

Degree of Interference: A General Framework For Causal Inference Under Interference

Yuki Ohnishi

*Department of Biostatistics
Yale School of Public Health
New Haven, Connecticut 06510, USA*

YUKI.OHNISHI@YALE.EDU

Bikram Karmakar

*Department of Statistics
University of Wisconsin-Madison
Madison, Wisconsin 53706, USA*

BKARMAKAR@WISC.EDU

Arman Sabbaghi

*Department of Statistics
Purdue University
West Lafayette, Indiana 47907, USA*

SABBAGHI@PURDUE.EDU

Editor: Joris Mooij

Abstract

One core assumption typically adopted for valid causal inference is that of no interference between experimental units, i.e., the outcome of an experimental unit is unaffected by the treatments assigned to other experimental units. This assumption can be violated in real-life experiments, which significantly complicates the task of causal inference. As the number of potential outcomes increases, it becomes challenging to disentangle direct treatment effects from “spillover” effects. Current methodologies are lacking, as they cannot handle arbitrary, unknown interference structures to permit inference on causal estimands. We present a general framework to address the limitations of existing approaches. Our framework is based on the new concept of the “degree of interference” (DoI). The DoI is a unit-level latent variable that captures the latent structure of interference. We also develop a data augmentation algorithm that adopts a blocked Gibbs sampler and Bayesian nonparametric methodology to perform inferences on the estimands under our framework. We illustrate the DoI concept and properties of our Bayesian methodology via extensive simulation studies and an analysis of a randomized experiment investigating the impact of a cash transfer program for which interference is a critical concern. Ultimately, our framework enables us to infer causal effects without strong structural assumptions on interference.

Keywords: Rubin Causal Model, Interference, Spillover effect, Bayesian Nonparametric, Dependent Dirichlet Process mixture

1. Introduction

The two components of the Stable Unit Treatment Value Assumption (SUTVA) (Rubin, 1974) are fundamental for causal inference for a treatment effect. The first component of SUTVA is that there are well-defined treatments for each experimental unit that yield well-defined outcomes, and the second is that no interference exists between units. The first

assumption is typically satisfied by the design of an experiment, whereas the second relates to the complicated phenomenon of interference and may not be true in certain experiments. Interference is said to arise when a particular experimental unit’s outcome is a function of both their assigned treatment as well as the treatments assigned to other experimental units. For example, interference can naturally arise in an experiment on the efficacy of a vaccine for a contagious disease, as one subject’s health outcome can depend on whether their close contacts received the vaccine. The existence of interference complicates causal inferences because one must disentangle direct treatment effects from the spillover effects of the treatments assigned to other units.

This paper proposes a Bayesian semiparametric methodology for causal inference under interference. It is designed for situations where only partial, though potentially multiple, insights are available regarding the interference. The existing methods for causal inference under interference, as we discuss in §2, either assume the interference structure or assess the robustness of inference when it is misspecified. In the latter case, the inference is generally sensitive to misspecification. In our proposed approach, a unit’s interference structure is modeled as a latent function of the unit’s features and other units’ treatments. This latent function identifies the direct and spillover effect. We develop a Bayesian nonparametric prior for this latent function and a fast data augmentation algorithm that adopts a blocked Gibbs sampler to draw posterior inferences for the causal effects. The code implementing our method is available in the online supplement.

We call the unit-level latent functions “degree of interference.” One of our main contributions is to formalize the notion of the degree of interference. Using this notion, we develop a context-aware modeling strategy that flexibly reflects our knowledge about the interference structure. In the next section, we review existing methods and highlight our contributions in their context.

2. Current Literature on Causal Inference Under Interference

The first foundational work towards causal inference under interference was by Hudgens and Halloran (2008), who proposed a design-based approach to identify and estimate direct and spillover effects in a two-stage hierarchical experimental design. The first stage of the design randomly assigns clusters of units to one of two treatment assignment strategies, and then the second stage assigns treatments within the clusters based on the selected strategies. Their methods assume that the interference is (1) partial, i.e., operates only within clusters, and (2) stratified, i.e., is specified via the proportion of treated units within clusters. Several important works have built upon the above assumptions (Liu and Hudgens, 2014; Kang and Imbens, 2016; Sussman and Airolidi, 2017; Basse and Feller, 2018; Imai et al., 2021; Ohnishi and Sabbaghi, 2022). However, in practice, the validity of the assumptions of partial interference, which requires multiple independent clusters, and stratified interference, which is typically a simplification for the unknown interference structure, may be questionable (Hudgens and Halloran, 2008, p. 836). When these assumptions fail, the above methodologies provide biased inferences (Sävje, 2023). In contrast, we utilize a Bayesian nonparametric prior for the “degree of interference”, specifically, the Dependent Dirichlet process mixture prior to flexibly model the unknown interference structure. These priors can data-adaptively separate

units into several clusters based on how they interfere with one another, thus accommodating a stratified nature of the interference structure.

Aronow and Samii (2017) introduced the idea of exposure mapping to move beyond the partial and stratified interference assumptions. Under this framework, a mapping is specified that relates the vector of treatment assignments for the other experimental units to a finite set of exposures. Then, an experimental unit’s outcome is specified as a function of its own treatment label and this exposure. In other words, an exposure map is supposed to capture the full nature of interference of a unit from all other units. Once an exposure map is well-specified, causal effects are defined in terms of comparisons of outcomes under different exposures. A line of research exists that uses exposure maps to provide causal inference under interference. Notably, Toulis and Kao (2013) developed randomization-based and Bayesian model-based inferential approaches for “peer influence effects” defined by fixing the specific number of neighbors who receive treatment. Basse et al. (2019) developed a framework for constructing valid randomization tests for general exposure contrasts. Other research works have addressed interference by specifying their interference mapping based on peers’ treatment status and attributes of a given network (Li et al., 2019; Zigler and Papadogeorgou, 2021; Forastiere et al., 2021; Papadogeorgou and Samanta, 2023).

As Basse and Airoidi (2018b) mention, causal inference under interference is impossible without assumptions about the interference structure. Thus, it is natural to use exposure maps specifying the structure. However, current causal inferential frameworks and methodologies based on exposure maps are limited in several respects. First, they require that the interference structure be fully characterized as a finite set of exposures *a priori* so that each exposure corresponds to a unique potential outcome. In practice, satisfying this requirement may be unrealistic because the interference structure is usually unknown, or partially known, and can be very complicated to understand. Second, as observed by Sävje (2023), a misspecified exposure mapping can lead to imprecise variance estimation and, thus, incorrect inference. The domain knowledge of the interference structure can mitigate this problem. Leung (2022) proposed a flexible model that permits treatments that are assigned to units “farther” from a particular unit to have non-vanishing, but smaller, effects on the unit’s response. Our proposal for the “degree of interference” concept generalizes the exposure map. Unlike other methods, it does not require precise knowledge of the interference structure. We specify a flexible semi-parametric model for the interference structure and define the potential outcomes as expected values of pseudo-potential outcomes that marginalize with respect to the uncertainty in the interference model. This approach allows for knowledge-aware specifications about interference and a fast Bayesian algorithm to infer the counterfactual outcomes.

Third, the exposure mapping framework simultaneously defines both the interference structure that determines a single unique potential outcome and exposure effects that determine causal quantities. However, since exposure mapping is typically problem-specific, exposure effects generally do not hold any meaning outside of the problem instance, even for the same treatment-outcome pair. Recent literature building upon the exposure mapping framework aims to unbiasedly estimate the exposure effects (e.g., Sävje (2023); Leung (2022)) or to construct the testing procedure for the correctness of the exposure map Hoshino and Yanagi (2023). In contrast, we define new causal estimands that are not bonded with the interference structure. Our estimands are generalizations of those suggested by Sävje

et al. (2021) as natural causal estimands for direct and spillover effects under interference. Sävje et al. (2021) demonstrated that the Horvitz-Thompson (HT) and Häjek estimators are consistent and derived their asymptotic distributions. However, these estimators lack theoretical guarantees when interference scales proportionally with the sample size, as outlined in Assumption 2 of Sävje et al. (2021). Our framework circumvents such restrictions by allowing us to model the interference structure and serves as a promising alternative in such a complex interference setting. Finally, our approach allows for continuous treatment factors and factors with many discrete levels. This flexibility is particularly advantageous in many practical settings, where such variations in treatment factors are commonplace.

The rest of the paper is organized as follows. We proceed in §3 to review the Rubin Causal Model and propose new spillover effect estimands under interference. §4 contains our new framework based on the “degree of interference” concept and our corresponding data augmentation algorithm for inferring causal estimands under interference. The frequentist properties of our method are demonstrated via extensive simulation studies in §5. In §6 we demonstrate our framework using case studies on cash transfer programs in Colombia and an experiment on stereolithography. §7 concludes the paper.

3. Background

3.1 The Rubin Causal Model and Interference Structures

Adopting the Rubin Causal Model (Rubin, 1974), consider N experimental units, indexed by $i = 1, \dots, N$, that correspond to physical objects at a particular point in time. Each experimental unit i has an observed set of pre-treatment covariates $\mathbf{X}_i \in \mathcal{X}$, with \mathcal{X} denoting the covariate space; typically, $\mathcal{X} = \mathbb{R}^d$ for d covariates. Treatment assignments are contained in the N -dimensional vector $\mathbf{Z} = (Z_1, \dots, Z_N)^\top$, where $Z_i \in \Omega$, the treatment assignment for unit i , can be binary, categorical or continuous, and Ω is the set of treatment levels. So, the vector $\mathbf{Z} \in \Omega^N$. For example, $\Omega = \mathbb{R}$ denotes a continuous treatment factor with “0” denoting the specified “control” level, and $\Omega = \{0, 1, \dots, K\}$ denotes a categorical treatment with $K + 1$ levels and with “0” again denoting control. Let $\mathbf{Z}_{-i} = (Z_1, \dots, Z_{i-1}, Z_{i+1}, \dots, Z_N)^\top$ be the sub-vector of \mathbf{Z} with the i th element removed; so, $\mathbf{Z}_{-i} \in \Omega^{N-1}$.

Next, consider the potential outcomes for unit i . To facilitate the exposition and indicate how, in the general setting of interference, unit i ’s potential outcomes depend on multiple entries in \mathbf{Z} , we write the potential outcomes for unit i as $Y_i(\mathbf{Z}) \equiv Y_i(Z_i, \mathbf{Z}_{-i})$. We let $\mathbf{Y} = \{Y_i(\mathbf{z}) : i = 1, \dots, N, \mathbf{z} \in \Omega^N\}$ denote the set of all potential outcomes. The observed outcomes are contained in vector $\mathbf{Y}^{\text{obs}} = (Y_1^{\text{obs}}, \dots, Y_N^{\text{obs}})^\top$ where $Y_i^{\text{obs}} = \sum_{\mathbf{z} \in \Omega^N} Y_i(\mathbf{z}) \mathbb{I}(\mathbf{Z} = \mathbf{z})$ for each i . Thus, only one outcome is observable to us among the many potential outcomes of unit i .

The probability distribution of \mathbf{Z} given covariates $\mathbf{X}_1, \dots, \mathbf{X}_N$ is known to us. We assume without loss of generality that the support of the probability distribution is equal to Ω^N .

Since the experimental units interfere, as is common in the literature, we assume throughout that they belong to a single, known network that governs their interference structure. Let \mathbf{A} denote the $N \times N$ matrix that captures how the units are related to one another. The specification of \mathbf{A} can vary depending on how one assumes interference between units to occur. For example, to describe the interference structure based just on unit adjacencies, \mathbf{A} can be defined simply via the adjacency matrix of the units. If instead one wishes to use the

shortest path between units to describe their interference structure, then \mathbf{A} is defined as the distance matrix of the units, with element (i, j) of \mathbf{A} being the shortest distance between units i and j . Later, in our definition of the degree of interference, \mathbf{A} becomes an argument of its functional form.

3.2 Causal Estimands

Throughout, we consider the finite-population perspective, where we view the N units are given, and an external, known randomization process assigns treatment to these units. All potential outcomes are defined as fixed but mostly missing for the unrealized treatment assignment.

We define two types of causal estimands—the first type is assignment-conditional effects, and the second is expected effects marginalized over a treatment assignment. Assignment-conditional causal estimands are defined by conditioning on a specific treatment assignment vector. We consider two effects for this type of estimand: average treatment and spillover effects.

Following Sävje et al. (2021), who used similar estimands in the case of a binary treatment, the unit-level assignment-conditional treatment effect for unit i is the contrast between its two potential outcomes corresponding to $Z_i = 0$ and $Z_i = z$ when the treatment assignments \mathbf{Z}_{-i} for the other experimental units are held fixed at $\mathbf{z}_{N-1} \in \Omega^{N-1}$. Denote this effect by $\tau_i^{\text{tr}}(Z_i = z, \mathbf{Z}_{-i} = \mathbf{z}_{N-1}) = Y_i(Z_i = z, \mathbf{Z}_{-i} = \mathbf{z}_{N-1}) - Y_i(Z_i = 0, \mathbf{Z}_{-i} = \mathbf{z}_{N-1})$. The average of the $\tau_i^{\text{tr}}(Z_i = z, \mathbf{Z}_{-i} = \mathbf{z}_{N-1})$ across the N experimental units corresponds to a finite-population average treatment effect.

Definition 1 (Conditional Average Treatment Effect) *An assignment-conditional average treatment effect (A-CATE) comparing a pair of treatment, a $z \in \Omega$ and vector $\mathbf{z}' \in \Omega^N$ is*

$$\tau_{A\text{-CATE}}(z, \mathbf{z}') = \frac{1}{N} \sum_{i=1}^N \tau_i^{\text{tr}}(Z_i = z, \mathbf{Z}_{-i} = \mathbf{z}'_{-i}),$$

where \mathbf{z}'_{-i} is an $(N - 1)$ -dimensional sub-vector of \mathbf{z}' with element i removed.

A-CATE is a natural extension of the estimand of Sävje et al. (2021) for the cases of categorical or continuous treatments. This effect quantifies the average change in the outcome as one unit's treatment changes from 0 to z when the treatments for all other units are fixed at \mathbf{z}'_{-i} .

We define our unit-level spillover effect that is inspired by the A-CATE. Specifically, for unit i with treatment z and for two treatment assignment vectors $\mathbf{z}'_{N-1}, \mathbf{z}^*_{N-1} \in \Omega^{N-1}$ for the other units, the unit-level spillover effect is defined as $\tau_i^{\text{sp}}(z, \mathbf{z}'_{N-1}, \mathbf{z}^*_{N-1}) = Y_i(Z_i = z, \mathbf{Z}_{-i} = \mathbf{z}'_{N-1}) - Y_i(Z_i = z, \mathbf{Z}_{-i} = \mathbf{z}^*_{N-1})$.

Definition 2 *An assignment-conditional average spillover effect (A-CASE) under a treatment $z \in \Omega$ for unit i and treatment vectors $\mathbf{z}'_N, \mathbf{z}^*_N \in \Omega^N$*

$$\tau_{A\text{-CASE}}(z, \mathbf{z}'_N, \mathbf{z}^*_N) = \frac{1}{N} \sum_{i=1}^N \tau_i^{\text{sp}}(z, \mathbf{z}'_{-i}, \mathbf{z}^*_{-i}),$$

where \mathbf{z}'_{-i} and \mathbf{z}^*_{-i} are $(N - 1)$ -dimensional sub-vectors of \mathbf{z}' and \mathbf{z}^* with the i th elements removed.

Our definition of the A-CASE corresponds to the average comparison of the units’ outcomes for a single treatment under two different cases, corresponding to the two different treatment vectors considered for the other units. Of special interest is the case with $z = 0$, $\mathbf{z}_N^* = (0, \dots, 0)^\top$, and \mathbf{z}' set at some other vector, where A-CASE captures a pure spillover effect.

Our second type of causal estimands is a set of three expected effects, which marginalize the assignment-conditional causal estimands over a probability distribution on the assignment vectors, i.e., the treatment assignment mechanism. Expected effects are useful in practice because the number of assignment-conditional effects grows rapidly as a function of N , whereas the number of expected effects is generally smaller.

Definition 3 (Expected Average Treatment and Spillover) *Let $\mathbb{E}_\pi(h(\mathbf{Z}))$ denote the expectation of a function $h(\mathbf{Z})$ of the assignment vector over the known probability mass/density function $\pi(\mathbf{z}) = p(\mathbf{z} \mid \mathbf{X}_1, \dots, \mathbf{X}_N)$. Then for each assignment-conditional average effect, we define the corresponding expected average effect based on the expectation of the assignment-conditional average effect over the assignment mechanism: $\tau_{E-ATE}(z; \pi) = \mathbb{E}_\pi\{\tau_{A-CATE}(z, \mathbf{Z})\}$ and $\tau_{E-ASE}(z; \pi) = \mathbb{E}_\pi\{\tau_{A-CASE}(z, \mathbf{Z}, \mathbf{0})\}$.*

Our version of this estimand, τ_{E-ATE} , is a natural extension of the estimand of Sävje et al. (2021) for categorical treatment factors with multiple levels and continuous factors. Our definitions of τ_{E-ASE} is motivated in a similar manner as that of τ_{E-ATE} . As noted by Sävje et al. (2021), the E-ATE—similarly, E-ASE—captures the expected average effect of changing the treatment of a single unit in the “current” experiment.

4. The Degree of Interference Framework and Bayesian Methodology

4.1 Definition and Assumptions for the Degree of Interference

The Degree of Interference (DoI) is defined in terms of the characteristics of a unit and the treatments for all other units. Our DoI framework addresses a limitation of exposure maps that, while an exposure map needs to be pre-specified, the DoI is an unknown latent function that represents the underlying interference structure between units, which is inferred from the data.

We let characteristics $\mathbf{T}_i \in \mathcal{T}$ for unit i contain, its covariates, \mathbf{X}_i , covariates for units other than i , \mathbf{X}_{-i} , and the adjacent network information of the other units. In all that follows, we let $\mathcal{T} = \mathcal{X} \times \mathcal{X}^{N-1} \times \mathcal{A}$, where \mathcal{A} denotes the set of matrices of networks, i.e., set of \mathbf{A} s.

Definition 4 *The DoI for unit i is a random function $G_i : \mathcal{T} \times \Omega^{N-1} \rightarrow \mathbb{R}$ that maps the unit’s characteristics, residing in \mathcal{T} , and treatments for all other units to a number.*

In a simple scenario with no individual or subgroup covariates, each unit’s “characteristics” only encode the network or grouping structure that dictates which other units can affect it. In this setting, G_i is a stochastic process indexed by an element $(\mathbf{A}, \mathbf{Z}_{-i}) \in \mathcal{A} \times \Omega^{N-1}$. For example, if it is reasonable to assume that the proportion of treated units in a neighborhood is the primary factor determining the spillover effect (i.e., the stratified interference assumption; (Hudgens and Halloran, 2008)), we may assume $G_i \sim N\left(\beta_0 + \beta_1 \frac{\sum_{j \in \mathcal{N}_i} Z_j A_{ij}}{|\mathcal{N}_i|}, \sigma^2\right)$, where \mathcal{N}_i

is the set of neighboring units for unit i . The core of the DoI framework allows the function G_i —and the corresponding parameters β_0 , β_1 , and σ —to be unknown, latent quantities inferred from the data rather than a fixed, pre-specified exposure map. For instance, one might be confident that the “proportion of treated neighbors” is relevant for interference but remain unsure whether it is indeed the correct functional form. In practice, the DoI framework can learn such dependencies from data, even if there is a more complex network relationship (e.g., higher-order neighbors). By contrast, the *exposure map* approach would require the analyst to explicitly fix a certain functional form (like a fraction of treated neighbors), potentially overlooking more nuanced patterns. Detailed discussions are provided in Section 4.4.

The DoI satisfies the following assumption which says that unit i ’s potential outcomes are identified from its own treatment and DoI.

Assumption 1 *For all $i \in \{1, \dots, N\}$, $\mathbf{T}_i \in \mathcal{T}$, $z \in \Omega$, and $\mathbf{z}_{N-1}, \mathbf{z}'_{N-1} \in \Omega^{N-1}$ in which $G_i(\mathbf{T}_i, \mathbf{z}_{N-1}) \stackrel{D}{=} G_i(\mathbf{T}_i, \mathbf{z}'_{N-1})$, where $\stackrel{D}{=}$ denotes equivalence in the marginal prior distributions of $G_i(\mathbf{T}_i, \mathbf{z}_{N-1})$ and $G_i(\mathbf{T}_i, \mathbf{z}'_{N-1})$, we have $Y_i(z, \mathbf{z}_{N-1}) = Y_i(z, \mathbf{z}'_{N-1})$.*

This assumption ultimately corresponds to equivalence in the potential outcomes according to DoIs that have the same marginal prior distributions. More formally, for all $i \in \{1, \dots, N\}$, $z \in \Omega$, $\mathbf{z}_{N-1} \in \Omega^{N-1}$, and $\mathbf{T}_i \in \mathcal{T}$, the equivalencies of the potential outcomes for unit i with $Y_i(z, \mathbf{z}_{N-1})$ are based on z and the marginal prior distribution on $G_i \equiv G(\mathbf{T}_i, \mathbf{z}_{N-1})$. We represent its conditional cumulative distribution function $F_{G(\mathbf{T}_i, \mathbf{z}_{N-1})}$.

Remark 5 (Randomness of DoI) *Assumption 1 is critical, as it implies that the DoI is effectively a latent representation of the interference structure for unit i with characteristics T_i when other units are assigned to \mathbf{z}_{N-1} . The distributional equivalence allows for additional flexibility of the DoI; it indicates the unknown uncertainty about the interference structure. Under the finite-population perspective, the treatment assignment is usually the only source of randomness for observables. Therefore, the existing design-based approach under the finite-population perspective (e.g., partial interference, stratified interference and exposure maps) inevitably defines the interference structure a priori even when there is limited knowledge about it. Under our framework, by contrast, the interference structure is an “unobserved” quantity, representing uncertainty for analysts. We incorporate the knowledge about interference and the associated uncertainty into our model-based inference developed in the following sections. Our DoI framework adopts a Bayesian perspective by introducing a random function G_i to capture how other units’ treatments and network structures affect unit i . Although $Y_i(\mathbf{z})$ is ultimately fixed, from the analyst’s point of view it remains unknown and is therefore modeled through a prior distribution over G_i . This approach allows us to incorporate uncertainty about the interference mechanism and update our beliefs based on observed data. In other words, the DoI framework posits a prior belief about how cross-unit effects flow through the latent function G_i , and subsequent inferences reflect how the data inform and refine our understanding of this interference channel. This uncertainty should be distinguished from the ones on the underlying network between units (Egami, 2021). We place priors on functional spaces, which are flexible enough to recover the latent DoIs that satisfy Assumption 1.*

In addition to Assumption 1, we consider the following assumption about the relation between the DoI function and potential outcomes.

Assumption 2 For each unit i , there exists auxiliary potential outcomes $\tilde{Y}_i(z, g)$, where $\tilde{Y}_i(z, G_i)$ and $\tilde{Y}_i(z, G'_i)$ have the same distribution when G_i and G'_i have the same marginal prior distributions. In contrast to the original potential outcome, the auxiliary potential outcomes have randomness induced by G_i , and the original potential outcome is:

$$Y_i(z, \mathbf{z}_{N-1}) = \mathbb{E}_{G(\mathbf{T}_i, \mathbf{z}_{N-1})} \{\tilde{Y}_i(z, G_i)\} = \int_{\mathbb{R}} \tilde{Y}_i(z, g) dF_{G_i(\mathbf{T}_i, \mathbf{z}_{N-1})}(g), \quad (1)$$

Under this assumption, our DoI framework naturally incorporates the uncertainty associated with interference dynamics into the process of inferring potential outcomes. In contrast, the exposure mapping framework assumes exact specification of the interference structure and thus is limiting. Further, an exact specification of the interference is often not possible in practice. Also, our framework under Assumptions 1 and 2 generalizes the exposure mapping framework since our framework does not assume a particular functional form for $F_{G_i(\mathbf{T}_i, \mathbf{z}_{N-1})}$ beforehand, as the exposure mapping does.

Equation (1) says that the uncertainty about G_i is incorporated by marginalizing the auxiliary potential outcomes over $G(\mathbf{T}_i, \mathbf{z}_{N-1})$ for the original potential outcomes with $\mathbf{Z}_{-i} = \mathbf{z}_{N-1}$. Also, it implies an inferential approach for $Y_i(z, \mathbf{z}_{N-1})$ via the Monte Carlo simulation: $Y_i(z, \mathbf{z}_{N-1}) \approx \sum_{m=1}^M \tilde{Y}_i(z, g_m)/M$ where $g_m \sim F_{G_i(\mathbf{T}_i, \mathbf{z}_{N-1})}$ for $m = 1, \dots, M$. The Monte Carlo procedure is introduced in §4.5. Finally, the unit-level assignment-conditional treatment effect is

$$\tau_i^{\text{tr}}(Z_i = z, \mathbf{Z}_{-i} = \mathbf{z}_{N-1}) = \int_{\mathbb{R}} \{\tilde{Y}_i(z, g) - \tilde{Y}_i(0, g)\} dF_{G_i(\mathbf{T}_i, \mathbf{z}_{N-1})}(g). \quad (2)$$

Unit-level assignment conditional spillover effects can also be expressed in terms of the distribution of the DoI. Equation (2) further clarifies the conceptual role of the DoI that it is effectively a low-dimensional compression of interference, and the direct effect is defined as the comparison of two auxiliary potential outcomes with different treatment assignments under the same DoI.

DoI, along with Assumptions 1 and 2, generalizes exposure mappings (Aronow and Samii, 2017). In particular, if $F_{G_i(\mathbf{T}_i, \mathbf{z}_{N-1})}$ is a prespecified, known function that maps to a finite set then the DoI is analogous to exposure mappings. For example, the exposure mapping for stratified interference (Hudgens and Halloran, 2008) would correspond to prespecifying $F_{G_i(\mathbf{T}_i, \mathbf{z}_{N-1})}$ as an indicator function of the number of treated units in the neighborhood of unit i : $F_{G_i(\mathbf{T}_i, \mathbf{z}_{N-1})}(g) = \mathbb{1}(\mathbf{z}_{N-1}^\top A_{-i,i} = g)$ for $g \in \{1, \dots, N-1\}$, where $A_{-i,i}$ is the i -th column vector of the adjacency matrix A with the i -th row removed. However, in contrast to existing approaches based on exposure maps, the DoI is not prespecified but is instead inferred from data via our Bayesian methodology in §4.2 and §4.3. Inferring the DoI is advantageous in principle, as imposing a misspecified interference structure and failing to modify it could yield biased inferences (Sävje, 2023).

Finally, the DoI framework allows causal inferences for experiments that satisfy the following new version of the latent unconfoundedness assumption.

Assumption 3 A treatment assignment is unconfounded under the DoI framework if the probability mass/density function in the treatment assignment mechanism does not depend on \mathbf{Y} or the $G_i(\mathbf{T}_i, \mathbf{z}_{-i})$ conditional on the \mathbf{T}_i , i.e., $p(\mathbf{z}|\mathbf{T}_1, \dots, \mathbf{T}_N, \mathbf{Y}, G_1(\mathbf{T}_1, \mathbf{z}'_{-1}), \dots, G_N(\mathbf{T}_N, \mathbf{z}'_{-N})) = p(\mathbf{z}|\mathbf{T}_1, \dots, \mathbf{T}_N)$, for all $\mathbf{T}_1, \dots, \mathbf{T}_N \in \mathcal{T}$, $\mathbf{z}, \mathbf{z}' \in \Omega^N$ and $G_1(\mathbf{T}_1, \mathbf{z}_{-1}), \dots, G_N(\mathbf{T}_N, \mathbf{z}_{-N})$.

While this assumption is less restrictive than a randomized treatment assignment, this assumption may not be valid in observational studies. In our case studies in §6, Assumption 3 is valid due to the design of the experiment.

4.2 Overview of the Bayesian Methodology

Bayesian inference is one of the modes of inference within the potential outcome framework, where we consider the missing potential outcomes as unobserved random variables, no different from unknown parameters. Specifically, Bayesian inference involves specifying a model for all random variables, including the potential outcomes, treatment, and covariates. Based on this model, one can infer causal estimates from the joint posterior distributions of the parameters and the missing potential outcomes, conditional on the observed data (Rubin, 1978).

The Bayesian approach not only provides a principled framework for analyzing causal studies but also provides a refined map of identifiability, clarifying what can be learned when causal estimates are not fully identifiable but instead weakly identifiable (i.e., when the likelihood functions of parameters and causal estimates have substantial regions of flatness) (Imbens and Rubin, 1997). In particular, issues of identification differ from those in the frequentist paradigm because if prior distributions are proper, then posterior distributions are always proper. Weak identifiability is reflected in the flatness of the posterior distribution and can be evaluated quantitatively (Gustafson, 2009). Li et al. (2022) gives further comprehensive reviews of Bayesian causal inference.

Our Bayesian methodology involves multiple imputations of the DoI for missing potential outcomes to derive the posterior distributions of the finite-population causal estimands. To formally describe this, let τ denote one of the causal estimands from §3.2, \mathbf{T} the $N \times P$ matrix of characteristics for all units (where the columns correspond to the concatenation of covariates and adjacent network information), \mathbf{Z} the $N \times 1$ vector of assigned treatments, \mathbf{Y}^{obs} the vector of observed outcomes, and \mathbf{Y}^{mis} the vector of all missing potential outcomes. The \mathbf{Y}^{mis} is determined according to the treatment assignment, with each missing potential outcome for unit i being either $Y_i(z'_i, \mathbf{z}_{-i})$ or $Y_i(z, \mathbf{z}'_{-i})$ with z'_i denoting an alternate treatment for unit i and \mathbf{z}'_{-i} denoting an alternate treatment vector for all units excluding unit i . In these cases, the DoI for unit i would be either $G_i(\mathbf{T}_i, \mathbf{z}_{-i})$ or $G_i(\mathbf{T}_i, \mathbf{z}'_{-i})$ respectively. We let $\mathbf{G}^{\text{o}} = (G_1(\mathbf{T}_1, \mathbf{Z}_{-1}), \dots, G_N(\mathbf{T}_N, \mathbf{Z}_{-N}))^\top$ and $\mathbf{G}^{\text{u}}(\mathbf{z}') = (G_1(\mathbf{T}_1, \mathbf{z}'_{-1}), \dots, G_N(\mathbf{T}_N, \mathbf{z}'_{-N}))^\top$ contain the DoIs for the realized and an unrealized treatment assignment vector \mathbf{z}' , respectively. None of the entries in \mathbf{G}^{o} or the $\mathbf{G}^{\text{u}}(\mathbf{z}')$ are observable. In accordance with the matrix \mathbf{Y}^{mis} , we generally define the matrix \mathbf{G}^{u} as the column-wise concatenation of all the $\mathbf{G}^{\text{u}}(\mathbf{z}')$ for treatment assignment vectors \mathbf{z}' that were not realized.

Standard types of point estimates of estimand τ are obtained via means, medians, or modes of the posterior distribution $p(\tau \mid \mathbf{T}, \mathbf{Z}, \mathbf{Y}^{\text{obs}})$, and standard interval estimates are obtained via the central credible intervals or highest posterior density intervals. As causal estimands are functions of observed and missing potential outcomes, we recognize that we can calculate this posterior distribution by integrating out \mathbf{Y}^{mis} , \mathbf{G}^{o} , and \mathbf{G}^{u} according to

$$\int p(\tau \mid \mathbf{T}, \mathbf{Z}, \mathbf{Y}^{\text{obs}}, \mathbf{G}^{\text{o}}, \mathbf{G}^{\text{u}}, \mathbf{Y}^{\text{mis}}) p(\mathbf{G}^{\text{o}}, \mathbf{G}^{\text{u}}, \mathbf{Y}^{\text{mis}} \mid \mathbf{T}, \mathbf{Z}, \mathbf{Y}^{\text{obs}}) d\mathbf{G}^{\text{o}} d\mathbf{G}^{\text{u}} d\mathbf{Y}^{\text{mis}}. \quad (3)$$

In order to sample from this posterior, we first derive the posterior distribution of the DoI and missing outcomes conditional on the observed data. Then, we calculate the posterior distribution of the estimand conditional on the observed data and imputations of the DoI and missing outcomes drawn from their posterior distributions. We perform all imputations and posterior calculations via Markov Chain Monte Carlo technique.

Four inputs are necessary to derive the posterior distributions of the causal estimands. (i) The first is knowledge of the assignment mechanism, as encoded in the probability mass/density function $p(\mathbf{Z} \mid \mathbf{T})$. As we consider designed experiments, this probability mass/density function is known. (ii) The second is the prior distribution for the DoIs conditional on \mathbf{T} and \mathbf{Z} , which we denote by the joint probability density function $p(\mathbf{G}^o, \mathbf{G}^u \mid \mathbf{T}, \mathbf{Z}, \phi)$ with a parameter vector ϕ . (iii) The third input is a model for the potential outcomes conditional on $\mathbf{T}, \mathbf{Z}, \mathbf{G}^o, \mathbf{G}^u$, which we represent via a joint probability density/mass function $p(\mathbf{Y}^{\text{obs}}, \mathbf{Y}^{\text{mis}} \mid \mathbf{T}, \mathbf{Z}, \mathbf{G}^o, \mathbf{G}^u, \theta)$ with parameter vector θ . (iv) The final input is the prior distribution for ϕ and θ , denoted by $p(\phi, \theta)$. We assume that ϕ and θ are distinct and do not share any parameters. We further assume the following cross-world independence condition about \mathbf{G} .

Assumption 4 \mathbf{G}^o and \mathbf{G}^u are a priori independent of each other given \mathbf{T}, \mathbf{Z} .

This assumption eliminates the residual correlation between any two cross-world potential degree of interference functions after adjusting for the measured information, \mathbf{T} and \mathbf{Z} . This assumption is necessary because the correlation structure between them cannot be learned from the data at hand. We describe our priors for the DoI, potential outcomes and the parameters (ϕ, θ) in §4.3.

Given those four inputs, we first express the complete data likelihood function under Assumption 3 and the assumption of independence in the statistical model according to

$$\begin{aligned} L(\theta, \phi \mid \mathbf{Y}^{\text{obs}}, \mathbf{Y}^{\text{mis}}, \mathbf{G}^o, \mathbf{G}^u, \mathbf{T}, \mathbf{Z}) &\propto p(\mathbf{Y}^{\text{obs}}, \mathbf{Y}^{\text{mis}}, \mathbf{G}^o, \mathbf{G}^u \mid \mathbf{T}, \mathbf{Z}, \theta, \phi) \\ &= p(\mathbf{Y}^{\text{obs}}, \mathbf{Y}^{\text{mis}}, \mathbf{G}^o, \mathbf{G}^u \mid \mathbf{T}, \theta, \phi) = \prod_{i=1}^N p(Y_i^{\text{obs}}, Y_i^{\text{mis}} \mid G_i^o, \mathbf{G}_i^u, \mathbf{T}_i, \theta) p(G_i^o, \mathbf{G}_i^u \mid \mathbf{T}_i, \phi). \end{aligned}$$

The posterior distribution of the model parameters $p(\phi, \theta \mid \mathbf{Y}^{\text{obs}}, \mathbf{T}, \mathbf{Z})$ is thus derived as

$$p(\phi, \theta \mid \mathbf{Y}^{\text{obs}}, \mathbf{T}, \mathbf{Z}) \propto p(\phi, \theta) \int \prod_{i=1}^N p(Y_i^{\text{obs}}, Y_i^{\text{mis}} \mid G_i^o, \mathbf{G}_i^u, \mathbf{T}_i, \theta) p(G_i^o, \mathbf{G}_i^u \mid \mathbf{T}_i, \phi) dG_i^o d\mathbf{G}_i^u dY_i^{\text{mis}}.$$

The posterior distribution of the model parameters based on the complete data likelihood function is derived in a straightforward manner as

$$p(\phi, \theta \mid \mathbf{Y}^{\text{obs}}, \mathbf{Y}^{\text{mis}}, \mathbf{G}^o, \mathbf{G}^u, \mathbf{T}, \mathbf{Z}) \propto p(\phi, \theta) \prod_{i=1}^N p(Y_i^{\text{obs}}, Y_i^{\text{mis}} \mid G_i^o, \mathbf{G}_i^u, \mathbf{T}_i, \theta) p(G_i^o, \mathbf{G}_i^u \mid \mathbf{T}_i, \phi).$$

Similarly, the posterior of the missing outcomes and DoIs conditional on the observed data and model parameters is $\prod_{i=1}^N p\{(\mathbf{Y}_i^{\text{mis}} \mid Y_i^{\text{obs}}, G_i^o, \mathbf{G}_i^u, \mathbf{T}_i, \theta) p(G_i^o, \mathbf{G}_i^u \mid Y_i^{\text{obs}}, \mathbf{T}_i, \phi)\}$.

We use the Gibbs sampler to obtain posterior draws for the imputations and thereby derive the posterior distributions of the causal estimands using (3) (Imbens and Rubin,

1997). Specifically, in each iteration of the sampler we alternate between drawing from the conditional posterior distributions of ϕ , θ , \mathbf{G}^o , \mathbf{G}^u , and \mathbf{Y}^{mis} given the other variables, respectively, and then use the draws of \mathbf{Y}^{mis} to obtain draws of the causal estimands from their respective posterior distributions. Details on the Gibbs sampler are in §4.5.

Remark 6 (Formalism from a Bayesian perspective) *Note that formalizing the DoI-based framework does not immediately imply a Bayesian perspective. We view our inferential framework as a model-assisted approach (Särndal et al., 1992; Basse and Airoidi, 2018a), and we use our semi-parametric model as an assistant to improve precision in estimating causal effects, yet continue to rely primarily on randomization (the design) as the fundamental source of valid causal inference. We may use the frequentist tool, such as fitting a semiparametric regression of the form $Y_i = X_i\beta + G_i + \epsilon_i$, where ϵ_i is the noise term, and G_i is a flexible, non-parametric function, which can be estimated as a function of covariates and treatment status of the neighboring units, $X_i, \mathbf{X}_{-i}, \mathbf{A}, \mathbf{Z}_{-i}$, using kernel regressions, spline-based approaches, local polynomial regressions, or similar methods. By flexibly capturing the interference mechanism, these methods enhance precision while design-based (or randomization-based) inference procedures preserve validity.*

We adopt a Bayesian perspective in this article because the Bayesian approach provides a coherent system for uncertainty quantification. Throughout this article, we also adopt a finite-population perspective for causal inference, where $Y_i(\mathbf{z})$ is treated as a fixed, non-random quantity. From the Bayesian perspective, the “model” is not merely about sampling from a super-population; rather, it structures our prior beliefs about outcomes in the finite population, enabling us to incorporate uncertainty about the interference mechanism (see discussions in Section 4.4).

4.3 Bayesian Semiparametric Models

The models in our Bayesian approach consist of a combination of parametric and nonparametric specifications. We adopt a semi-parametric model for the Y -model, with the particular form specified based on the specific setting under consideration. We shall discuss the specification of the Y -model later in our simulation studies and case study. In those studies, we set $\theta = (\beta, \lambda)^\top$, where β contain regression coefficients associated with the covariates and λ is a scale parameter in the regression. The prior for (β, λ) is the Normal-Inverse-Gamma distributions, $N(0, \sigma_1^2)IG(a_1, b_1)$ with σ_1^2 , a_1 and b_1 being the variance, shape and scale parameters respectively.

The interference structure captured by the DoI is generally complex, hence unlikely to be adequately represented by parametric models. As such, we utilize a Bayesian nonparametric prior for the G -model, specifically, the Dependent Dirichlet process mixture prior (DDPM). As indicated by equation (2), the estimation of the causal estimands is reduced to estimation of the probability density function of the DoI, and this further justifies our use of the DDPM as it is a natural choice for density estimation under the Bayesian paradigm. In addition to its flexibility, the clustering property of the DDPM, in terms of its potential ability to automatically separate units into several clusters based on how they interfere with one another, provides another advantage that motivates our adoption of this nonparametric model. Quintana et al. (2022) provides a comprehensive review of the DDPM.

A random probability measure H drawn from a Dirichlet Process, $\text{DP}(\alpha, H_0)$, with concentration parameter $\alpha > 0$ and base probability measure H_0 over a measurable space (Φ, \mathcal{B}) is such that for any finite partition (B_1, \dots, B_k) of \mathcal{B} , $(H(B_1), \dots, H(B_k)) \sim \text{Dir}(\alpha H_0(B_1), \dots, \alpha H_0(B_k))$, where $\text{Dir}(\alpha_1, \dots, \alpha_k)$ denotes the Dirichlet distribution with parameters $\alpha_1, \dots, \alpha_k > 0$ (Ferguson, 1974). The concentration parameter α determines the number of clusters in the Dirichlet Process. Our DDPM prior for the probability density function of the G_i is specified as

$$(G_1, \dots, G_N) \mid H_{\mathbf{T}, \mathbf{Z}} \stackrel{\text{ind}}{\sim} p(G \mid H_{\mathbf{T}, \mathbf{Z}}),$$

where $p(G \mid H_{\mathbf{T}, \mathbf{Z}}) = \int \eta(G \mid \phi) dH_{\mathbf{T}, \mathbf{Z}}(\phi)$ with $\eta(\cdot \mid \phi)$ being a continuous density function for every $\phi \in \Phi$, and $H_{\mathbf{T}, \mathbf{Z}} \sim \text{DP}(\alpha, H_0)$ for some $\alpha > 0$ and a base measure H_0 (Maceachern, 2000), where $\stackrel{\text{ind}}{\sim}$ indicates independently distributed variables. The subscript of $H_{\mathbf{T}, \mathbf{Z}}$ indicates the dependence of the measure on \mathbf{T} and \mathbf{Z} . As shown below, we include covariate information in the DDPM by allowing the locations in the stick-breaking construction of the $\text{DP}(\alpha, H_0)$ to depend on covariates. We consider using common weights across the values of DoI, so that our DDPM is a “single-weights” DDPM, for its simplicity and ease of implementation (which follows as it ignores the dependence of the weights on covariates). The representation of the stick-breaking process is $H_{\mathbf{T}, \mathbf{Z}} = \sum_{k=1}^{\infty} w_k \delta_{\phi_k(\mathbf{T}, \mathbf{Z})}(\cdot)$, $w_k = v_k \prod_{l < k} (1 - v_l)$, and $v_l \stackrel{\text{ind}}{\sim} \text{Beta}(1, \alpha)$. Ultimately, our G -model is

$$p(G_i \mid \mathbf{T}_i, \mathbf{Z}_{-i}, \phi) \propto \sum_{k=1}^{\infty} w_k (\sigma_k^2)^{-1/2} \exp\left(-\frac{1}{2\sigma_k^2} (G_i - \mu(\mathbf{T}_i, \mathbf{Z}_{-i}, \gamma_k))^2\right) \quad (4)$$

where the atoms $\phi_k = (\gamma_k, \sigma_k^2)$ and the weights w_k are nonparametrically specified via the $\text{DP}(\alpha, H_0)$. This is an infinite mixture of Normal distributions with $\mu(\mathbf{T}_i, \mathbf{Z}_{-i}, \gamma_k)$ being the location parameter of each component that is a function of the covariates of unit i , the i th row of A , and treatment assignments excluding that of unit i .

Finally, we specify the priors for α and H_0 . We use the $\text{Gamma}(a, b)$ prior for α with shape a and scale b (Escobar and West, 1995), and obtain posterior draws of α via a Metropolis-Hastings step in our Gibbs sampler. Next, we take H_0 as a Normal-Inverse-Gamma conjugate $\text{N}(\mu_0, \sigma_0^2)\text{IG}(a_0, b_0)$ with a_0 and b_0 being the shape and scale parameter respectively. Our simulation studies and case studies illustrate how the values for all the hyperparameters are set.

Remark 7 (The role of μ with domain knowledge about interference) *The DDPM offers a straightforward way to model distributions that evolve depending on covariates. In our case, using the DDPM to represent the uncertainty of the interference structure (as detailed in §4.1) allows this uncertainty to evolve based on characteristics such as the network matrix, treatment variables, and covariates, which are typically extremely high-dimensional. Therefore, following Maceachern (1999), we advocate for specifying the dependence on the uncertainty distribution by lower-dimensional summaries of these characteristics, incorporating the analysts’ domain knowledge about interference. In this context, we specify μ to incorporate our possibly limited understanding of the interference structure. By specifying μ , researchers can leverage their domain expertise to develop network features that are expected*

to influence interference. On the other hand, while specifying μ_k is necessary, the precise choice of μ_k is not as critical. This is because the resulting marginal distribution for G_i forms a flexible, infinite mixture distribution, as detailed in Equation (4). The key advantage of adopting the DDPM for the G -model lies in this nonparametric nature. In our simulation studies in §5, we have verified that even basic forms of μ_k are effective in handling complex interference structures.

4.4 Structural Assumptions for Inference and Prior Elicitation

The DoI framework is conceptually straightforward, gives us a clear interpretation of the treatment effects under interference, and extends conventional exposure maps by allowing itself to be unknown, as discussed in §4.1. However, this flexibility comes with a trade-off, necessitating an additional structural assumption between \mathbf{Y} and \mathbf{G} . Specifically, when G_i is fully latent, Assumption 1 remains non-restrictive because any function mapping each value of the treatment assignment vector to a unique real number will satisfy the assumption. This is true because the cardinality of the support of the treatment assignment vector is always less than or equal to that of the real numbers. As §4.2 implies, statistical inference regarding the causal effects requires some modeling assumptions between \mathbf{Y} and \mathbf{G} .

This structural assumption can be tailored to specific applications, allowing the integration of prior knowledge into the inference model. For example, one can view each unit’s outcome Y_i as determined by its own treatment Z_i and a latent function G_i that collects how other units’ treatments, covariates, and the network structure affect Y_i . Formally, $Y_i = f_i(Z_i, G_i(X_i, \mathbf{X}_{-i}, \mathbf{A}, \mathbf{Z}_{-i}), \epsilon_i)$, where ϵ_i is an exogenous noise term and G_i is nonparametric, flexibly characterizing the interference channel from other units. Assumption 1 (equating potential outcomes whenever G_i coincides in distribution) then says that if two configurations of $\{X_i, \mathbf{X}_{-i}, \mathbf{A}, \mathbf{Z}_{-i}\}$ yield the same effective distribution for G_i , they are indistinguishable to Y_i . From a more classical statistical modeling viewpoint, the exposure map may be interpreted as a *fixed effect* (i.e., deterministic function of the observed data), while the DoI can be viewed as a *random effect* that encapsulates additional, possibly unobserved, sources of variation. Consequently, if we specialize f_i to be a linear function in its arguments, then the DoI framework reduces to a linear mixed model, with G_i playing the role of a random component that depends on the network structure and other units’ treatments.

A key advantage of our DoI framework, with its structural assumption, over traditional exposure mapping is the clear separation of roles. As Sävje (2023) notes, exposure mapping in existing literature serves a dual purpose: defining direct and spillover effects and imposing assumptions on the interference structure. This dual role necessitates defining both the interference structure, which determines a unique potential outcome, and exposure effects, which determine causal quantities. Given that exposure mapping is problem-specific, its effects often lack scientific meaning outside the given problem. Moreover, misspecification of these maps typically leads to biased estimators (Sävje, 2023).

This separation enables a unified procedure in our inferential approach, allowing for the incorporation of case-specific interference structures in estimating causal effects without altering the effects being estimated. This is because our causal estimands are generally defined in terms of treatment assignments and are not dependent on exposures. Secondly, we incorporate knowledge about interference structures by specifying priors on the DoI,

capturing their complexities and uncertainties, and yielding valid posterior distributions for finite-population causal estimands. The Bayesian methodology propagates uncertainties about interference structures using an imputation-based approach, as outlined in §4.2. Our approach distinguishes the interference structure from the effects of interest, as our causal estimands are generally defined in terms of treatment assignments and are not dependent on exposures.

To motivate the idea of prior elicitation given prior investigation or knowledge about the interference structure, we discuss how to incorporate this knowledge in the form of the DoI. Suppose we know that the interference structure satisfies the stratified interference condition (Hudgens and Halloran, 2008), which states that spillover effects are determined solely by the proportion of treated units within the same cluster as each unit. In this case, we may specify the function $\mu(\mathbf{T}_i, \mathbf{Z}_{-i}, \gamma)$ for unit i as

$$\mu(\mathbf{T}_i, \mathbf{Z}_{-i}, \gamma_k) = \gamma_k \frac{\sum_{j \in \mathcal{G}_i} Z_j}{|\mathcal{G}_i|},$$

where \mathcal{G}_i is the connected subgroup to which unit i belongs, and $|\mathcal{G}_i|$ is the size of \mathcal{G}_i . If we have more detailed knowledge that, for all i , unit i is affected by its direct neighbors and m -th order neighbors in the same connected graph, but is affected by them differently, we may specify

$$\mu(\mathbf{T}_i, \mathbf{Z}_{-i}, \gamma_k) = \sum_{s=1}^m \gamma_{s,k} \frac{\sum_{j \in \mathcal{G}_i^{(s)}} Z_j}{|\mathcal{G}_i^{(s)}|},$$

where $\mathcal{G}_i^{(s)}$ denotes the set of s -th order neighbors (neighbors at distance s) of unit i , and $|\mathcal{G}_i^{(s)}|$ is its corresponding size. If we wish to include additional demographic information about the neighbors, we may adjust the function accordingly (e.g., the number of teenagers, etc.).

If individuals have even deeper insights into the interference structure and know that spillover effects are determined by certain metrics of the neighbors that represent the local importance of units (e.g., centrality, eigenvalues, PageRank metrics (Newman, 2010), etc.), as is often the case in social network analysis, we may incorporate these metrics into the specification of $\mu(\mathbf{T}_i, \mathbf{Z}_{-i}, \gamma_k)$. For example,

$$\mu(\mathbf{T}_i, \mathbf{Z}_{-i}, \gamma_k) = \gamma_{1,k} \sum_{j \neq i} P_j \exp(-\gamma_{2,k} d(i, j)),$$

where P_j represents the importance metric of unit j , and $d(i, j)$ represents the (potentially weighted) distance between units i and j in the network. The term $\exp(-\gamma_{2,k} d(i, j))$ acts as a decaying factor for the influence of unit j on unit i .

Example 1 (Linear-in-Means Model) *The Linear-in-Means (LIM) model is a social interaction model that is commonly used in the economic literature (Bramoullé et al., 2009; Chin, 2019; Bramoullé et al., 2020; Leung, 2022). Assuming $|\gamma| < 1$, the model defines potential outcomes $Y_i(\mathbf{Z})$ through*

$$\mathbf{Y} = \alpha + \beta \mathbf{Z} + \gamma \tilde{\mathbf{A}} \mathbf{Y} + \delta \tilde{\mathbf{A}} \mathbf{Z} + \boldsymbol{\epsilon},$$

where $\boldsymbol{\epsilon} = (\epsilon_1, \dots, \epsilon_N)^\top$ and $\tilde{\mathbf{A}}$ is obtained by row-normalizing \mathbf{A} (divide each row by its sum). By the matrix identity $(\mathbf{I} - \gamma\tilde{\mathbf{A}})^{-1} = \sum_{k=0}^{\infty} \gamma^k \tilde{\mathbf{A}}^k$, one obtains the unit-level reduced form of the model (Bramoullé et al., 2009):

$$Y_i = \tilde{\alpha} + \tilde{\beta}Z_i + \sum_{k=0}^{\infty} \tilde{\gamma}_k (\tilde{\mathbf{A}}^{k+1}\mathbf{Z})_i + \tilde{\epsilon}_i, \quad (5)$$

where $\tilde{\alpha} = \alpha/(1 - \gamma)$, $\tilde{\beta} = \beta$, $\tilde{\gamma}_k = (\beta\gamma + \delta)\gamma^k$, $\tilde{\epsilon} = \sum_{k=0}^{\infty} \gamma^k \mathbf{A}^k \boldsymbol{\epsilon}$, and $(\tilde{\mathbf{A}}^{k+1}\mathbf{Z})_i$ is the i -th coordinate of $\tilde{\mathbf{A}}^{k+1}\mathbf{Z}$. The third term indicates that the impact of treatments assigned to k -neighbors is exponentially down-weighted by $\tilde{\gamma}_k$. Our additive semi-parametric model specification is a general form of the LIM model because (5) is expressed as a linear combination of the direct effects ($\tilde{\beta}Z_i$ and the term that corresponds to Z_i in $\sum_{k=0}^{\infty} \tilde{\gamma}_k (\tilde{\mathbf{A}}^{k+1}\mathbf{Z})_i$) and the infinite sum of the spillover effects (the terms that correspond to \mathbf{Z}_{-i} in $\sum_{k=0}^{\infty} \tilde{\gamma}_k (\tilde{\mathbf{A}}^{k+1}\mathbf{Z})_i$), and the infinite sum of the spillover can correctly be specified by the infinite mixture of the G -model (4).

Remark 8 (Extension to more general interaction structures) Although we focus here on network-based adjacency information (which typically encodes pairwise edges between units), the DoI framework is not limited to strictly pairwise (graphical) interactions. In principle, one can generalize the set \mathcal{A} to represent hypergraphs or other higher-order structures in which “edges” (hyperedges) connect subsets of more than two units.

Concretely, suppose we have a hypergraph $\mathcal{H} = \{e_1, e_2, \dots, e_K\}$, where each hyperedge $e_k \subseteq \{1, 2, \dots, N\}$ is a subset of units. The usual adjacency matrix \mathbf{A} in a graph may be replaced by an incidence structure for hypergraphs, for instance, an incidence matrix $\mathbf{B} \in \{0, 1\}^{N \times K}$ where

$$B_{i,k} = \begin{cases} 1, & \text{if unit } i \in e_k, \\ 0, & \text{otherwise.} \end{cases}$$

Incorporating such a higher-order structure into the DoI framework means that each unit i has access to information not just about pairwise connections with others but also about membership in shared hyperedges. Thus, the characteristic set \mathcal{T} can be expanded to capture hypergraph membership. Formally, we can write $\mathcal{T} = \mathcal{X} \times \mathcal{X}^{N-1} \times \mathcal{H}$, where \mathcal{H} denotes the set of all possible hypergraph incidence structures (analogous to how \mathcal{A} previously denoted the set of all adjacency matrices).

Then, the DoI for unit i can be modified to $G_i : (\mathcal{X} \times \mathcal{X}^{N-1} \times \mathcal{H}) \times \Omega^{N-1} \rightarrow \mathbb{R}$, so that when evaluating G_i , one accounts for all relevant hyperedges e_k that involve unit i (or subsets thereof that might exert a joint influence on i). For example, if one models a situation where interference depends on whether multiple distinct units are jointly exposed to certain treatments, these higher-order relationships can be captured by the hyperedges in \mathcal{H} . For instance, one may specify a mean function

$$\mu(\mathbf{T}_i, \mathbf{Z}_{-i}, \gamma_j) = \sum_{k=1}^K \gamma_{j,k} \mathbf{1}\{i \in e_k\} \frac{\sum_{j \in e_k, j \neq i} Z_j}{|e_k|},$$

which measures the proportion of other treated units in each hyperedge e_k . This flexibility allows the DoI framework to model a richer class of interference structures than what can be

described by a standard adjacency matrix, accommodating any scenario in which the influence on a given unit depends on group-level attributes or actions of other units, rather than strictly pairwise interactions.

4.5 Details on the Gibbs Sampler

A discussion on efficient MCMC algorithms for DDPM models is provided by Escobar (1994). Our Y -model is generally specified by $\boldsymbol{\theta}$ as $p(Y_i(z, g_i) \mid \mathbf{T}_i, G_i = g_i, \boldsymbol{\theta})$. For our DDPM, we utilize an approximated blocked Gibbs sampler based on a truncation of the stick-breaking representation of the DP proposed by Ishwaran and Zarepour (2000). This algorithm proceeds by first selecting a conservative upper bound on the number K of clusters. Let $C_i \in \{1, \dots, K\}$ denote the latent class indicators for unit i . We specify a Multinomial distribution $C_i \sim \text{MN}(\mathbf{w})$ on C_i , where $\mathbf{w} = (w_1, \dots, w_K)^\top$ contains the weights from the DDPM. Conditional on $C_i = k$, equation (4) is simplified to a single Normal component. Ishwaran and James (2001) demonstrated that an accurate approximation to the exact DP is obtained as long as K is sufficiently large. To ensure this, we ran several MCMC iterations with different values of K and increased K after an iteration if all clusters were occupied. We terminated this process when the number of occupied clusters was less than K . Our entire algorithm proceeds as follows; throughout, ‘ $\mid -$ ’ indicates conditional on all others.

1. Given $\boldsymbol{\theta}$, $\boldsymbol{\phi}$ and $\mathbf{C} = (C_1, \dots, C_N)$, draw each G_i^o from

$$p(G_i^o \mid -) \propto p\left(Y_i^{\text{obs}} \mid \mathbf{T}_i, G_i^o, \boldsymbol{\theta}\right) p(G_i^o \mid \mathbf{T}_i, \mathbf{Z}_{-i}, \boldsymbol{\phi}_{C_i}).$$

2. Given $\boldsymbol{\phi}$, \mathbf{w} and G_i^o , draw each C_i from $p(C_i = k \mid -) \propto w_k p(G_i^o \mid \mathbf{T}_i, \mathbf{Z}_{-i}, \boldsymbol{\phi}_k)$.
3. Let $w'_K = 1$. Given α , \mathbf{C} , draw w'_k for $k \in \{1, \dots, K-1\}$ from $w'_k \sim \text{Beta}\left(1 + \sum_{i:C_i=k} 1, \alpha + \sum_{i:C_i>k} 1\right)$. Then, update $w_k = w'_k \prod_{j<k} (1 - w'_j)$ for $k = 1, \dots, K$.
4. Given \mathbf{C} and \mathbf{w}' , draw α from $p(\alpha \mid -) \propto p(\alpha) \prod_{k=1}^K f(w'_k \mid 1 + \sum_{i:C_i=k} 1, \alpha + \sum_{i:C_i>k} 1)$ where f is the pdf of w'_k , the beta distribution.
5. Given \mathbf{C} and \mathbf{G}^o , draw $\boldsymbol{\phi}_k$ from $p(\boldsymbol{\phi}_k \mid -) \propto H_0(\boldsymbol{\phi}_k) \prod_{i:C_i=k} p(G_i^o \mid \mathbf{T}_i, \mathbf{Z}_{-i}, \boldsymbol{\phi}_k)$.
6. Given $\boldsymbol{\phi}$, \mathbf{C} and \mathbf{G}^o , draw $\boldsymbol{\theta}$ from $p(\boldsymbol{\theta} \mid -) \propto p(\boldsymbol{\theta}) \prod_{i=1}^N p(Y_i^{\text{obs}} \mid \mathbf{T}_i, G_i^o, \boldsymbol{\theta})$.

The imputations of the $\mathbf{Y}_i^{\text{mis}}$ are generated from the Y -model using the G_i^o in Step 1 and the $\boldsymbol{\theta}$ in Step 6. Specifically, for $i = 1, \dots, N$:

7. Given G_i^o in Step 1 and $\boldsymbol{\theta}$ in Step 6, draw from $Y_i(z', \mathbf{z}_{-i}) \sim p(Y_i \mid Z_i = z', \mathbf{T}_i, G_i^o, \boldsymbol{\theta})$.

To obtain posterior draws of the spillover effect estimands, we impute the $Y_i(z, \mathbf{z}'_{-i})$ by using the posterior draws of $G_i^u = G(\mathbf{T}_i, \mathbf{z}')$ and $\boldsymbol{\theta}$.

8. Given \mathbf{w} in step 3, draw C_i^u from $C_i^u \sim \text{MN}(\mathbf{w})$.
9. Given C_i^u and $\boldsymbol{\phi}$, draw G_i^u from $G_i^u \sim \text{N}(\mu(\mathbf{T}_i, \mathbf{z}'_{-i}, \gamma_{C_i^u}), \sigma_{C_i^u}^2)$.
10. Given G_i^u in step 9 and $\boldsymbol{\theta}$ in step 6, draw from $Y_i(z, \mathbf{z}'_{-i}) \sim p(Y_i \mid Z_i = z, \mathbf{T}_i, G_i^u, \boldsymbol{\theta})$.

This Gibbs sampler ultimately enables us to impute all of the missing potential outcomes that define the assignment-conditional effects, and yields approximations for the expected effects. Additional technical details regarding these points and each step are provided in the Appendix.

5. Simulation Studies

We perform simulation studies to evaluate the performance of our Bayesian methodology under our DoI framework with respect to that of the Horvitz-Thompson (HT) estimator in the case of a Bernoulli trial. The HT estimator was selected as the comparison method in our evaluation because Sävje et al. (2021) and Li and Wager (2022) have demonstrated that it is a consistent estimator for the E-ATE under a Bernoulli trial under interference and that it possesses desirable large sample properties.

To compare these methods, we assess the bias, mean square error (MSE), and coverage of a procedure based on M simulated datasets by calculating $\sum_{n=1}^{N_{\text{sim}}} (\tau - \hat{\tau}_n) / N_{\text{sim}}$, $\sum_{n=1}^{N_{\text{sim}}} (\tau - \hat{\tau}_n)^2 / N_{\text{sim}}$ and $\sum_{n=1}^{N_{\text{sim}}} \mathbb{1}(\hat{\tau}_n^l \leq \tau \leq \hat{\tau}_n^u) / N_{\text{sim}}$ respectively, where τ denotes the true causal estimand, and $\hat{\tau}_n$, $\hat{\tau}_n^l$ and $\hat{\tau}_n^u$ are the estimates of the causal estimand, 2.5% lower endpoint of the interval estimator, and 95% upper endpoint of the interval estimator in dataset $n = 1, \dots, N_{\text{sim}}$. Here, we take the posterior mean of the causal estimand as our Bayesian point estimator and the 95% central credible interval as the interval estimator. For the HT estimator under the Bernoulli trial, Sävje et al. (2021) gave the variance estimator and conservative confidence interval. As the true value of E-ATE is hard to obtain in closed form, we approximated it in our simulation study via the Monte Carlo average of 1000 draws for each generated dataset and graph.

We consider seven different data-generating mechanisms and two different underlying networks of experimental units. The seven data-generating mechanisms that we consider are:

1. $Y_i = \mathbf{X}_i^\top \beta_1 + Z_i \tau + \psi_1 \sum_{j \in \mathcal{N}_i} Z_j A_{ij} / (|\mathcal{N}_i| + 1) + \epsilon_i$,
2. $Y_i = \mathbf{X}_i^\top \beta_1 + Z_i \tau + \psi_1 \sum_{j \in \mathcal{N}_i} P_j Z_j A_{ij} / (|\mathcal{N}_i| + 1) + \epsilon_i$,
3. $Y_i = \left(\mathbf{X}_i^\top \beta_1 + \psi_1 \sum_{j \in \mathcal{N}_i} P_j Z_j A_{ij} / (|\mathcal{N}_i| + 1) \right) \exp \left(-\psi_2 Z_i \sum_{j \in \mathcal{N}_i} P_j Z_j A_{ij} / (|\mathcal{N}_i| + 1) \right) + \epsilon_i$,
4. $Y_i = \left(\mathbf{X}_i^\top \beta_1 + Z_i \tau + \psi_1 \sum_{j \in \mathcal{N}_i} P_j Z_j A_{ij} / (|\mathcal{N}_i| + 1) \right) \exp \left(-\psi_2 \sum_{j \in \mathcal{N}_i} P_j Z_j A_{ij} / (|\mathcal{N}_i| + 1) \right) + \epsilon_i$,
5. $Y_i = \left(\mathbf{X}_i^\top \beta_1 + Z_i \tau + \psi_1 \sum_{j \in \mathcal{N}_i} P_j Z_j A_{ij} / (|\mathcal{N}_i| + 1) \right) \cos \left(\pi \psi_2 \sum_{j \in \mathcal{N}_i} P_j Z_j A_{ij} / (|\mathcal{N}_i| + 1) \right) + \epsilon_i$,
6. $Y_i = \mathbf{X}_i^\top \beta_1 + \psi_2 U_i + Z_i \tau U_i + \psi_1 U_i \sum_{j \in \mathcal{N}_i} Z_j A_{ij} / (|\mathcal{N}_i| + 1) + \epsilon_i$,
7. $Y_i = \mathbf{C}_i^\top \beta_2 + Z_i \tau + \psi_1 \sum_{j \in \mathcal{N}_i} Z_j A_{ij} / (|\mathcal{N}_i| + 1) + \epsilon_i$,

where $\beta_1 = (-1, 1.5)^\top$, $\beta_2 = (-1, 1.5, 0.5, -0.5, 0.6, -0.6, 0.8, -0.8, 1.0, -1.0)^\top$, $\tau = 3$, $\psi_1 = 2$, $\psi_2 = 0.5$, $\mathbf{X}_i \in \mathbb{R}^2$ has first entry $X_{i,1} = 1$ and second entry $X_{i,2} \sim N(0, 1)$, $\mathbf{C}_i \in \mathbb{R}^{10}$ has first entry $C_{i,1} = 1$ and other entries $C_{i,j} \sim N(0, 1)$ for $j = 2, \dots, 10$, $Z_i \sim \text{Bernoulli}(0.5)$, $\epsilon_i \sim N(0, 1)$ is an additive error term, A_{ij} is the (i, j) entry of the adjacency matrix, \mathcal{N}_i is the set of neighboring units for unit i , P_i denotes the eigenvector centrality measure (Newman, 2010, p. 169), and $U_i \sim N(1.0, 0.5)$ represents an unmeasured covariate. The P_i is considered as a proxy for the ‘‘importance’’ of unit i on the network. We assume that this measure quantifies the interference structure and spillover effects for each unit. This metric evaluates the importance of a node in a graph, based not only on the number of its direct connections (degree) but also on the centrality of its connected nodes. By accounting for these indirect influences, it effectively captures the complex spillover effects common in real-world applications, such as social network analysis, where interference issues frequently arise. Scenario 1 corresponds to the stratified interference structure. Including P_i in Scenarios

2–5 creates more complicated interference structures because relative scores are assigned to all nodes within the network, based on the principle that connections to high-scoring nodes contribute more significantly to a node’s score than identical connections to low-scoring nodes. Untreated units are down-weighted by 0.1 when computing P_i . Furthermore, the treatment effect in Scenario 3 has a nonlinear multiplicative form in that the treatment variable Z_i appears in the exponential function. Additionally, Scenarios 4 and 5 have additive and multiplicative treatment and spillover effects. Scenario 6 examines the interference pattern involving an unmeasured variable U_i , which is not used for inference. In this scenario, the outcome Y_i depends on U_i both directly and through its role in modifying the effects of treatment and spillover. It is important to note that U_i is not necessarily an unmeasured confounder, as the treatment assignment does not depend on U_i . Finally, Scenario 7 examines the performance of the proposed method for a moderate-sized dimension of covariates.

To see the performance of the methodologies under different underlying networks, we use two different random graph models to generate the underlying networks: the Erdős–Rényi and Barabási–Albert graphs (Frieze and Karoński, 2016). In all seven scenarios, we simulate $N_{\text{sim}} = 200$ datasets for each sample size of $N = 300, 500, 1000, 2000, 3000$. We ran the MCMC algorithm for 2000 iterations after a burn-in of 2000. We chose the iteration numbers after experimentation verifying they deliver stable results over multiple runs.

For scenario 1–6, the Y -model is specified by $\boldsymbol{\theta} = (\beta, \lambda)^\top$ as: $Y_i(z) \sim N(\mathbf{X}_i^\top \boldsymbol{\beta}_z + G_i(\mathbf{T}_i, \mathbf{Z}_{-i}, \lambda_z), \sigma_z^2)$. In the G -model we set $\mu(\mathbf{T}_i, \mathbf{Z}_{-i}, \gamma_k) = \gamma_{k,1} + \gamma_{k,2} \sum_{j \in \mathcal{N}_i} Z_j A_{ij}$. The prior distributions for the parameters in this case are $\boldsymbol{\beta}_z \sim N((0, 0)^\top, 10^2 \mathbf{I}_2)$, where \mathbf{I}_2 is the 2×2 identity matrix, $\lambda_z \sim \text{IG}(0.1, 0.1)$, $\gamma_{k,1}, \gamma_{k,2} \sim N(0, 10^2)$ and $\sigma_k^2 \sim \text{IG}(0.1, 0.1)$. For scenario 1, the inference model is well specified for the additive interference structure of the Y -model as well as the G -model in §4.4. For scenarios 2, the Y -model has a correct additive form but the location function μ_k of the G -model is simplified compared to the true data-generating process. Specifically, it considers only the treatment status of the neighboring units and ignores the spillover effects from k -neighbors for $k \geq 2$ in its specification. This scenario assesses the robustness of our methodologies for a specific, simple choice of μ_k . For scenarios 3–5, the inference model is misspecified in that the data-generating models are multiplicative of treatment and spillover effect terms (exp and cos), whereas the inferential model is additive. These four scenarios assess the robustness of our approach to structural misspecification. Scenario 6 further evaluates the robustness of the proposed methodology in the presence of an unmeasured covariate. For Scenario 7, the lower-dimensional covariates X_i are replaced with higher-dimensional covariates C_i and the dimension of $\boldsymbol{\beta}_z$ increases accordingly.

Figures 1 and 2 summarize the MSE of the two methods. These figures show that the proposed approach outperforms the HT estimator for all scenarios in both networks. The MSE converges at the same rate across all scenarios, indicating that the proposed approach is as efficient as the frequentist estimator. The proposed approach consistently performs well, even under scenarios of model misspecifications (scenarios 2 to 5), demonstrating its efficiency and robustness. In particular, the successful results in Scenario 2 highlight that the specific choice of the location function μ_k is less critical, as discussed in §4.3. This is attributed to the nonparametric nature of the DPMM, represented by the infinite mixture of distributions characterized by different atoms, which compensates for possibly incorrect specification of μ_k . In the challenging scenario, Scenario 6, which involves an unmeasured covariate, the MSE slightly increases for both the HT estimator and the proposed method. However, the proposed




Figure 1: MSE of the HT estimator and our proposed Bayesian methodology for the Erdős-Rényi graphs $ER(N, p)$ with different sparsity parameters p .

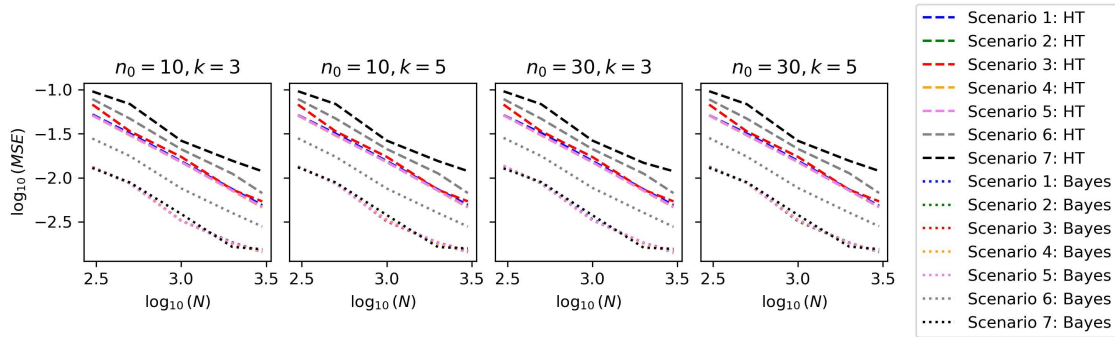


Figure 2: MSE of the HT estimator and our proposed Bayesian methodology for the Barabási–Albert graphs $BA(N, n_0, k)$ with different sparsity parameters (n_0, k) ; n_0 is the initial number of nodes and k is the number of edges a new node attaches to existing nodes.

method continues to demonstrate superior performance, highlighting the reliability of the DoI framework in practical contexts where the data may not be sufficiently rich. Another interesting observation is that when we have higher-dimensional covariates in Scenario 7, the improvement of our methodology over the traditional HT estimator becomes more pronounced. This is because, in high-dimensional settings, the increased auxiliary information enables our model-assisted approach to capture more of the outcome variation and thus reduce variance more effectively, whereas HT estimator does not fully exploit such auxiliary information.

Further evaluation of bias and coverage for the treatment effects and all metrics for the spillover effects are provided in the Appendix. Those results also demonstrate the efficiency of the proposed method under interference.

6. Case Studies

6.1 Cash Transfer Programs in Colombia

The methodology developed in §4 is applied to a randomized experiment investigating the impact of a conditional cash transfer program, treatment, on students' attendance rates (Barrera-Osorio et al., 2011). The study was conducted in two regions of Bogota, Colombia: San Cristobal and Suba. Households with school children, ranging from one to five children per household, were recruited for the experiment. Within each household, school children were randomly assigned to either enroll in the cash transfer program or not, by a stratified randomization approach based on locality (San Cristobal/Suba), school type (public/private), gender, and grade level.

Multiple children within a household had the potential to receive the treatment, and within each stratum, all children had an equal chance of being treated. Additionally, given the known treatment assignment, Barrera-Osorio et al. (2011) examined both the direct effect of the program and the possible spillover effect on siblings within the same household. Enrolled students were eligible for cash subsidies if they attended school at least 80% of the time in a given month. Enrolling one student in the program could influence the attendance rate of their sibling(s) in the same household, either positively or negatively. Our objective is to employ the proposed methodology to analyze these direct and spillover effects in the two regions of Bogota, Colombia.

Formally, for student i , we denote $Y_i \in \{0, 1\}$ where $Y_i = 1$ if the student is eligible for cash subsidies, i.e., their attendance rate is over 80%. We also let Z_i denote a treatment indicator that equals 1 if student i was enrolled in the program, and 0 otherwise. We consider the interference occurs within the same household and thus denote $A_{i,j}$ as the adjacency that equals 1 if student i and j are in the same household (siblings), and 0 otherwise. \mathcal{N}_i denotes the number of siblings of unit i . Additionally, \mathbf{X}_i represents the pre-treatment covariates, including the student's age, grade level, gender, household head's age, single-parent household indicator, household size, household's poverty score, household's income status, locality, and the number of students in the household who participated in the lottery. We establish three cluster types based on the two regions and household sizes: households in Suba with two students, households in San Cristobal with two students, and households in San Cristobal with three or more students. Overall, our analysis comprises 1012 households containing 2135 students.

We adopt the original treatment randomization probability as the treatment allocation strategy, indicated by the treatment allocation probabilities $\pi = (\pi_{\text{SC}}, \pi_{\text{Su}}) = (0.628, 0.449)$ for the two localities San Cristobal and Suba respectively. The primary estimands of interest are E-ATE $\tau_{\text{E-ATE}}(1; \pi)$ and E-ASE $\tau_{\text{E-ASE}}(0; \pi)$ for the current treatment allocation strategy π , see §3.2. In the given context, the spillover effect measures the discrepancy in the attendance rate of an unenrolled student under the allocation strategy π when none of their siblings is enrolled.

Following the discussion in §4.4, we use a probit model for the Y -model: $Y_i(z) \sim \text{Bernoulli}(\Psi(\mathbf{X}_i^\top \boldsymbol{\beta}_z + G_i(\mathbf{T}_i, \mathbf{Z}_{-i})))$, where $\Psi(\cdot)$ denotes the standard normal cumulative distribution function. The details of the Gibbs sampler for the binary outcome are provided in the Appendix. We here customize the components of G_i to reflect our knowledge about the interference: $\mu(\mathbf{T}_i, \mathbf{Z}_{-i}, \gamma_k) = \gamma_{k,1} \sum_{j \in \mathcal{N}_i} Z_j A_{ij} / |\mathcal{N}_i| + \gamma_{k,2} \left(\text{Grade}_i - \sum_{j \in \mathcal{N}_i} \text{Grade}_j Z_j A_{ij} / |\mathcal{N}_i| \right)$,

Table 1: Causal estimands and their estimates of posterior means, standard deviation, percentiles, and interval length under the current treatment allocation strategy π . The superscript $_{nb,k}$ denotes the spillover effects on students with k siblings. The superscript $_{rt,p}$ denotes the spillover effects on students for whom $p\%$ of their siblings are treated. E-ATE, expected average treatment effects; E-ASE, expected average spillover effects.

Estimands	Mean	SD	2.5%	Median	97.5%	Length
E-ATE $_{\pi}$	0.011	0.007	-0.003	0.011	0.025	0.028
E-ASE $_{\pi}$	-0.003	0.006	-0.014	-0.003	0.009	0.023
E-ASE $_{\pi}^{nb,1}$	-0.003	0.006	-0.014	-0.003	0.01	0.024
E-ASE $_{\pi}^{nb,2}$	0.0	0.017	-0.033	0.0	0.036	0.069
E-ASE $_{\pi}^{rt,33}$	-0.001	0.026	-0.053	0.0	0.051	0.105
E-ASE $_{\pi}^{rt,50}$	-0.0	0.008	-0.016	-0.001	0.016	0.032
E-ASE $_{\pi}^{rt,66}$	0.004	0.029	-0.051	0.0	0.061	0.113

where Grade_i denotes the grade of student i . The intuition of this specification is that students are more encouraged to enroll in the program if more or fewer siblings are treated, and they are likely influenced if their siblings are close in grade. We use the same proper and weakly informative priors, $\beta_z \sim N((0, 0)^\top, 10^2 \mathbf{I}_2)$, where \mathbf{I}_2 is the 2×2 identity matrix, $\gamma_{k,1}, \gamma_{k,2} \sim N(0, 10^2)$ and $\sigma_k^2 \sim \text{IG}(0.01, 0.01)$. We present a sensitivity analysis with respect to different prior specifications in the following section.

Table 1 presents the estimates of E-ATE and E-ASE under the current treatment allocation strategy π . We also present the conditional E-ASE for those with 1 or 2 siblings and those whose siblings' treated proportion is 33%, 50%, or 66%. We observe that the median of E-ATE is positive and those of E-ASE are mostly negative, implying the potential positive direct treatment effect and negative spillover effects. The spillover effects tend to be slightly larger as the number of siblings or the proportion of treated siblings increases. Table 2 presents the estimates of the E-ASE under different treatment allocation strategies $\pi' \in \{(0.1, 0.1), (0.3, 0.3), (0.5, 0.5), (0.7, 0.7), (0.9, 0.9)\}$, $\tau_{\text{E-ASE}}(0; \pi')$. We observe that the effect is always negative. As the proportion of treated units increases, the effects slightly get larger.

Finally, when performing Bayesian analyses with weakly identifiable models, it is important to investigate the robustness of the results with respect to the prior specifications so as to make inferences more reliable. Considering different prior specifications, we examined the robustness of our analyses; the results are provided in the Appendix.

6.2 Inference on the Degree of Interference in an Experiment on Stereolithography

Additive manufacturing (AM), or three-dimensional (3D) printing, refers to a manufacturing technology in which material is deposited and solidified additively based on a computer-aided design (CAD) model to manufacture the corresponding physical product. A significant challenge with AM technologies is dimensional accuracy control of a printed product. Shape deviations inevitably arise in additively manufactured products due to the physics associated with the rapid heating and cooling in any AM process. Several studies have been performed to

Table 2: The spillover effects and their estimates of posterior means, standard deviation, percentiles, and interval length under different treatment allocation strategies π' . E-ASE $_p$ denotes expected average spillover effects under the treatment allocation strategy $\pi' = (p, p)$.

Estimands	Mean	SD	2.5%	Median	97.5%	Length
E-ASE $_{\pi}$	-0.003	0.006	-0.014	-0.003	0.009	0.023
E-ASE $_{0.1}$	-0.011	0.008	-0.029	-0.011	0.004	0.033
E-ASE $_{0.3}$	-0.009	0.007	-0.023	-0.008	0.005	0.029
E-ASE $_{0.5}$	-0.006	0.007	-0.019	-0.007	0.006	0.025
E-ASE $_{0.7}$	-0.004	0.006	-0.017	-0.004	0.008	0.025
E-ASE $_{0.9}$	-0.002	0.006	-0.015	-0.002	0.01	0.024

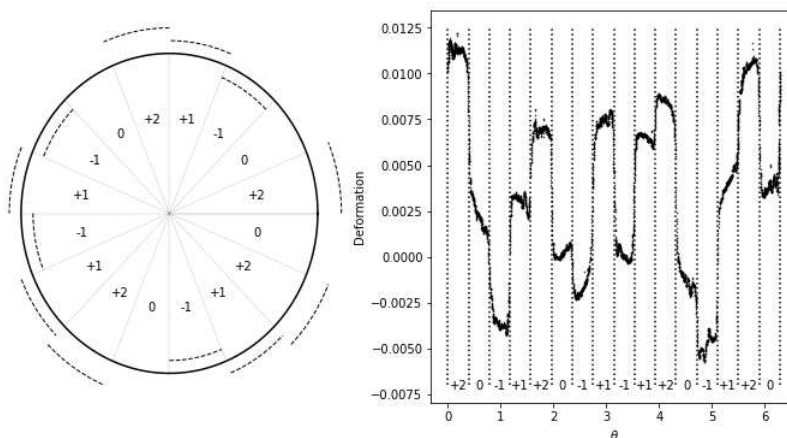


Figure 3: Experimental design (left) and observed deformation (right). Dashed lines on the left figure represent assigned compensations.

investigate different strategies for dimensional accuracy control. One strategy is to implement a compensation plan for shape deviations, which corresponds to a modification of the CAD model such that when the modified model is manufactured the corresponding physical product should exhibit reduced shape deviations compared to the original printed product (Huang et al., 2015).

Sabbaghi et al. (2014) studied shape deviations in a stereolithography process. Their experiment considered in-plane deviations for the top layers of four circular cylinders with negligible heights. The cylinder was of nominal radius 0.5". The objective of their experiment was to determine whether interference would arise in the shape deviations for the points on the cylinders under discretized compensation plans. The experimental units in this context are the individual points on the top boundary of a printed product. The covariates for each unit correspond to the location of the unit on a cylinder under the polar coordinate system (i.e., the angle $\psi_i \in (0, 2\pi)$ in radians). The treatment is a continuous treatment factor, corresponding to the amount of compensation (i.e., the addition or subtraction of material) applied to an experimental unit. Each of the potential outcomes $Y_i(\mathbf{Z})$ for unit i is defined as the difference between the potential and nominal radius at ψ_i under compensation plan \mathbf{Z} applied to the product containing unit i . The experimental design corresponds to a restricted Latin square design on a cylinder. The cylinder was divided into 16 equal-sized sections of $\pi/8$ radians, and one of four levels of the continuous treatment factor was applied to each, with the levels set based on the size of the cylinder. The compensations for the 0.5" cylinder were $-0.004''$, $0''$, $0.004''$, and $0.008''$. The left panel of Figure 3 summarizes the experimental design for the 0.5" cylinder. The right panel of this figure summarizes the shape deviation data of all the measured points.

We proceed to apply our DoI framework to analyze the $N = 6159$ datapoints from the cylinder with nominal radius $r = 0.5''$. This application is challenging in the sense that we are interested in the unit-level effects. We first estimate the unit-level assignment-conditional total effects. These are similar in spirit to the effective treatments considered by Sabbaghi et al. (2014). However, a major difference is that Sabbaghi et al. (2014) defined the estimand in terms of the parameters of their Bayesian model whereas we define the estimand according to the finite-sample perspective, i.e., in terms of observed and missing potential outcomes such as $\tau_i^{\text{tot}}(\mathbf{z}) = Y_i(z_i, \mathbf{z}_{-i}) - Y_i(0, \mathbf{0})$. The left panel of Figure 4 summarizes the inferences on the total effects obtained from our methodology. These inferences match the effective treatments presented in Figure 5 of (Sabbaghi et al., 2014). Considering the fact that their analysis was based on parametric Bayesian models tailored to this particular AM data set, we conclude that our DoI framework and Bayesian approach successfully captures the effective treatment of compensation as the total effect of compensation in a conceptually more straightforward and general manner. In addition, our DDPM is much more flexible than the parametric models considered by Sabbaghi et al. (2014), with the flexibility arising from the infinite mixture of kernels with atoms that depend on \mathbf{A} . Instead of the common choice of the adjacency matrix, we employed an inverse-distance matrix that was defined as

$$A_{ij} = \begin{cases} \frac{1}{|\psi_i - \psi_j|} & \text{if } |\psi_i - \psi_j| \leq \frac{\pi}{8} \\ 0 & \text{otherwise.} \end{cases}$$

The set of neighboring units for unit i is defined as $\mathcal{N}_i = \{j : A_{ij} \neq 0\}$. Our potential outcome model was specified according to $Y_i(z) = \mathbf{X}_i^\top \beta_{Z_i} + G_i(\mathbf{T}_i, \mathbf{Z}_{-i})$, with $\mu(\mathbf{X}_i, \mathbf{A}_i, \mathbf{Z}_{-i}, \gamma_k) =$

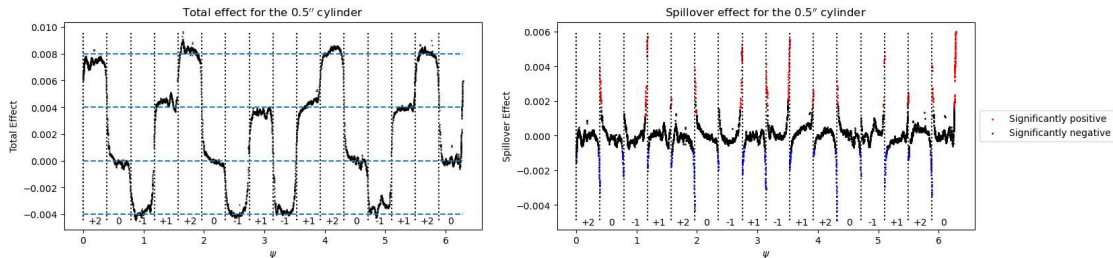


Figure 4: Summary of the total effects (left) and spillover effects (right) for all points on the 0.5'' cylinder as inferred under our DoI framework and Bayesian methodology.

$\gamma_k \sum_{j \in \mathcal{N}_i} Z_j A_{ij}$. Following Sabbaghi et al. (2014) we set $\mathbf{X}_i = (1, \cos \psi_i)^\top$ and specify our prior distributions for the model parameters as $\beta_{Z_i} \sim N((0, 0)^\top, 10\mathbf{I}_2)$, where \mathbf{I}_2 is the 2×2 identity matrix, for $Z_i \in \{-0.004'', 0, 0.004'', 0.008''\}$, $\gamma_k \sim N(0, 10)$, $\sigma_k^2 \sim \text{IG}(1, 10^{-7})$, $\sigma^2 \sim \text{IG}(1, 10^{-7})$.

In addition to the difference in perspectives, our approach possesses additional advantages. First and foremost, we need only use a single data set to estimate the total effect, whereas Sabbaghi et al. (2014) used two distinct sets of experimental data to assess which units on the cylinders have negligible interference. In particular, their first data set consisted of control data from previously manufactured cylinders for which every unit received zero compensation. They computed the posterior distribution of the model parameters and formed the posterior predictive distribution of the potential outcomes based on these control data. After performing these inferences, they then used the data from the compensation experiment to assess interference in terms of how the units' effective treatments differed from that to which they were physically assigned. This approach is similar to that of Hudgens and Halloran (2008) in the sense that they prepared multiple clusters of units by their designs and assessed interference by comparing the outcomes in different clusters. In our approach, we only need the data from the compensated products, and not control data on products that did not have a compensation plan, to infer the effects of interference.

As discussed in §3.2, the unit-level spillover effect is $\tau_i^{\text{SP}}(Z_i, \mathbf{Z}_{-i}, \mathbf{Z}'_{-i}) = Y_i(Z_i, \mathbf{Z}_{-i}) - Y_i(Z_i, \mathbf{Z}'_{-i})$ for two different assignment vectors \mathbf{Z} and \mathbf{Z}' with $Z_i = Z'_i$. Following the definition of negligible interference by Sabbaghi et al. (2014, p. 1403, 1405), we consider spillover effects in which, for a given \mathbf{Z} , we set $\mathbf{Z}' = Z_i \mathbf{1}_{N-1}$, where $\mathbf{1}_{N-1}$ is the $(N-1)$ -dimensional vector of ones. Our unit-level spillover estimand in this case is formally expressed as $\tau_i^{\text{SP}}(Z_i, \mathbf{Z}_{-i}, \mathbf{Z}'_{-i}) = Y_i(Z_i, \mathbf{Z}_{-i}) - Y_i(Z_i, Z_i \mathbf{1}_{N-1})$. The right panel of Figure 4 presents the spillover effect of all units as inferred from our Bayesian method. The central 95% posterior intervals for significantly positive (negative) units are always positive (negative), and never contain zero, which indicates that these units receive nonzero spillover effects. We observe that our Bayesian method indicates that approximately 90% of units, primarily in the central regions of the 16 sections of the cylinders, exhibit negligible interference. This result is consistent with that obtained by Sabbaghi et al. (2014), and was obtained in a more general and straightforward manner compared to Sabbaghi et al. (2014).




Figure 5: Summary of the DoIs for all points on the 0.5'' cylinder.

Another advantage of our Bayesian method compared to the work of Sabbaghi et al. (2014) is the clustering property of the DDPM. Figure 5 summarizes the DoI of all experimental units as inferred from our model, with different colors representing different clusters of units. The clusters of units are determined by the posterior modes of the latent class indicators C_i . Different clusters exhibit different interference structures, which are governed by the atoms of the DDPM model. As seen in this figure, most of the units in the middle of the sections belong to one major cluster (Cluster 6, indicated as green). Units on the edges of the sections belong to other minor clusters and exhibit more significant interference, which corresponds to the fact that in this experiment these units should have different interference structures compared to that of the major cluster.

Each atom of the DDPM model has its own covariate-dependent location parameter, so that each component corresponds to a single parametric model with a different location parameter. As we obtain multiple clusters, it is clear that modeling the interference structure with a single parametric model, as was done by Sabbaghi et al. (2014), is not appropriate. Our methodology automatically captures the interference structure and separates the units into different clusters based on the interference effects that they receive. Of the total $N = 6159$ units, 4458 (or approximately 70%) belong to the major cluster. Clusters 3, 4 and 8 appear only on the edge of the sections in which the absolute difference in their compensation level versus that of the neighboring section is $0.008''$. As such, units in clusters 3, 4 and 8 are highly affected by other units through interference. Finally, we see that the cluster membership changes depending on the distance of a unit from the edge of the section, which gradates the interference structures within sections.

7. Concluding Remarks

This article proposed our new notion of the Degree of Interference (DoI) that flexibly handles interference. We presented several new types of causal estimands and developed a fast data augmentation based Bayesian semiparametric methodology for estimating them. In simulation studies across different scenarios with complicated interference structures

our approach exhibited better finite-sample performance than that of the conventional HT estimator in terms of MSE. The results provided in Appendix B also highlight the superiority of our methodology in terms of bias and coverage probability. Our two case studies demonstrate that our DoI framework and Bayesian methodology yield more insights and knowledge on interference in a straightforward manner.

Our methodology should complement existing methodologies rather than replace them. For instance, while our approach introduces modeling assumptions, existing methods (Sävje, 2023; Hoshino and Yanagi, 2023) are fully nonparametric, which may appeal to analysts reluctant to make modeling assumptions. Conversely, analysts may find the conservative variance estimate provided by nonparametric methods too large and seek efficiency gains through our proposed method.

A possible future direction to build upon this framework would be to extend our framework to observational studies where the treatment assignment mechanisms take unknown functional forms.

Appendix A. Detailed Steps of Gibbs Sampler

In this section we present the detailed steps of the Gibbs sampler developed in §4.5. First, our implementation of the Gibbs sampler is outlined below.

1. Initialize parameters $\boldsymbol{\phi}^{(0)}, \boldsymbol{\theta}^{(0)}$.
2. Repeat the following steps. For $t = 0, 1, \dots$
 - (a) Draw $\mathbf{G}^{\text{o},(t+1)} \sim p(\mathbf{G} \mid \mathbf{T}, \mathbf{Y}^{\text{obs}}, \boldsymbol{\phi}^{(t)}, \boldsymbol{\theta}^{(t)})$.
 - (b) Draw $(\boldsymbol{\phi}^{(t+1)}, \boldsymbol{\theta}^{(t+1)}) \sim p(\boldsymbol{\phi}, \boldsymbol{\theta} \mid \mathbf{T}, \mathbf{Y}^{\text{obs}}, \mathbf{G}^{\text{o},(t+1)})$.
 - (c) Draw $\mathbf{G}^{\text{u},(t+1)} \sim p(\mathbf{G} \mid \mathbf{T}, \mathbf{Y}^{\text{obs}}, \mathbf{G}^{\text{o},(t+1)}, \boldsymbol{\phi}^{(t+1)}, \boldsymbol{\theta}^{(t+1)})$.
 - (d) Draw $\mathbf{Y}^{\text{mis},(t+1)} \sim p(\mathbf{Y}^{\text{mis}} \mid \mathbf{T}, \mathbf{Y}^{\text{obs}}, \mathbf{G}^{\text{o},(t+1)}, \mathbf{G}^{\text{u},(t+1)}, \boldsymbol{\phi}^{(t+1)}, \boldsymbol{\theta}^{(t+1)})$.
 - (e) Calculate the causal estimand based on the observed and imputed data.

The detailed steps are provided below.

1. For the units with $C_i = k$ and $Z_i = z$, draw G_i^o from $N(\{\lambda_z(\gamma_0 + \gamma_k \sum_{j \in \mathcal{N}_i} Z_j A_{ij}) + (Y_i - X_i \beta_z) \sigma_k^2\} / (\lambda_z + \sigma_k^2), \lambda_z \sigma_k^2 / (\lambda_z + \sigma_k^2))$
2. For each unit, draw C_i from $P(C_i = k \mid -) = \frac{w_k N(G_i \mid \gamma_0 + \gamma_k \sum_{j \in \mathcal{N}_i} Z_j A_{ij}, \sigma_k^2)}{\sum_{k=1}^K w_k N(G_i \mid \gamma_0 + \gamma_k \sum_{j \in \mathcal{N}_i} Z_j A_{ij}, \sigma_k^2)}$
3. Update w_k 's as Step 3 in §4.5.
4. We adopt the Metropolis Hasting step with a proposal draw $\alpha^* \sim \text{Gamma}(1, 1)$. Given \mathbf{w}' and \mathbf{C} , we accept the draw α^* with probability r where

$$r = \min \left\{ 1, \frac{p(\alpha^* \mid 1, 1) \prod_{k=1}^K f(w'_k \mid 1 + \sum_{i: C_i=k} 1, \alpha^* + \sum_{i: C_i > k} 1)}{p(\alpha \mid 1, 1) \prod_{k=1}^K f(w'_k \mid 1 + \sum_{i: C_i=k} 1, \alpha + \sum_{i: C_i > k} 1)} \right\}.$$

We reject the draw α^* with probability $1 - r$ and accept the draw α , obtained in the previous iteration. Note that $p(\alpha \mid a, b)$ denotes the prior distribution of α , which we specify by the Gamma distribution with parameters $a = 1, b = 1$ and f is specified in Step 4 in §4.5.

5. (a) If $N_k = \sum_{i=1}^N \mathbb{1}(C_i = k) > 0$, draw σ_k^2 from $\text{IG}(2 + 0.5N_k, 1 + 0.5s_k)$ where $s_k = \sum_{i:C_i=k} (G_i^o - \gamma_0 - \gamma_k \sum_{j \in \mathcal{N}_i} Z_j A_{ij})^2$. If $N_k = 0$, then draw σ_k^2 from the prior $\text{IG}(2, 1)$.
 - (b) If $N_k > 0$, draw γ_k from $\text{N}((\mathbf{M}_k^\top \mathbf{M}_k + \sigma_k^2/10 \times I_2)^{-1} \mathbf{M}_k^\top \mathbf{G}_k^o, \sigma_k^2(\mathbf{M}_k^\top \mathbf{M}_k + \sigma_k^2/10 \times I_2)^{-1})$ where \mathbf{G}_k^o is a subvector of \mathbf{G}^o for units with $C_i = k$ and \mathbf{M}_k is an $N_k \times 2$ matrix with the row vector $(1, \sum_{j \in \mathcal{N}_i} Z_j A_{ij})$ for units with $C_i = k$. If $N_k = 0$, then draw from $\text{N}((0, 0)^\top, \sigma_k^2/10 \times I_2)$.
6. (a) Draw λ_z from $\text{IG}(2 + N_z/2, 1 + 0.5 \sum_{i:Z_i=z} (Y_i - X_i \beta_z - G_i^o)^2)$.
 - (b) Draw β_z from $\text{N}((\mathbf{X}_z^\top \mathbf{X}_z + \lambda_z/10 \times I_2)^{-1} \mathbf{X}_z^\top (\mathbf{Y}_z - \mathbf{G}_z^o), \lambda_z(\mathbf{X}_z^\top \mathbf{X}_z + \lambda_z/10 \times I_2)^{-1})$, where \mathbf{G}_z^o , \mathbf{X}_z and \mathbf{Y}_z are respectively a submatrix and a subvector of \mathbf{G}^o , \mathbf{X} and \mathbf{Y} for units with $Z_i = z$.
- 7-10. Follow the procedures in §4.5 for assignment-conditional effects or the steps provided in the next section for the expected effects, depending on which of the effects (conditional effects or expected effects) one is interested in.

A.1 Estimation of Expected Effects

We present an extension to infer the expected effects. It suffices to derive the procedures for estimating $\mathbb{E}_{\mathbf{z}_{-i} \sim \pi} [Y_i(z, \mathbf{z}_{-i})]$ where π is the marginalized distribution of \mathbf{z}_{-i} , which depends on the experimental design. For notational convenience, the outer integral below is taken with respect to the counting measure π for the distribution of \mathbf{z}_{-i} .

$$\begin{aligned} \mathbb{E}_{\mathbf{z}_{-i} \sim \pi} [Y_i(z, \mathbf{z}_{-i})] &= \mathbb{E}_{\mathbf{z}_{-i} \sim \pi} \left[Y_i(z, \mathbf{z}_{-i}) \int_{\mathcal{G}} dF_{G_i | \mathbf{z}_{-i} = \mathbf{z}_{-i}}(g) \right] = \mathbb{E}_{\mathbf{z}_{-i} \sim \pi} \left[\int_{\mathcal{G}} Y_i(z, \mathbf{z}_{-i}) dF_{G_i | \mathbf{z}_{-i} = \mathbf{z}_{-i}}(g) \right] \\ &= \int_{\Omega^{N-1}} \int_{\mathcal{G}} Y_i(z, g) dF_{G_i | \mathbf{z}_{-i} = \mathbf{z}_{-i}}(g) d\pi(\mathbf{z}_{-i}) \approx \frac{1}{M} \sum_{m=1}^M Y_i(z, g_i^{(m)}) \end{aligned}$$

$g_i^{(m)}$ is obtained by the following two steps. For $m = 1, \dots, M$, (1) $\mathbf{z}^m \sim \pi$, and then (2) $g_i^{(m)} \sim F_{G_i | \mathbf{z}_{-i} = \mathbf{z}_{-i}^m}(g)$ where \mathbf{z}_{-i}^m is a subvector of \mathbf{z}^m with the i -th entry removed. We infer the expected effects by following the same steps from step 1 to step 6 in §4.5, and then we proceed as:

7. Draw $\mathbf{z}^* \sim \pi$.
8. Given \mathbf{w} in step 3, draw C_i^* from $C_i^* \sim \text{MN}(\mathbf{w})$.
9. (a) Given C_i^* , γ and σ^2 , draw G_i^* from $G_i^* \sim \text{N}(\mu(\mathbf{T}_i, \mathbf{z}_{-i}^*, \gamma_{C_i^*}), \sigma_{C_i^*}^2)$.
 - (b) Given C_i^* , γ and σ^2 , draw G_i^0 from $G_i^0 \sim \text{N}(\mu(\mathbf{T}_i, \mathbf{0}^{N-1}, \gamma_{C_i^*}), \sigma_{C_i^*}^2)$.
10. (a) Given G_i^* in step 9(a) and θ in step 6, draw from $Y_i(z, \mathbf{z}_{-i}^*) \sim P(Y_i | Z_i = z, \mathbf{T}_i, G_i^*, \theta)$ for $z = 0, 1$.
 - (b) Given G_i^0 in step 9(b) and θ in step 6, draw from $Y_i(0, \mathbf{0}^{N-1}) \sim P(Y_i | Z_i = 0, \mathbf{T}_i, G_i^0, \theta)$.

A.2 Extention to Binary Outcomes

When we adopt a Probit model as in our illustrative analyses in §6.1, we draw the sample of β_z as follows. Given $G_i^o = g_i$ and β_z in the previous iteration, we first draw an auxiliary

variable v_i for unit i with $Z_i = z$ from

$$v_i \sim \begin{cases} \text{TN}(X_i\beta_z + g_i, 1, 0, \infty) & \text{if } Y_i = 1, \\ \text{TN}(X_i\beta_z + g_i, 1, -\infty, 0) & \text{if } Y_i = 0, \end{cases}$$

where $\text{TN}(\mu, \sigma^2, l, u)$ denotes the truncated normal distribution with the mean, variance, lower bound, and upper bound parameters. Then, we draw from $\beta_z \sim \text{N}(M, V)$, where $V = (H_0^{-1} + \mathbf{X}_z^\top \mathbf{X}_z)^{-1}$, $M = V(H_0^{-1}h_0 + \mathbf{X}_z^\top \mathbf{v}_z)$ with \mathbf{v}_z being a subvector of $\mathbf{v} = (v_1, \dots, v_N)$ for units with $Z_i = z$. The prior distribution of β_z is specified here by $\beta_z \sim \text{N}(h_0, H_0)$.

Appendix B. Simulation Results

This section presents the simulation evaluation not presented in §5; bias and coverage for E-ATE, and bias, MSE, and coverage for E-ASE. We consider the same data-generating mechanisms as in §5. Table 3 - 8 present the results.

For E-ATE, the DoI approach outperforms the HT estimator in terms of MSE and bias for almost all scenarios of both networks, even under severe model misspecifications (Scenario 2 – 5), and in the presence of an unmeasured covariate. The MSE decreases as N increases, and the coverage probability is well-calibrated with 95% probability. These observations demonstrate the efficiency and robustness of our DoI approach. The HT estimator sometimes exhibits large bias when N is small, whereas the Bayesian methodology consistently works well. Both methodologies exhibit smaller bias as N increases. They are well-calibrated in that they yield intervals with nearly 95% coverage across nearly all conditions for the E-ATE, indicating the robustness of our methodologies for misspecifications. Another advantage of the DoI methodology is that it provides a principled estimation procedure for the spillover effects. Overall, the DoI methodology yields good bias, MSE, and coverage for the E-ASE across all scenarios.

DEGREE OF INTERFERENCE

Table 3: Evaluation metrics of simulations studies (Scenario 1). The data-generating process (DGP) and the underlying network between units are provided in §5.

DGP	Network (parameters)	N	HT (E-ATE)			DoI (E-ATE)			DoI (E-ASE)		
			Bias	MSE	Coverage	Bias	MSE	Coverage	Bias	MSE	Coverage
Scenario 1	ER($p = 0.01$)	300	-0.0061	0.0513	0.99	0.002	0.0136	0.97	-0.0013	0.0074	0.97
		500	-0.0055	0.0321	0.975	-0.0018	0.0086	0.95	0.0019	0.0077	0.99
		1000	-0.0113	0.0156	0.98	-0.0032	0.0033	0.99	0.0018	0.0068	0.995
	ER($p = 0.03$)	300	-0.0052	0.0515	0.99	0.0025	0.0133	0.965	-0.0133	0.0205	0.98
		500	-0.0064	0.0319	0.98	-0.0017	0.0088	0.96	-0.0096	0.019	0.975
		1000	-0.0117	0.0155	0.98	-0.0029	0.0033	0.995	-0.0144	0.0242	0.965
	ER($p = 0.05$)	300	-0.005	0.0512	0.99	0.0028	0.0134	0.955	-0.0037	0.0347	0.975
		500	-0.0057	0.032	0.98	-0.0018	0.0087	0.955	-0.0377	0.034	0.985
		1000	-0.0115	0.0155	0.98	-0.0027	0.0033	0.995	-0.0179	0.0338	0.975
	BA($n_0 = 10, k = 3$)	300	-0.0056	0.0522	0.99	0.0027	0.0133	0.975	-0.0051	0.0098	0.98
		500	-0.0061	0.0318	0.98	-0.0016	0.0087	0.95	-0.0104	0.0071	0.985
		1000	-0.0116	0.0154	0.98	-0.0026	0.0033	0.99	-0.0035	0.0039	0.97
	BA($n_0 = 10, k = 5$)	300	-0.0058	0.0513	0.995	0.0009	0.0134	0.965	-0.0285	0.0185	0.975
		500	-0.0065	0.0316	0.98	-0.0017	0.0087	0.955	-0.0056	0.0129	0.985
		1000	-0.0116	0.0155	0.98	-0.003	0.0033	0.99	-0.0088	0.0074	0.945
	BA($n_0 = 30, k = 3$)	300	-0.0052	0.0513	0.99	0.0022	0.0136	0.96	-0.0182	0.0104	0.965
		500	-0.0059	0.0322	0.975	-0.0022	0.0087	0.95	-0.0075	0.006	0.975
		1000	-0.0115	0.0155	0.98	-0.0026	0.0034	0.99	-0.0054	0.0042	0.95
	BA($n_0 = 30, k = 5$)	300	-0.0045	0.0511	0.99	0.0021	0.0134	0.96	-0.005	0.0115	0.985
		500	-0.006	0.0318	0.98	-0.0013	0.0087	0.955	0.0012	0.0078	0.975
		1000	-0.0118	0.0155	0.98	-0.003	0.0033	0.995	0.0006	0.007	0.945

Table 4: Evaluation metrics of simulations studies (Scenario 2).

DGP	Network (parameters)	N	HT (E-ATE)			DoI (E-ATE)			DoI (E-ASE)		
			Bias	MSE	Coverage	Bias	MSE	Coverage	Bias	MSE	Coverage
Scenario 2	ER($p = 0.01$)	300	-0.0047	0.0508	0.99	0.0031	0.0134	0.97	0.0044	0.0074	0.96
		500	-0.0058	0.0309	0.98	-0.0015	0.0086	0.955	0.0053	0.0077	0.985
		1000	-0.0106	0.0152	0.98	-0.0026	0.0033	0.995	0.005	0.0069	0.995
	ER($p = 0.03$)	300	-0.0046	0.051	0.99	0.0019	0.0131	0.965	-0.0067	0.0209	0.98
		500	-0.0059	0.0311	0.98	-0.0005	0.0087	0.95	-0.0021	0.0198	0.985
		1000	-0.0111	0.0152	0.98	-0.0025	0.0033	0.995	-0.0079	0.0248	0.95
	ER($p = 0.05$)	300	-0.0048	0.051	0.99	0.003	0.0134	0.955	0.0054	0.035	0.975
		500	-0.0058	0.0312	0.98	-0.0018	0.0085	0.95	-0.0275	0.0337	0.985
		1000	-0.0111	0.0153	0.98	-0.003	0.0033	0.99	-0.0098	0.0347	0.965
	BA($n_0 = 10, k = 3$)	300	-0.0048	0.0511	0.99	0.0022	0.0131	0.97	0.0001	0.01	0.98
		500	-0.006	0.0305	0.98	-0.0014	0.0086	0.95	-0.0071	0.007	0.98
		1000	-0.0109	0.0151	0.98	-0.003	0.0033	0.985	-0.0008	0.004	0.97
	BA($n_0 = 10, k = 5$)	300	-0.0052	0.051	0.99	0.0007	0.0134	0.96	-0.0204	0.0182	0.975
		500	-0.0062	0.0304	0.98	-0.0015	0.0086	0.955	-0.0011	0.013	0.98
		1000	-0.0109	0.0151	0.98	-0.003	0.0033	0.995	-0.0061	0.0075	0.95
	BA($n_0 = 30, k = 3$)	300	-0.0049	0.051	0.99	0.0032	0.0136	0.955	-0.0136	0.0103	0.97
		500	-0.0061	0.0305	0.98	-0.0011	0.0086	0.95	-0.0046	0.006	0.975
		1000	-0.0108	0.015	0.98	-0.0027	0.0033	0.99	-0.0031	0.0042	0.945
	BA($n_0 = 30, k = 5$)	300	-0.0048	0.0509	0.99	0.0027	0.0132	0.965	0.0009	0.0116	0.98
		500	-0.006	0.0305	0.98	-0.0016	0.0088	0.955	0.0058	0.0079	0.985
		1000	-0.0109	0.0151	0.98	-0.0023	0.0032	0.995	0.0032	0.0071	0.925

Table 5: Evaluation metrics of simulations studies (Scenario 3).

DGP	Network (parameters)	N	HT (E-ATE)			DoI (E-ATE)			DoI (E-ASE)		
			Bias	MSE	Coverage	Bias	MSE	Coverage	Bias	MSE	Coverage
Scenario 3	ER($p = 0.01$)	300	0.0033	0.0678	0.91	0.0034	0.0132	0.975	0.0049	0.0074	0.97
		500	0.0048	0.0333	0.945	-0.0017	0.0086	0.96	0.0089	0.0078	0.975
		1000	-0.0029	0.0168	0.95	-0.0022	0.0033	0.99	0.01	0.0069	0.995
	ER($p = 0.03$)	300	0.0034	0.066	0.915	0.0038	0.0133	0.965	-0.0017	0.021	0.975
		500	0.0045	0.0329	0.945	-0.0018	0.0088	0.95	0.0031	0.0194	0.975
		1000	-0.0035	0.0167	0.955	-0.0024	0.0033	0.99	-0.0036	0.0246	0.95
	ER($p = 0.05$)	300	0.0032	0.066	0.915	0.0027	0.0133	0.965	0.0095	0.0349	0.975
		500	0.0046	0.0328	0.945	-0.0013	0.0087	0.95	-0.0223	0.0335	0.99
		1000	-0.0035	0.0166	0.95	-0.0031	0.0033	0.995	-0.0045	0.0347	0.965
	BA($n_0 = 10, k = 3$)	300	0.0032	0.0675	0.91	0.0029	0.013	0.975	0.003	0.0099	0.98
		500	0.0044	0.0336	0.94	-0.0017	0.0087	0.95	-0.0051	0.007	0.975
		1000	-0.0034	0.0173	0.955	-0.0021	0.0033	0.995	0.0004	0.004	0.97
	BA($n_0 = 10, k = 5$)	300	0.0027	0.0676	0.915	0.0014	0.0133	0.975	-0.0188	0.0183	0.98
		500	0.0042	0.0336	0.94	-0.0017	0.0088	0.95	0.0011	0.013	0.975
		1000	-0.0034	0.0173	0.955	-0.0025	0.0032	0.995	-0.0041	0.0075	0.94
	BA($n_0 = 30, k = 3$)	300	0.0031	0.0675	0.91	0.0039	0.0136	0.955	-0.0111	0.0102	0.965
		500	0.0043	0.0336	0.945	-0.0012	0.0087	0.95	-0.0025	0.006	0.965
		1000	-0.0033	0.0172	0.955	-0.0026	0.0033	0.99	-0.0016	0.0042	0.95
	BA($n_0 = 30, k = 5$)	300	0.0031	0.0675	0.91	0.0021	0.0131	0.97	0.0027	0.0115	0.98
		500	0.0044	0.0336	0.94	-0.001	0.0087	0.955	0.0074	0.008	0.985
		1000	-0.0033	0.0172	0.955	-0.003	0.0033	0.995	0.0052	0.0071	0.925

Table 6: Evaluation metrics of simulations studies (Scenario 4).

DGP	Network (parameters)	N	HT (E-ATE)			DoI (E-ATE)			DoI (E-ASE)		
			Bias	MSE	Coverage	Bias	MSE	Coverage	Bias	MSE	Coverage
Scenario 4	ER($p = 0.01$)	300	-0.0049	0.0508	0.99	0.0034	0.0133	0.97	0.0007	0.0074	0.96
		500	-0.0062	0.0302	0.98	-0.0016	0.0087	0.95	-0.0064	0.0078	0.98
		1000	-0.0114	0.0152	0.975	-0.0032	0.0034	0.99	-0.0129	0.0069	0.985
	ER($p = 0.03$)	300	-0.0056	0.0507	0.985	0.0011	0.0136	0.955	-0.0235	0.0212	0.98
		500	-0.007	0.0298	0.98	-0.0021	0.009	0.95	-0.0192	0.0194	0.975
		1000	-0.0117	0.0152	0.975	-0.0033	0.0035	0.99	-0.0311	0.0252	0.965
	ER($p = 0.05$)	300	-0.0053	0.0503	0.99	0.0026	0.0137	0.965	-0.0128	0.0346	0.965
		500	-0.0072	0.0298	0.98	-0.0025	0.009	0.95	-0.048	0.0353	0.98
		1000	-0.0121	0.0154	0.975	-0.0036	0.0036	0.99	-0.0357	0.0349	0.975
	BA($n_0 = 10, k = 3$)	300	-0.0053	0.0509	0.99	0.0019	0.0133	0.97	-0.0061	0.0098	0.98
		500	-0.0063	0.0301	0.98	-0.0004	0.0086	0.95	-0.0137	0.0071	0.98
		1000	-0.0111	0.0151	0.98	-0.0027	0.0033	0.995	-0.0054	0.004	0.965
	BA($n_0 = 10, k = 5$)	300	-0.0055	0.0506	0.99	0.0007	0.0133	0.975	-0.0271	0.0187	0.975
		500	-0.0065	0.0301	0.98	-0.0026	0.0086	0.945	-0.0072	0.0129	0.985
		1000	-0.0112	0.0151	0.98	-0.0029	0.0034	0.995	-0.0099	0.0074	0.95
	BA($n_0 = 30, k = 3$)	300	-0.0051	0.0509	0.985	0.003	0.0136	0.965	-0.0208	0.0106	0.965
		500	-0.0063	0.0302	0.98	-0.0011	0.0087	0.945	-0.0105	0.0061	0.965
		1000	-0.0108	0.015	0.98	-0.0029	0.0033	0.995	-0.007	0.0043	0.95
	BA($n_0 = 30, k = 5$)	300	-0.005	0.0507	0.99	0.0021	0.0132	0.97	-0.0065	0.0116	0.985
		500	-0.0063	0.0302	0.98	-0.002	0.0088	0.955	-0.0001	0.008	0.98
		1000	-0.0111	0.0151	0.98	-0.0034	0.0034	0.995	-0.0008	0.007	0.94

DEGREE OF INTERFERENCE

Table 7: Evaluation metrics of simulations studies (Scenario 5).

DGP	Network (parameters)	N	HT (E-ATE)			DoI (E-ATE)			DoI (E-ASE)		
			Bias	MSE	Coverage	Bias	MSE	Coverage	Bias	MSE	Coverage
Scenario 5	ER($p = 0.01$)	300	-0.0047	0.0508	0.99	0.0029	0.013	0.97	0.0041	0.0073	0.97
		500	-0.0058	0.0308	0.98	-0.0011	0.0086	0.955	0.0059	0.0078	0.98
		1000	-0.0107	0.0152	0.98	-0.0027	0.0032	0.99	0.0036	0.0068	0.99
	ER($p = 0.03$)	300	-0.0048	0.0511	0.99	0.0023	0.0133	0.97	-0.0067	0.0207	0.98
		500	-0.0063	0.0308	0.98	-0.0018	0.0087	0.95	-0.0027	0.0193	0.97
		1000	-0.0113	0.0153	0.98	-0.0035	0.0033	0.985	-0.0088	0.0244	0.96
	ER($p = 0.05$)	300	-0.0049	0.0511	0.99	0.0026	0.0133	0.97	0.0056	0.0351	0.975
		500	-0.0062	0.0309	0.98	-0.0016	0.0086	0.955	-0.0275	0.0336	0.985
		1000	-0.0115	0.0154	0.98	-0.003	0.0033	0.995	-0.0099	0.034	0.97
	BA($n_0 = 10, k = 3$)	300	-0.0049	0.0511	0.99	0.0026	0.0132	0.965	0.0008	0.0098	0.98
		500	-0.006	0.0305	0.98	-0.0009	0.0086	0.955	-0.0074	0.007	0.985
		1000	-0.0109	0.0151	0.98	-0.0027	0.0033	0.995	-0.001	0.0039	0.965
	BA($n_0 = 10, k = 5$)	300	-0.0052	0.051	0.99	-0.0	0.0134	0.965	-0.0201	0.0182	0.975
		500	-0.0063	0.0304	0.98	-0.002	0.0085	0.95	-0.0011	0.0129	0.985
		1000	-0.011	0.0151	0.98	-0.0032	0.0033	0.99	-0.0063	0.0075	0.945
	BA($n_0 = 30, k = 3$)	300	-0.0049	0.051	0.99	0.0032	0.0137	0.96	-0.0131	0.0104	0.975
		500	-0.0061	0.0305	0.98	-0.0017	0.0087	0.94	-0.0046	0.006	0.965
		1000	-0.0108	0.015	0.98	-0.0024	0.0033	0.995	-0.0029	0.0042	0.955
	BA($n_0 = 30, k = 5$)	300	-0.0048	0.0509	0.99	0.0019	0.0131	0.965	-0.0008	0.0116	0.985
		500	-0.0061	0.0305	0.98	-0.0013	0.0087	0.945	0.0063	0.008	0.98
		1000	-0.0109	0.0151	0.98	-0.0032	0.0033	0.98	0.0032	0.007	0.93

Table 8: Evaluation metrics of simulations studies (Scenario 6).

DGP	Network (parameters)	N	HT (E-ATE)			DoI (E-ATE)			DoI (E-ASE)		
			Bias	MSE	Coverage	Bias	MSE	Coverage	Bias	MSE	Coverage
Scenario 6	ER($p = 0.01$)	300	0.0068	0.0775	0.99	0.008	0.0279	0.99	-0.0062	0.0087	0.985
		500	0.0004	0.0468	0.98	-0.0016	0.0172	0.99	0.0039	0.0103	0.98
		1000	-0.0099	0.0213	0.975	0.0017	0.0076	0.995	0.0069	0.0132	0.955
	ER($p = 0.03$)	300	0.0065	0.0777	0.99	0.0053	0.0277	0.99	-0.0128	0.0302	0.985
		500	0.0006	0.047	0.99	-0.0018	0.0172	0.99	-0.0035	0.0325	0.97
		1000	-0.01	0.0214	0.975	0.0025	0.0076	1.0	-0.0222	0.0339	0.975
	ER($p = 0.05$)	300	0.0068	0.0783	0.99	0.0076	0.0283	0.995	-0.0361	0.0398	0.98
		500	0.0007	0.047	0.99	-0.0009	0.0176	0.99	-0.0221	0.041	0.985
		1000	-0.01	0.0214	0.975	0.0025	0.0077	1.0	-0.0293	0.0547	0.965
	BA($n_0 = 10, k = 3$)	300	0.0066	0.0779	0.99	0.0088	0.0279	1.0	-0.0034	0.0182	0.97
		500	0.001	0.0466	0.99	-0.0014	0.0176	0.99	0.0095	0.0116	0.97
		1000	-0.0098	0.021	0.975	0.0016	0.0076	1.0	0.0013	0.0065	0.97
	BA($n_0 = 10, k = 5$)	300	0.0068	0.078	0.99	0.0077	0.028	0.99	-0.0236	0.0265	0.975
		500	0.0012	0.0469	0.99	-0.0005	0.0173	0.99	-0.0016	0.0222	0.95
		1000	-0.0092	0.0212	0.975	0.0026	0.0076	1.0	0.0007	0.0094	0.99
	BA($n_0 = 30, k = 3$)	300	0.0068	0.0779	0.99	0.0074	0.0283	0.985	-0.0057	0.0124	0.98
		500	0.0011	0.0466	0.985	-0.0019	0.0176	0.99	-0.0117	0.0092	0.975
		1000	-0.0095	0.0211	0.975	0.0023	0.0077	1.0	-0.0043	0.0062	0.97
	BA($n_0 = 30, k = 5$)	300	0.0071	0.0772	0.99	0.0066	0.0282	0.985	-0.0079	0.0192	0.975
		500	0.0011	0.0468	0.985	-0.0023	0.0174	0.99	-0.0217	0.014	0.98
		1000	-0.0098	0.0213	0.975	0.002	0.0077	1.0	-0.0054	0.0086	0.97

Table 9: Evaluation metrics of simulations studies (Scenario 7).

DGP	Network (parameters)	N	HT (E-ATE)			DoI (E-ATE)			DoI (E-ASE)		
			Bias	MSE	Coverage	Bias	MSE	Coverage	Bias	MSE	Coverage
Scenario 7	ER($p = 0.01$)	300	0.0486	0.0958	1.0	0.0074	0.013	0.985	0.0008	0.0064	0.99
		500	-0.0267	0.0684	0.96	0.0039	0.0088	0.955	0.0041	0.0069	0.975
		1000	0.0105	0.0263	0.995	0.0038	0.0039	0.955	-0.0007	0.0101	0.955
	ER($p = 0.03$)	300	0.0476	0.0953	1.0	0.0073	0.0129	0.98	-0.0215	0.0229	0.96
		500	-0.0269	0.0684	0.96	0.004	0.0089	0.94	-0.0147	0.0198	0.98
		1000	0.0105	0.0263	0.995	0.0045	0.0038	0.965	0.0059	0.021	0.975
	ER($p = 0.05$)	300	0.0471	0.0955	1.0	0.0086	0.0128	0.985	-0.0039	0.0355	0.975
		500	-0.027	0.0683	0.96	0.0038	0.009	0.935	-0.0363	0.0338	0.975
		1000	0.0108	0.0265	0.995	0.004	0.0038	0.965	-0.015	0.0373	0.965
	BA($n_0 = 10, k = 3$)	300	0.0478	0.0952	1.0	0.0082	0.0129	0.985	-0.002	0.011	0.975
		500	-0.027	0.0683	0.96	0.0028	0.0088	0.955	-0.0052	0.0071	0.97
		1000	0.0105	0.0263	0.995	0.0038	0.0039	0.96	0.0062	0.0036	0.975
	BA($n_0 = 10, k = 5$)	300	0.0475	0.0956	1.0	0.007	0.0131	0.98	-0.0151	0.0192	0.96
		500	-0.0271	0.0685	0.96	0.0033	0.0089	0.945	-0.0054	0.0125	0.975
		1000	0.0106	0.0264	0.995	0.0049	0.0038	0.965	0.0017	0.0066	0.935
	BA($n_0 = 30, k = 3$)	300	0.0485	0.0955	1.0	0.0071	0.0128	0.98	-0.0145	0.0092	0.98
		500	-0.0273	0.0681	0.96	0.0034	0.0089	0.945	-0.0075	0.006	0.97
		1000	0.0105	0.0264	0.995	0.004	0.0038	0.96	-0.0025	0.0036	0.97
	BA($n_0 = 30, k = 5$)	300	0.0479	0.0957	1.0	0.0074	0.013	0.985	-0.0068	0.0129	0.955
		500	-0.0275	0.0682	0.96	0.0026	0.0089	0.955	-0.008	0.0105	0.955
		1000	0.0108	0.0262	0.995	0.0037	0.0038	0.965	-0.0117	0.0062	0.945

Table 10: Sensitivity Analysis for different prior specifications (Prior 1).

Estimands	Mean	SD	2.5%	Median	97.5%	Length
E-ATE $_{\pi}$	0.011	0.007	-0.002	0.011	0.026	0.028
E-ASE $_{\pi}$	-0.004	0.006	-0.016	-0.004	0.008	0.024
E-ASE $_{\pi}^{nb,1}$	-0.004	0.006	-0.016	-0.004	0.009	0.026
E-ASE $_{\pi}^{nb,2}$	-0.002	0.018	-0.036	-0.004	0.033	0.069
E-ASE $_{\pi}^{rt,33}$	-0.003	0.027	-0.055	0.0	0.05	0.105
E-ASE $_{\pi}^{rt,50}$	-0.003	0.008	-0.019	-0.004	0.013	0.033
E-ASE $_{\pi}^{rt,66}$	-0.0	0.03	-0.058	0.0	0.06	0.118

Appendix C. Sensitivity Analyses for the Case Study in the Main Text

Throughout the manuscript, we use proper, although weakly informative, conjugate prior distributions for all parameters. In this section, we examine the robustness of the results concerning the prior specifications. We present sensitivity analyses with the following prior specifications: Table 11–13 report the results with the following priors:

Prior 1: $\beta_z, \sim N((0, 0)^\top, 20^2 \mathbf{I}_2)$, $\gamma_{k,1}, \gamma_{k,2} \sim N(0, 20^2)$, $\sigma_k^2 \sim IG(0.1, 0.1)$.

Prior 2: $\beta_z, \sim N((0, 0)^\top, 5^2 \mathbf{I}_2)$, $\gamma_{k,1}, \gamma_{k,2} \sim N(0, 5^2)$, $\sigma_k^2 \sim IG(0.1, 0.1)$.

Prior 3: $\beta_z, \sim N((0, 0)^\top, 10^2 \mathbf{I}_2)$, $\gamma_{k,1}, \gamma_{k,2} \sim N(0, 10^2)$, $\sigma_k^2 \sim IG(0.01, 0.01)$.

Prior 4: $\beta_z, \sim N((0, 0)^\top, 10^2 \mathbf{I}_2)$, $\gamma_{k,1}, \gamma_{k,2} \sim N(0, 10^2)$, $\sigma_k^2 \sim IG(1.0, 1.0)$.

Compared with Table 1, the results change only slightly, validating the robustness of our analyses.

Table 11: Sensitivity Analysis for different prior specifications (Prior 2).

Estimands	Mean	SD	2.5%	Median	97.5%	Length
E-ATE $_{\pi}$	0.01	0.006	-0.002	0.009	0.023	0.025
E-ASE $_{\pi}$	-0.0	0.006	-0.011	-0.0	0.012	0.023
E-ASE $_{\pi}^{nb,1}$	-0.0	0.006	-0.012	0.0	0.013	0.024
E-ASE $_{\pi}^{nb,2}$	0.002	0.017	-0.033	0.0	0.036	0.069
E-ASE $_{\pi}^{rt,33}$	0.002	0.026	-0.048	0.0	0.054	0.101
E-ASE $_{\pi}^{rt,50}$	0.0	0.008	-0.015	0.0	0.016	0.031
E-ASE $_{\pi}^{rt,66}$	0.003	0.03	-0.052	0.0	0.065	0.117

Table 12: Sensitivity Analysis for different prior specifications (Prior 3).

Estimands	Mean	SD	2.5%	Median	97.5%	Length
E-ATE $_{\pi}$	0.011	0.007	-0.002	0.011	0.026	0.028
E-ASE $_{\pi}$	-0.003	0.006	-0.015	-0.003	0.01	0.025
E-ASE $_{\pi}^{nb,1}$	-0.003	0.007	-0.015	-0.003	0.01	0.025
E-ASE $_{\pi}^{nb,2}$	-0.0	0.018	-0.033	0.0	0.036	0.069
E-ASE $_{\pi}^{rt,33}$	-0.001	0.026	-0.052	0.0	0.051	0.103
E-ASE $_{\pi}^{rt,50}$	-0.002	0.008	-0.017	-0.002	0.015	0.032
E-ASE $_{\pi}^{rt,66}$	0.001	0.03	-0.056	0.0	0.06	0.116

Table 13: Sensitivity Analysis for different prior specifications (Prior 4).

Estimands	Mean	SD	2.5%	Median	97.5%	Length
E-ATE $_{\pi}$	0.012	0.007	-0.001	0.011	0.027	0.028
E-ASE $_{\pi}$	-0.003	0.006	-0.015	-0.003	0.008	0.024
E-ASE $_{\pi}^{nb,1}$	-0.003	0.006	-0.016	-0.003	0.009	0.025
E-ASE $_{\pi}^{nb,2}$	-0.002	0.018	-0.036	-0.004	0.033	0.069
E-ASE $_{\pi}^{rt,33}$	-0.002	0.027	-0.053	0.0	0.052	0.104
E-ASE $_{\pi}^{rt,50}$	-0.003	0.008	-0.019	-0.004	0.013	0.032
E-ASE $_{\pi}^{rt,66}$	-0.001	0.029	-0.058	0.0	0.059	0.117

References

- Peter M. Aronow and Cyrus Samii. Estimating average causal effects under general interference, with application to a social network experiment. *Annals of Applied Statistics*, 11: 1912–1947, 12 2017. ISSN 19417330. doi: 10.1214/16-AOAS1005.
- Felipe Barrera-Osorio, Marianne Bertrand, Leigh L. Linden, and Francisco Perez-Calle. Improving the design of conditional transfer programs: Evidence from a randomized education experiment in colombia. *American Economic Journal: Applied Economics*, 3(2):167–195, April 2011. doi: 10.1257/app.3.2.167. URL <https://www.aeaweb.org/articles?id=10.1257/app.3.2.167>.
- G. W. Basse, A. Feller, and P. Toulis. Randomization tests of causal effects under interference. *Biometrika*, 106:487–494, 6 2019. ISSN 14643510. doi: 10.1093/biomet/asy072.
- Guillaume Basse and Avi Feller. Analyzing two-stage experiments in the presence of interference. *Journal of the American Statistical Association*, 113:41–55, 1 2018. ISSN 1537274X. doi: 10.1080/01621459.2017.1323641.
- Guillaume W Basse and Edoardo M Airoidi. Model-assisted design of experiments in the presence of network-correlated outcomes. *Biometrika*, 105(4):849–858, 08 2018a. ISSN 0006-3444. doi: 10.1093/biomet/asy036. URL <https://doi.org/10.1093/biomet/asy036>.
- Guillaume W. Basse and Edoardo M. Airoidi. Limitations of design-based causal inference and a/b testing under arbitrary and network interference. *Sociological Methodology*, 48: 136–151, 8 2018b. ISSN 14679531. doi: 10.1177/0081175018782569.
- Yann Bramoullé, Habiba Djebbari, and Bernard Fortin. Peer effects in networks: A survey. *Annual Review of Economics*, 12(1):603–629, 2020. doi: 10.1146/annurev-economics-020320-033926. URL <https://doi.org/10.1146/annurev-economics-020320-033926>.
- Yann Bramoullé, Habiba Djebbari, and Bernard Fortin. Identification of peer effects through social networks. *Journal of Econometrics*, 150(1):41–55, 2009. ISSN 0304-4076. doi: <https://doi.org/10.1016/j.jeconom.2008.12.021>. URL <https://www.sciencedirect.com/science/article/pii/S0304407609000335>.
- Alex Chin. Regression adjustments for estimating the global treatment effect in experiments with interference. *Journal of Causal Inference*, 7(2):20180026, 2019. doi: doi:10.1515/jci-2018-0026. URL <https://doi.org/10.1515/jci-2018-0026>.
- Naoki Egami. Spillover effects in the presence of unobserved networks. *Political Analysis*, 29 (3):287–316, 2021. doi: 10.1017/pan.2020.28.
- Michael D. Escobar. Estimating normal means with a dirichlet process prior. *Journal of the American Statistical Association*, 89(425):268–277, 1994. ISSN 01621459. URL <http://www.jstor.org/stable/2291223>.

- Michael D. Escobar and Mike West. Bayesian density estimation and inference using mixtures. *Journal of the American Statistical Association*, 90(430):577–588, 1995. doi: 10.1080/01621459.1995.10476550. URL <https://www.tandfonline.com/doi/abs/10.1080/01621459.1995.10476550>.
- Thomas S. Ferguson. Prior Distributions on Spaces of Probability Measures. *The Annals of Statistics*, 2(4):615–629, 1974. doi: 10.1214/aos/1176342752. URL <https://doi.org/10.1214/aos/1176342752>.
- Laura Forastiere, Edoardo M. Airoldi, and Fabrizia Mealli. Identification and estimation of treatment and interference effects in observational studies on networks. *Journal of the American Statistical Association*, 116:901–918, 2021. ISSN 1537274X. doi: 10.1080/01621459.2020.1768100.
- Alan Frieze and Michał Karoński. *Introduction to Random Graphs*. Cambridge University Press, Cambridge, 2016. ISBN 978-1-107-11850-8. doi: 10.1017/CBO9781316339831. URL <https://doi.org/10.1017/CBO9781316339831>.
- Paul Gustafson. What are the limits of posterior distributions arising from nonidentified models, and why should we care? *Journal of the American Statistical Association*, 104(488):1682–1695, 2009. doi: 10.1198/jasa.2009.tm08603. URL <https://doi.org/10.1198/jasa.2009.tm08603>.
- Tadao Hoshino and Takahide Yanagi. Randomization test for the specification of interference structure, 2023. URL <https://arxiv.org/abs/2301.05580>.
- Qiang Huang, Jizhe Zhang, Arman Sabbaghi, and Tirthankar Dasgupta. Optimal offline compensation of shape shrinkage for three-dimensional printing processes. *IIE Transactions*, 47(5):431–441, 2015. doi: 10.1080/0740817X.2014.955599. URL <https://doi.org/10.1080/0740817X.2014.955599>.
- Michael G. Hudgens and M. Elizabeth Halloran. Toward causal inference with interference. *Journal of the American Statistical Association*, 103:832–842, 6 2008. ISSN 01621459. doi: 10.1198/016214508000000292.
- Kosuke Imai, Zhichao Jiang, and Anup Malani. Causal inference with interference and noncompliance in two-stage randomized experiments. *Journal of the American Statistical Association*, 116:632–644, 2021. ISSN 1537274X. doi: 10.1080/01621459.2020.1775612.
- Guido W. Imbens and Donald B. Rubin. Bayesian inference for causal effects in randomized experiments with noncompliance. *The Annals of Statistics*, 25(1):305–327, 1997. ISSN 00905364. URL <http://www.jstor.org/stable/2242722>.
- Hemant Ishwaran and Lancelot F James. Gibbs sampling methods for stick-breaking priors. *Journal of the American Statistical Association*, 96(453):161–173, 2001. doi: 10.1198/016214501750332758. URL <https://doi.org/10.1198/016214501750332758>.
- Hemant Ishwaran and Mahmoud Zarepour. Markov chain monte carlo in approximate dirichlet and beta two-parameter process hierarchical models. *Biometrika*, 87(2):371–390, 2000. ISSN 00063444. URL <http://www.jstor.org/stable/2673470>.

- Hyunseung Kang and Guido Imbens. Peer encouragement designs in causal inference with partial interference and identification of local average network effects, 2016. <https://arxiv.org/abs/1609.04464>.
- Michael P. Leung. Causal inference under approximate neighborhood interference. *Econometrica*, 90(1):267–293, 2022. doi: <https://doi.org/10.3982/ECTA17841>. URL <https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA17841>.
- Fan Li, Peng Ding, and Fabrizia Mealli. Bayesian causal inference: A critical review, 2022. URL <https://arxiv.org/abs/2206.15460>. arXiv:2206.15460.
- Shuangning Li and Stefan Wager. Random graph asymptotics for treatment effect estimation under network interference. *The Annals of Statistics*, 50(4):2334–2358, 2022. doi: 10.1214/22-AOS2191. URL <https://doi.org/10.1214/22-AOS2191>.
- Xinran Li, Peng Ding, Qian Lin, Dawei Yang, and Jun S. Liu. Randomization inference for peer effects. *Journal of the American Statistical Association*, 114(528):1651–1664, 2019. doi: 10.1080/01621459.2018.1512863. URL <https://doi.org/10.1080/01621459.2018.1512863>.
- Lan Liu and Michael G. Hudgens. Large sample randomization inference of causal effects in the presence of interference. *Journal of the American Statistical Association*, 109:288–301, 2014. ISSN 1537274X. doi: 10.1080/01621459.2013.844698.
- Steven Maceachern. Dependent nonparametric processes. *Proceedings of the Section on Bayesian Statistical Science, American Statistical Association*, pages 50–55, 01 1999.
- Steven Maceachern. Dependent dirichlet processes. *Technical Report*, 2000. [http://www.gatsby.ucl.ac.uk/~sim\\$porbanz/talks/MacEachern2000.pdf](http://www.gatsby.ucl.ac.uk/~sim$porbanz/talks/MacEachern2000.pdf).
- Mark Newman. *Networks: An Introduction*. Oxford University Press, Inc., USA, 2010. ISBN 0199206651.
- Yuki Ohnishi and Arman Sabbaghi. A Bayesian Analysis of Two-Stage Randomized Experiments in the Presence of Interference, Treatment Nonadherence, and Missing Outcomes. *Bayesian Analysis*, pages 1–30, 2022. URL <https://doi.org/10.1214/22-BA1347>. , doi.org/10.1214/22-BA1347.
- Georgia Papadogeorgou and Srijata Samanta. Spatial causal inference in the presence of unmeasured confounding and interference, 2023. arXiv:2303.08218.
- Fernando A. Quintana, Peter Müller, Alejandro Jara, and Steven N. MacEachern. The Dependent Dirichlet Process and Related Models. *Statistical Science*, 37(1):24–41, 2022. doi: 10.1214/20-STS819. URL <https://doi.org/10.1214/20-STS819>.
- Donald B Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of educational Psychology*, 66(5):688, 1974.
- Donald B. Rubin. Bayesian Inference for Causal Effects: The Role of Randomization. *The Annals of Statistics*, 6(1):34–58, 1978. doi: 10.1214/aos/1176344064. URL <https://doi.org/10.1214/aos/1176344064>.

- Arman Sabbaghi, Tirthankar Dasgupta, Qiang Huang, and Jizhe Zhang. Inference for deformation and interference in 3D printing. *The Annals of Applied Statistics*, 8(3):1395–1415, 2014. doi: 10.1214/14-AOAS762. URL <https://doi.org/10.1214/14-AOAS762>.
- Carl-Erik Särndal, Bengt Swensson, and Jan Wretman. *Model Assisted Survey Sampling*. Springer New York, NY, 1992.
- Daniel L. Sussman and Edoardo M. Airoldi. Elements of estimation theory for causal effects in the presence of network interference, 2017.
- Fredrik Sävje. Causal inference with misspecified exposure mappings: separating definitions and assumptions. *Biometrika.*, 03 2023. ISSN 1464-3510. URL <https://doi.org/10.1093/biomet/asad019>. doi.org/10.1093/biomet/asad019.
- Fredrik Sävje, Peter M. Aronow, and Michael G. Hudgens. Average treatment effects in the presence of unknown interference. *Annals of Statistics*, 49:673–701, 2021. ISSN 21688966. doi: 10.1214/20-AOS1973.
- Panos Toulis and Edward Kao. Estimation of causal peer influence effects. In *Proceedings of the 30th International Conference on Machine Learning*, volume 28, Atlanta, Georgia, USA, 17–19 Jun 2013. PMLR. URL <https://proceedings.mlr.press/v28/toulis13.html>. pages: 1489–1497.
- Corwin M. Zigler and Georgia Papadogeorgou. Bipartite Causal Inference with Interference. *Statistical Science*, 36(1):109–123, 2021. doi: 10.1214/19-STS749. URL <https://doi.org/10.1214/19-STS749>.