tilde{\gamma}_i)' f_t + \beta_u \nu_{it} + \varepsilon_{it}. \quad (26)$$

Define $\tilde{\delta}_{it} = \delta_{it} + \beta_u \lambda_d$, $\tilde{\beta}_0 = \beta_0 + \beta_u \lambda_0$, $\tilde{\beta} = \beta + \beta_u \lambda_x$, $\tilde{\omega}_\gamma = (1 + \beta_u \lambda_f)\omega_\gamma$, and $\tilde{\varepsilon}_{it} = \beta_u \nu_{it} + \varepsilon_{it}$. Model (13) can then be re-written as:

$$\begin{aligned} Y_{it} &= \tilde{\delta}_{it} D_{it} + \tilde{\beta}_0 + X'_{it} \tilde{\beta} + (\tilde{\omega}_\gamma \cdot \tilde{\gamma}_i)' f_t + \tilde{\varepsilon}_{it}, \\ \tilde{\varepsilon}_{it} &\sim \mathcal{N}(0, \tilde{\sigma}^2), \\ \tilde{\sigma}^2 &= \beta_u^2 c_1^2 + \sigma_Y^2. \end{aligned} \quad (27)$$

The original parameter vector is $\theta = \{\beta_0, \beta, \omega, \Gamma, F, \sigma_Y^2, \beta_u, \lambda_0, \lambda_d, \lambda_x, \lambda_f\}$, and the re-parameterized vector is:

$$\phi = \{\tilde{\beta}_0, \tilde{\beta}, \tilde{\omega}, \Gamma, F, \tilde{\sigma}^2, \beta_u, \lambda_0, \lambda_d, \lambda_x, \lambda_f\}, \quad (28)$$

where $\phi_I = \{\tilde{\beta}_0, \tilde{\beta}, \tilde{\omega}, \Gamma, F, \tilde{\sigma}^2\}$ and $\phi_N = \{\beta_u, \lambda_0, \lambda_d, \lambda_x, \lambda_f\}$. One can verify that the data likelihood only depends on ϕ_I :¹⁰

$$\begin{aligned} f(Y^{obs}|\phi) &\propto (\sqrt{\sigma^2 + \beta_u^2 c_1^2})^{-(NT - \sum_{i,t} D_{it})} \\ &\quad \exp\left[-\frac{1}{2(\sigma^2 + \beta_u^2 c_1^2)} \sum_{D_{it}=0} (y_{it} - (\beta_0 + \beta_u \gamma_0) - X'_{it}(\beta + \beta_u \gamma_x) - (\omega_\gamma (1 + \beta_u \gamma_f) \cdot \tilde{\gamma}_i)' f_t)^2\right] \\ &= \tilde{\sigma}^{-(NT - \sum_{i,t} D_{it})} \exp\left[-\frac{1}{2\tilde{\sigma}^2} \sum_{D_{it}=0} (y_{it} - \tilde{\beta}_0 - X'_{it} \tilde{\beta} - (\tilde{\omega}_\gamma \cdot \tilde{\gamma}_i)' f_t)^2\right] \\ &\propto f(Y^{obs}|\phi_I). \end{aligned} \quad (29)$$

Now, we derive the posteriors for parameters in the re-parameterized model (27). Denote the one-to-one mapping of the transparent re-parameterization by $\phi = g(\theta)$. Prior distributions assigned to the original parameter vector θ imply priors assigned to ϕ are:

$$f_\phi(\phi) = f_\theta(g^{-1}(\phi))|J|, \quad (30)$$

¹⁰For counterfactual imputation-based approach, only observations under control status are used for model fitting.

where $|J|$ is the determinant of the Jacobian $\theta = g^{-1}(\phi)$. Thus, posterior distributions of ϕ can be written as:

$$\begin{aligned}\pi(\phi|Y^{obs}) &\propto f(Y^{obs}|\phi)f_{\phi}(\phi) \\ &= f(Y^{obs}|\phi)f_{\theta}(g^{-1}(\phi))|J| \\ &= f(Y^{obs}|\phi)|1 + \beta_u\lambda_f|^{-r} \\ &\quad \prod_{j=0}^p \pi(\tilde{\beta}_j - \beta_u\lambda_x|\tau_{\tilde{\beta}_j}^2)\pi(\tau_{\tilde{\beta}_j}^2|\xi_{\tilde{\beta}}^2) \prod_{k=1}^r \pi\left(\frac{\tilde{\omega}_{\gamma_j}}{1 + \beta_u\lambda_f}|\tau_{\gamma_j}^2\right)\pi(\tau_{\gamma_j}^2|\xi_{\gamma}^2)\pi(\xi_{\gamma}^2) \\ &\quad \pi(\beta_u|\tau_{\beta_u}^2)\pi(\tau_{\beta_u}^2|\xi_{\beta}^2)\pi(\tilde{\Gamma})\pi(F)\pi(\tilde{\sigma}^2 - \beta_u^2c_1^2)\pi(\lambda_0, \lambda_d)\pi(\lambda_x)\pi(\lambda_f),\end{aligned}\tag{31}$$

where $\pi(\cdot)$ denotes the prior distribution of the corresponding original parameter. Then, we can recover the posteriors of the original parameter vector θ from the posteriors (31) of ϕ . Since prior distributions are assigned to θ , the implied priors are not conjugate for some parameters in ϕ . Therefore, we develop a Metropolis–Hastings step within Gibbs sampling procedure to simulate the posteriors summarized as Algorithm 1. Details of the MCMC algorithm can be found in Appendix 2 of the Supplementary Material.

Remark. In Step 10 of Algorithm 1, we estimate the individual treatment effect for observations under treatment ($D_{it} = 1$) using $\delta_{it} = \tilde{\delta}_{it} - \beta_u \cdot \lambda_d$, where $\tilde{\delta}_{it} = Y_{it} - Y_{it}(0)$, with $Y_{it}(0)$ representing the imputed counterfactual outcome under control. This step demonstrates the “frequentist analogy” of the proposed BSA. To see this, suppose the treatment effects are homogeneous ($\delta_{it} \equiv \delta$). In that case, the identified parameter is $\tilde{\delta} = \delta + \beta_u \cdot \lambda_d$. Researchers can obtain an interval estimate for δ by fixing the sensitivity parameters at specific values and varying these values. The BSA approach, rather than selecting specific values for the sensitivity parameters, marginalizes the treatment effect estimate over the posterior distributions of β_u and λ_d .

Algorithm 1 A Metropolis–Hastings within Gibbs sampler.

At iteration ($h + 1$):

1. Update coefficients $\tilde{\beta}_0, \tilde{\beta}$,
2. Update factor loadings $\tilde{\gamma}_i$ for $i \in \{1, \dots, N\}$,
3. Update factors f_t for $t \in \{1, \dots, T\}$,
4. Update the square root of prior variance for factor loadings $\tilde{\omega}_{\gamma}$,
5. Update error variance in the re-parameterized model $\tilde{\sigma}^2$,
6. Update sensitivity parameters λ_0, λ_{x_j} for $j \in \{1, \dots, p\}$, λ_d, λ_f and β_u ,
7. Update prior variances $\tau_{\tilde{\beta}_j}^2, j \in \{0, 1, \dots, p\}$, $\tau_{\beta_u}^2, \tau_{\gamma_j}^2$ for $j \in \{1, \dots, r\}$,
8. Update hyperparameters ξ_{β}^2 and ξ_{γ}^2 ,
9. Update counterfactual outcome $Y_{it}(0)$ and individual treatment effect $\delta_{it} = Y_{it} - Y_{it}(0)$ for $D_{it} = 1$,
10. Average over δ_{it} for $D_{it} = 1$ as an estimate of ATT.

Repeat Steps 1-10 at the next iteration until the Markov chains converge.

4. Monte Carlo Studies

In this section, we conduct a series of Monte Carlo studies to investigate the behavior of the proposed BSA (BSA-LFM) in terms of key frequentist performance metrics. We focus on the ATT ($\sum_{i,t} D_{it}\delta_{it} / \sum_{i,t} D_{it}$) as a quantity of interest and investigate the coverage rate of the posterior credible intervals, as well as their average lengths. In particular, we are interested in how well the proposed method performs against alternatives in terms of the aforementioned bias–variance trade-off. Specifically, our analysis focuses on whether we can (1) improve the coverage rate for the true ATT compared to the LFM (DM-LFM; Pang *et al.* 2022) and (2) obtain narrower credible intervals while

maintaining the same false positive rate (measured by coverage rate when true ATT is 0) using the proposed shrinkage priors when there exists an unobserved confounder.

We generate simulated datasets with varying sample sizes according to the following DGP:

$$\begin{aligned} Y_{it} &= 3 + \delta_{it}D_{it} + X_{1it} + 3X_{2it} + \alpha_i + \xi_t + \gamma'_{it}f_t + U_{it} + \varepsilon_{it}, \\ U_{it} &= 0.3 - 0.5D_{it} + 0.1 \sum_{k=1}^5 X_{kit} + 0.1(\alpha_i + \xi_t + \gamma'_{it}f_t) + \nu_{it}. \end{aligned} \quad (32)$$

The DGP (32) assumes that there are three time-varying confounders: X_1 , X_2 , and U that determine the value of potential outcomes. The first two are observed while the last one is unobserved. Their corresponding coefficients are 1, 3, and 1. α_i and ξ_t are additive two-way random effects, and $\gamma'_{it}f_t$ is the latent factor term with two factors. All of them are drawn from independent and identical standard normal distributions, i.e., $\alpha_i \sim N(0, 1)$, $\xi_t \sim N(0, 1)$, $\gamma_{ki} \sim N(0, 1)$, and $f_{kt} \sim N(0, 1)$ for $k \in \{1, 2\}$. We also assume that five covariates are observed: X_1 , X_2 , X_3 , X_4 , and X_5 , and only the first two have non-zero effects on potential outcomes. The observed covariates X_1 to X_5 are also generated from independent and identical standard normal distributions. For the unobserved confounder U , we assume that it is correlated with all observed covariates X_1 to X_5 , treatment indicator D , and additive two-way random effects as well as the latent factor term. ε_{it} and ν_{it} are normal errors with mean 0 and variance 1 and 0.25, respectively. Therefore, X_3 , X_4 , and X_5 implicitly determine the value of potential outcomes through U . Values of sensitivity parameters (each element of λ_x and λ_f) are all set to 0.1. We keep them small so that they are not on the margin of their corresponding prior distributions. Finally, we set the value of the treatment effect δ_{it} . For coverage rate and length of credible intervals, it suffices to study the case where ATT equals 0. For power analysis, we set a variety of values for ATT . We assume that individual treatment effects are heterogeneous; however, for a given value of the ATT , there is no variation across simulated samples.

Existing literature finds that the performance of factor models for causal inference depends on the number of pre-treatment periods as well as the number of control units (Pang *et al.* 2022; Samartsidis *et al.* 2020; Xu 2017). For the design of Monte Carlo studies, we consider several combinations of the number of units ($N \in \{50, 100\}$) and time periods ($T \in \{30, 50, 90\}$). We set the proportion of the treated units to be 20%, whose treatment status switches from 0 to 1 in the post-treatment periods. For the remaining control units, their treatment statuses are always equal to 0. We regard the last ten periods as post-treatment periods. For each combination of N and T , we generate 1,000 simulated data samples using the DGP described above.

For model fitting, we include all observed covariates (X_1 to X_5) in the regression models and regard the true number of factors as known.¹¹ We compare estimation results based on four different model specifications: (1) DM-LFM (w/ U) the ideal, “oracle” case where the unobserved confounder U is observed, (2) DM-LFM (w/o U) the naïve case where we ignore the unobserved U , (3) BSA-LFM (w/ Shrinkage Prior) BSA with the proposed shrinkage prior for β_u , and (4) BSA-LFM (w/ Normal Prior) BSA with normal prior for β_u . For shrinkage prior, we assign a diffuse prior $\xi_{\beta_u}^2 \sim \text{Gamma}(0.001, 0.001)$ to the hyperparameter, while for the normal prior, we assign $\beta_u \sim N(0, 10)$. The results are displayed in Tables 2 and 3.

Table 2 displays a coverage rate of nominal 95% credible intervals. As expected, DM-LFM (w/ U) has a coverage rate close to the nominal value for different number of units. DM-LFM (w/o U), however, produces a coverage rate substantially smaller than the nominal value, and the problem of under-coverage worsens as the number of units increases. BSA-DM-LFM with shrinkage or normal prior has a coverage rate equal to 100%, indicating that they are always conservative. Table 3 reports average lengths of the nominal 95% credible intervals. For both the ideal case and naïve case, the average length of CI

¹¹Readers may refer to Pang *et al.* (2022) for details about how to conduct selection of number of factors with shrinkage priors.

Table 2. Coverage rate of 95% credible intervals.

Model	N = 50			N = 100		
	T = 30	T = 50	T = 90	T = 30	T = 50	T = 90
DM-LFM (w/ U)	0.933	0.933	0.934	0.910	0.943	0.959
DM-LFM (w/o U)	0.566	0.539	0.515	0.317	0.242	0.182
BSA-LFM (w/ Shrinkage Prior)	1.000	1.000	1.000	1.000	1.000	1.000
BSA-LFM (w/ Normal Prior)	1.000	1.000	1.000	1.000	1.000	1.000

Table 3. Average length of 95% credible intervals.

Model	N = 50			N = 100		
	T = 30	T = 50	T = 90	T = 30	T = 50	T = 90
DM-LFM (w/ U)	1.131	1.024	0.960	0.792	0.718	0.674
DM-LFM (w/o U)	1.152	1.045	0.983	0.805	0.733	0.690
BSA-LFM (w/ Shrinkage Prior)	4.351	4.340	4.344	4.296	4.313	4.315
BSA-LFM (w/ Normal Prior)	5.389	5.385	5.382	5.355	5.356	5.353

decreases at the order of the squared root of the sample size. For BSA-LFM, the average length of CI decreases only slightly as the number of units increases.¹²

The conservativeness of the credible intervals and their asymptotic behavior arise from the two sources of uncertainty contributing to these intervals. While sampling variability decreases as sample size increases, the “scenario uncertainty” about the extent of unobserved confounding, which is reflected on the specification of sensitivity parameters, remains unchanged (Gustafson *et al.* 2010). If the substantive prior knowledge about the sensitivity parameter values is available, it is possible to reduce the width of the credible interval by changing priors assigned to the sensitivity parameters, e.g., by reducing the prior variance of λ_d .

We also conduct a power analysis on BSA-LFM with varying sample sizes and magnitudes of treatment effect and compare its performance with different priors with that of DM-LFM. Figure 1 presents the results. We find that the power curves of the ideal and naïve cases look very similar for different number of units, except that it is symmetric about 0 for the ideal case while shifted toward right for the naïve case due to the negative bias. We have two notable findings about BSA-LFM. First, with shrinkage prior, the power curve looks more symmetric about 0 even if the posterior median indicates a negative bias. Second, for a given number of units, as the magnitude of true ATT increases, the percentage of credible intervals excluding 0 becomes larger. Compared to normal prior, BSA-LFM with shrinkage prior may produce a larger percentage of credible intervals excluding 0 for a given combination of the number of units and treatment effect, which is consistent with the evidence that on average BSA-LFM with shrinkage prior produces narrower credible interval than that with normal prior. Combined with the results above, we may conclude that BSA-LFM with shrinkage prior can be more efficient to detect a non-zero ATT compared to that with normal prior while maintaining same false positive rate when some observed covariates have zero effect on potential outcomes.

¹²For BSA-LFM with and without shrinkage priors, the coverage rate of 95% credible intervals is both 100%. Therefore, we conduct additional Monte Carlo studies and compare the coverage rate of 50% credible intervals. The results are reported in Appendix 3.2 of the Supplementary Material.

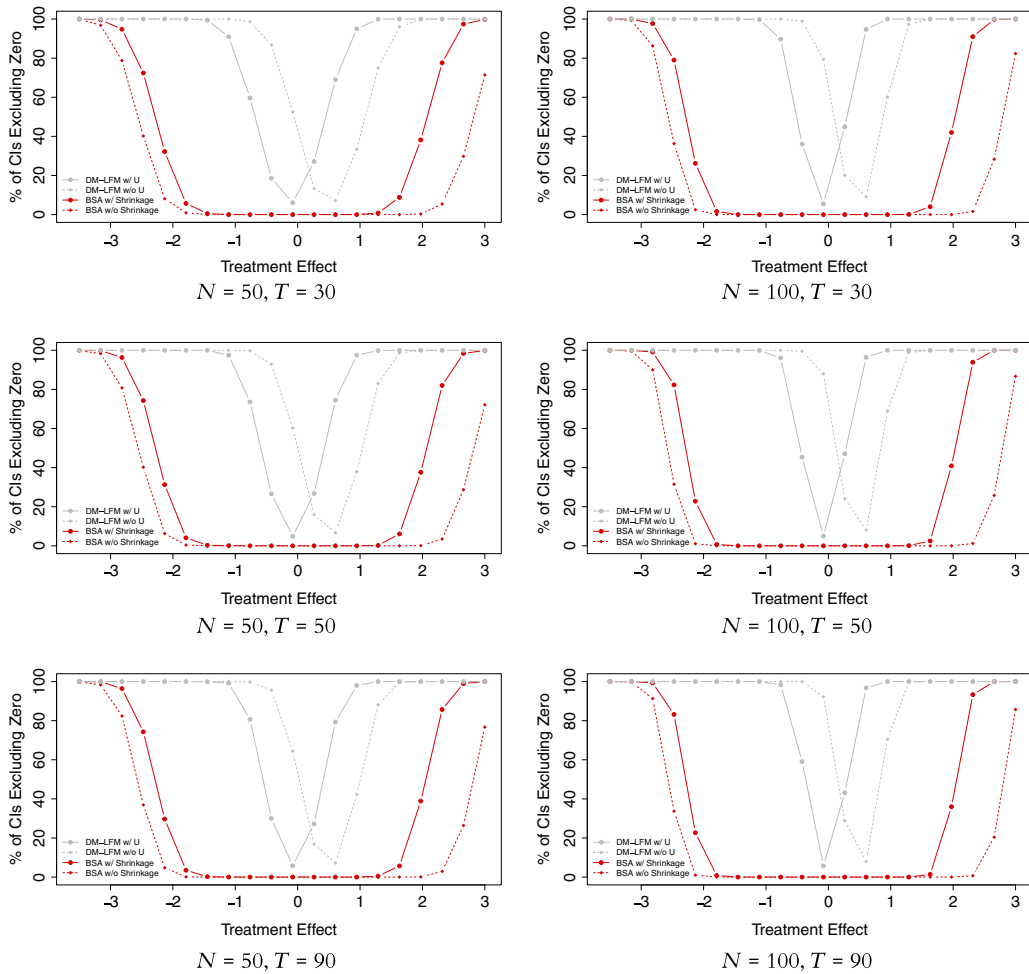


Figure 1. Percentage of nominal 95% credible intervals excluding zero.

5. Empirical Application

To illustrate its applicability, we apply our model to the empirical section of Scheve and Stasavage (2012), which studies the causal effect of war and democracy on taxation of inherited wealth, an important debate in the field of political economy. The authors implemented a difference-in-differences (DID) model specification to analyze an original dataset that includes 19 countries that spans 180 years (1821–2000). They found that mass warfare has a significant positive effect on the top margin rate of inheritance tax, while democracy does not have a significant effect. In this replication, we follow the article and regard mass warfare as a binary treatment indicator which equals 1 if a country engaged in any interstate wars and 2% of the total population was serving in the military. Democracy measured in the expansion of suffrage right is regarded as an observed time-varying confounder. We consider the following model specification:¹³

$$Y_{it} = \delta_{it} D_{i,t} + X'_{it} \beta + \beta_u U_{it} + \alpha_i + \xi_t + \lambda'_i f_t + \varepsilon_{it}, \quad (33)$$

¹³Note that in our model, both the treatment indicator and observed confounders are contemporaneous with the outcome variable. For consistency and comparison with the original findings in Scheve and Stasavage (2012), we lag the treatment and observed confounders for one year.




Figure 2. The effect of war on taxation.

Note: On the left panel, dashed lines indicate posterior 95% credible intervals for ATT estimates based on the naïve model (blue) and the proposed BSA (red).

where X_{it} is a vector of (lagged) time-varying confounders that include: a binary indicator of universal male suffrage, a binary indicator of left executive, and GDP per capita. Lagged outcome variable is also controlled. The original statistical model only adds TWFEs to control for unobserved confounders. In some alternative model specifications, Scheve and Stasavage (2012) control for country-specific time trends to approximate common shocks with unit-heterogeneous effects. Here, we add a latent factor term with ten factors and shrinkage priors for factor selection.

As noted by the original authors, the treatment effect of mass warfare in their baseline model could be biased due to the existence of unobserved time-varying confounders. For example, it is possible that a country's belief about its ability to finance wars by taxing the rich and selection into war are positively correlated, biasing the estimate of the treatment effect in the positive direction. To address this concern, we apply our proposed sensitivity analysis. We regard U_{it} as an unobserved continuous confounder, and investigate how it affects the magnitude and significance of the causal effects. For comparison, we also fit a naïve statistical model which does not include U_{it} but otherwise is exactly the same. The estimation results are displayed in Figure 2.

As expected, the point estimates (posterior means) are very close between our proposed model ("w/ BSA") and the naïve model ("w/o BSA"), both in terms of time-varying treatment effects (right plot, red and blue solid lines) and overall average effects (left plot, solid circles). However, the resulting 95% credible interval is wider when we consider the existence of U_{it} , especially for the first two post-treatment periods (red and blue dotted lines). For average treatment effect, while the 95% credible interval based on the naïve model does not cover 0, the interval based on the proposed model does, indicating that the estimated effect becomes insignificant when we consider the uncertainty due to the existence of an unobserved confounder.

In some cases, researchers may be more interested in hypothetical scenarios corresponding to particular values of the sensitivity parameters. In the taxation example, if the ability to tax the rich to finance war is the unobserved confounder, then β_u should be positive. Posteriors of the sensitivity parameters offer useful, principled guidelines for this type of analysis, by way of providing probabilistic benchmarks for such specific values.¹⁴ To examine the posterior dependence between the treatment effect and the sensitivity parameters, we fix the values of β_u and λ_d at different deciles of their posterior distributions and re-estimate the treatment effect. The results are shown in Figure 3. We find that when

¹⁴Due to limited space, we provide a simulated example on BSA conditional on a specific value of sensitivity parameters β_u or λ_d with conditional posteriors of the other sensitivity parameter in Appendix 3.1 of the Supplementary Material.

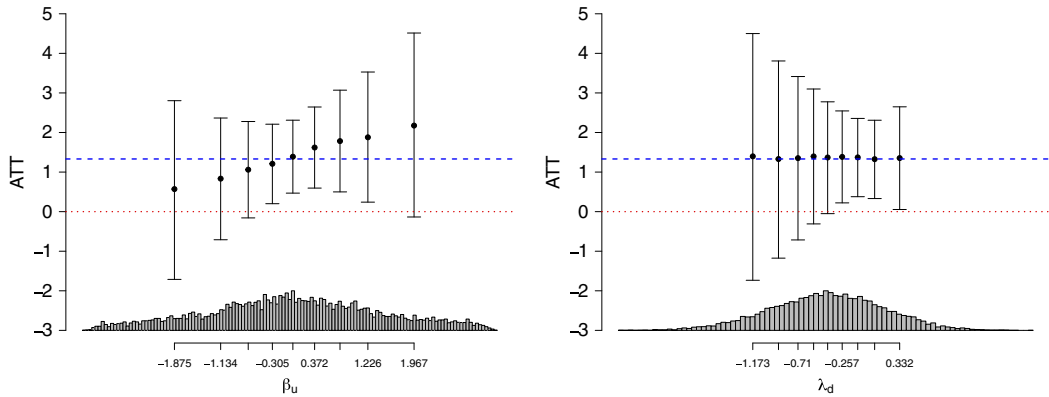


Figure 3. Sensitivity analysis conditional on posterior deciles of β_u and λ_d .

Note: In each plot, the red dotted line indicates zero, while the blue dotted line represents the naïve estimate of the ATT, i.e., the posterior mean of the ATT without Bayesian sensitivity analysis. Each tick on the x -axis represents a decile of the corresponding sensitivity parameter's posterior distribution. The histogram on the bottom represents the posterior marginal density of the sensitivity parameter.




Figure 4. Sensitivity contour of the ATT estimate.

Note: In this plot, the value at (0, 0) corresponds to the ATT estimate without applying BSA. Other values represent the ATT estimates obtained under different combinations of sensitivity parameter values. The ranges of β_u and λ_d shown in the plot correspond to their 95% posterior credible intervals.

the value of β_u or λ_d is closer to 0, the corresponding 95% CI is narrower, indicating less uncertainty when confounding strength of U_{it} is weaker. Indeed, the 95% CI may exclude 0 when either β_u or λ_d is close to 0. Besides, the point estimate of the treatment effect increases as β_u changes from negative to positive, while the value of λ_d barely affects the point estimate.

Finally, researchers may also be interested in assessing the robustness of their findings under extreme or “worst-case” scenarios, as is common in conventional frequentist sensitivity analysis. For instance, one may ask how large the sensitivity parameters must be to explain away the estimated ATT. To illustrate this, we present a contour plot (Figure 4) that shows how the simultaneous variation of β_u and λ_d corresponds to different values of the ATT estimate. In parallel, the posterior distributions of the sensitivity parameters can help researchers assess how plausible a given scenario on the sensitivity contour is in light of the observed data.

6. Conclusion

In this article, we propose a novel BSA method for treatment effect estimation with causal panel data models, especially the linear factor model, based on counterfactual imputation, assessing the sensitivity of the causal effect estimate to the possible existence of unit-specific time-varying unobserved confounding. Through appropriate prior specification, our approach is essentially a Bayesian analog of “benchmarking” sensitivity parameters by observed pre-treatment confounders, making the results of sensitivity analysis more interpretable.

Another advantage of our proposed method is its computation efficiency. When there exists an unobserved time-varying confounder, the augmented linear factor model with sensitivity parameters is by design nonidentified. We utilize the transparent parameterization (Gustafson 2005) to develop an efficient MCMC algorithm. Furthermore, the shrinkage prior assigned to the sensitivity parameter induces substantial indirect learning, making the model more efficient to detect a non-zero treatment effect compared to a diffuse normal prior, as is shown in the Monte Carlo studies.

Our proposed method also has several limitations. First, although the posterior distributions of the sensitivity parameters provide insight into how plausible different degrees of violation of the no unobserved confounding assumption may be, specifying priors is not necessarily easier than defining a range of plausible sensitivity values. As such, researchers must rely on substantive knowledge when selecting priors for the sensitivity parameters. Second, the method relies on parametric assumptions for the outcome model and the unobserved confounder. While the use of linear models reflects a pragmatic modeling choice in practice, such assumptions may limit flexibility in capturing complex DGPs. Finally, although our proposed method is quite flexible to summarize the influence of the unobserved confounder, the Bayesian approach is inherently more computationally intensive than typical frequentist approaches, which may constrain its practical applicability to relatively small datasets. Future research aimed at overcoming these limitations is worth pursuing.

Acknowledgments. We thank Xun Pang, Anton Strezhnev, Ye Wang, Yiqing Xu, Soichiro Yamauchi, members of the Kim Research Group at MIT, and participants at the 2022 APSA Annual Meeting, the 2023 Asian Political Methodology Meeting, the 2023 Annual Meeting of the Society for Political Methodology, and seminars at Peking University School of International Studies and the University of Tokyo. We are especially grateful to Editors-in-Chief Daniel Hopkins and Brandon Stewart and two anonymous reviewers for their incredibly helpful comments. All remaining errors are our own.

Funding Statement. This research is supported by the Political Methodology Lab at MIT.

Data Availability Statement. Replication data and code for this article (Liu and Yamamoto 2025) have been published in the Political Analysis Harvard Dataverse at <https://doi.org/10.7910/DVN/HL7ZYO>.

Supplementary Material. For supplementary material accompanying this paper, please visit <https://doi.org/10.1017/pan.2025.10029>.

References

- Abadie, A., A. Diamond, and J. Hainmueller. 2010. “Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California’s Tobacco Control Program.” *Journal of the American Statistical Association* 105 (490): 493–505.
- Abadie, A., A. Diamond, and J. Hainmueller. 2015. “Comparative Politics and the Synthetic Control Method.” *American Journal of Political Science* 59 (2): 495–510.
- Altonji, J. G., T. E. Elder, and C. R. Taber. 2005. “Selection on Observed and Unobserved Variables: Assessing the Effectiveness of Catholic Schools.” *Journal of Political Economy* 113 (1): 151–184.
- Arkhangelsky, D., and G. W. Imbens. 2022. “Doubly Robust Identification for Causal Panel Data Models.” *The Econometrics Journal* 25 (3): 649–674.
- Athey, S., M. Bayati, N. Doudchenko, G. Imbens, and K. Khosravi. 2021. “Matrix Completion Methods for Causal Panel Data Models.” *Journal of the American Statistical Association* 116 (536): 1716–1730.
- Bai, J. 2009. “Panel Data Models with Interactive Fixed Effects.” *Econometrica* 77 (4): 1229–1279.
- Belmonte, M. A. G., G. Koop, and D. Korobilis. 2014. “Hierarchical Shrinkage in Time-Varying Parameter Models.” *Journal of Forecasting* 33 (1): 80–94.
- Cameron, A. C., and D. L. Miller. 2015. “A Practitioner’s Guide to Cluster-Robust Inference.” *Journal of Human Resources* 50 (2): 317–372.

- Cinelli, C., and C. Hazlett. 2020. "Making Sense of Sensitivity: Extending Omitted Variable Bias." *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 82 (1): 39–67.
- Diegert, P., M. A. Masten, and A. Poirier. 2022. "Assessing Omitted Variable Bias When the Controls are Endogenous." Preprint, [arXiv:2206.02303](https://arxiv.org/abs/2206.02303).
- Frühwirth-Schnatter, S., and H. Wagner. 2010. "Stochastic Model Specification Search for Gaussian and Partial Non-Gaussian State Space Models." *Journal of Econometrics* 154 (1): 85–100.
- Gelfand, A. E., and S. K. Sahu. 1999. "Identifiability, Improper Priors, and Gibbs Sampling for Generalized Linear Models." *Journal of the American Statistical Association* 94 (445): 247–253.
- Gill, J. 2012. "Bayesian Methods in Political Science: Introduction to the Virtual Issue." *Political Analysis* 20 (V3): 1–9.
- Greenland, S. 2003. "The Impact of Prior Distributions for Uncontrolled Confounding and Response Bias: A Case Study of the Relation of Wire Codes and Magnetic Fields to Childhood Leukemia." *Journal of the American Statistical Association* 98 (461): 47–54.
- Gustafson, P., L. C. McCandless, A. R. Levy, and S. Richardson. 2010. "Simplified Bayesian Sensitivity Analysis for Mismeasured and Unobserved Confounders." *Biometrics* 66 (4): 1129–1137.
- Gustafson, P. 2005. "On Model Expansion, Model Contraction, Identifiability and Prior Information: Two Illustrative Scenarios Involving Mismeasured Variables." *Statistical Science* 20 (2): 111–140.
- Gustafson, P., and L. C. McCandless. 2018. "When Is a Sensitivity Parameter Exactly that?" *Statistical Science* 33 (1): 86–95.
- Hollenbach, F. M., J. M. Montgomery, and A. Crespo-Tenorio. 2019. "Bayesian Versus Maximum Likelihood Estimation of Treatment Effects in Bivariate Probit Instrumental Variable Models." *Political Science Research and Methods* 7 (3): 651–659.
- Imbens, G. W. 2003. "Sensitivity to Exogeneity Assumptions in Program Evaluation." *American Economic Review* 93 (2): 126–132.
- Imbens, G. W., and D. B. Rubin. 1997. "Estimating Outcome Distributions for Compliers in Instrumental Variables Models." *The Review of Economic Studies* 64 (4): 555–574.
- Keele, L., and K. M. Quinn. 2017. "Bayesian Sensitivity Analysis for Causal Effects from 2×2 Tables in the Presence of Unmeasured Confounding with Application to Presidential Campaign Visits." *Annals of Applied Statistics* 11 (4): 1974–1997.
- Klinenberg, D. 2022. "Synthetic Control with Time Varying Coefficients a State Space Approach with Bayesian Shrinkage." *Journal of Business & Economic Statistics* 41 (4): 1065–1076.
- Lin, D. Y., B. M. Psaty, and R. A. Kronmal. 1998. "Assessing the Sensitivity of Regression Results to Unmeasured Confounders in Observational Studies." *Biometrics* 54 (3): 948–963.
- Liu, L., and T. Yamamoto. 2025. "Replication Data for: Bayesian Sensitivity Analysis for Unmeasured Confounding in Causal Panel Data Models." Harvard Dataverse, V1. <https://doi.org/10.7910/DVN/HL7ZYO>
- Manski, C. F., and J. V. Pepper. 2018. "How Do Right-to-Carry Laws Affect Crime Rates? Coping with Ambiguity Using Bounded-Variation Assumptions." *Review of Economics and Statistics* 100 (2): 232–244.
- McCandless, L. C., P. Gustafson, and A. Levy. 2007. "Bayesian Sensitivity Analysis for Unmeasured Confounding in Observational Studies." *Statistics in Medicine* 26 (11): 2331–2347.
- McCandless, L. C., P. Gustafson, and A. R. Levy. 2008. "A Sensitivity Analysis Using Information about Measured Confounders Yielded Improved Uncertainty Assessments for Unmeasured Confounding." *Journal of Clinical Epidemiology* 61 (3): 247–255.
- Mebane Jr, W. R., and P. Poast. 2013. "Causal Inference without Ignorability: Identification with Nonrandom Assignment and Missing Treatment Data." *Political Analysis* 21 (2): 233–251.
- Oster, E. 2019. "Unobservable Selection and Coefficient Stability: Theory and Evidence." *Journal of Business & Economic Statistics* 37 (2): 187–204.
- Pang, X., L. Liu, and Y. Xu. 2022. "A Bayesian Alternative to Synthetic Control for Comparative Case Studies." *Political Analysis* 30 (2): 269–288.
- Park, T., and G. Casella. 2008. "The Bayesian Lasso." *Journal of the American Statistical Association* 103 (482): 681–686.
- Rambachan, A., and J. Roth. 2023. "A more Credible Approach to Parallel Trends." *Review of Economic Studies* 90 (5): 2555–2591.
- Rosenbaum, P. R., and P. R. Rosenbaum. 2002. *Observational Studies*. New York, NY: Springer New York.
- Rubin, D., X. Wang, L. Yin, and E. Zell. 2008. "Bayesian Causal Inference: Approaches to Estimating the Effect of Treating Hospital Type on Cancer Survival in Sweden Using Principal Stratification." In *Handbook of Applied Bayesian Analysis*, edited by Anthony O'Hagan and Mike West, 679–708, Oxford University Press.
- Rubin, D. B. 1974. "Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies." *Journal of Educational Psychology* 66 (5): 688.
- Rubin, D. B. 1976. "Inference and Missing Data." *Biometrika* 63 (3): 581–592.
- Rubin, D. B. 1978. "Bayesian Inference for Causal Effects: The Role of Randomization." *The Annals of Statistics* 6 (1): 34–58.
- Samartsidis, P., S. R. Seaman, S. Montagna, A. Charlett, M. Hickman, and D. De Angelis. 2020. "A Bayesian Multivariate Factor Analysis Model for Evaluating an Intervention by Using Observational Time Series Data on Multiple Outcomes." *Journal of the Royal Statistical Society Series A: Statistics in Society* 183 (4): 1437–1459.
- Scheve, K., and D. Stasavage. 2012. "Democracy, War, and Wealth: Lessons from Two Centuries of Inheritance Taxation." *American Political Science Review* 106 (1): 81–102.

- VanderWeele, T. J. 2008. "Sensitivity Analysis: Distributional Assumptions and Confounding Assumptions." *Biometrics* 64 (2): 645–649.
- VanderWeele, T. J., and P. Ding. 2017. "Sensitivity Analysis in Observational Research: Introducing the e-Value." *Annals of Internal Medicine* 167 (4): 268–274.
- Xu, Y. 2017. "Generalized Synthetic Control Method: Causal Inference with Interactive Fixed Effects Models." *Political Analysis* 25 (1): 57–76.
- Ye, T., L. Keele, R. Hasegawa, and D. S. Small. 2024. "A Negative Correlation Strategy for Bracketing in Difference-in-Differences." *Journal of the American Statistical Association* 119 (547): 2256–2268.