

A Bayesian Analysis of Two-Stage Randomized Experiments in the Presence of Interference, Treatment Nonadherence, and Missing Outcomes

Yuki Ohnishi* and Arman Sabbaghi†

Abstract. Three critical issues for causal inference that often occur in modern, complicated experiments are interference, treatment nonadherence, and missing outcomes. A great deal of research efforts has been dedicated to developing causal inferential methodologies that address these issues separately. However, methodologies that can address these issues simultaneously are lacking. We propose a Bayesian causal inference methodology to address this gap. Our methodology extends existing causal frameworks and methods, specifically, two-staged randomized experiments and the principal stratification framework. In contrast to existing methods that invoke strong structural assumptions to identify principal causal effects, our Bayesian approach uses flexible distributional models that can accommodate the complexities of interference and missing outcomes, and that ensure that principal causal effects are weakly identifiable. We illustrate our methodology via simulation studies and a re-analysis of real-life data from an evaluation of India’s National Health Insurance Program. Our methodology enables us to identify new active causal effects that were not identified in past analyses. Ultimately, our simulation studies and case study demonstrate how our methodology can yield more informative analyses in modern experiments with interference, treatment nonadherence, missing outcomes, and complicated outcome generation mechanisms.

Keywords: Bayesian causal inference, noncompliance, principal stratification, Rubin causal model, two-stage randomized design, missing not at random.

1 Introduction

Causal inference is a fundamental consideration across a wide range of domains in science, technology, engineering, and medicine (Pearl, 2009; Imbens and Rubin, 2015). A traditional gold standard for performing causal inference is the classical randomized experiment (Rubin, 2008; Imbens and Rubin, 2015). In this type of experiment, a great deal of control and precautions can be taken so as to eliminate events that would introduce instabilities and biases in causal inferences. However, modern experiments, e.g., social experiments, can become so complicated that it may be difficult to institute such control and precautions. Three significant sources of complications that are increasingly

arXiv: [2110.10216](https://arxiv.org/abs/2110.10216)

*Purdue University, West Lafayette, IN, USA, yohnishi@purdue.edu

†Purdue University, West Lafayette, IN, USA, sabbaghi@purdue.edu

of interest are interference among experimental units, nonadherence/noncompliance of experimental units to their assigned treatments, and unintended missing outcomes of experimental units. Interference exists if the outcome of an experimental unit depends not only on its assigned treatment, but also on the assigned treatments for other units. It arises when limited controls are placed on the interactions of experimental units with one another, or when competition for a limited set of resources exists, during the course of an experiment. Treatment nonadherence frequently occurs in human subject experiments, as it can be unethical to force an individual to take their assigned treatment. Clinical trials in particular typically have subjects that do not adhere to their assigned treatments due to adverse side effects or intercurrent events (Little et al., 2012). Missing outcomes commonly occur in human studies. For example, respondents may refuse to report sensitive outcomes (e.g., income) after receiving a treatment in a study (Little and Rubin, 2002, p. 3). Failing to account for interference, nonadherence, and missing outcomes in modern experiments will generally yield unstable and biased inferences on treatment effects.

A great deal of research efforts has been dedicated over the past three decades to developing causal inferential methodologies that address the first two issues separately. Hudgens and Halloran (2008) first introduced the concept of the two-stage randomized design for performing causal inference in the presence of interference when all experimental units adhere to their assigned treatments. In this design, experimental units belong to clusters, and randomizations are performed at both the cluster-level and experimental unit-level. Specifically, the clusters are first randomly assigned different probabilities for treatment assignment of their constituent experimental units, and then treatments are randomly assigned to the units within the clusters (with treatment assignment performed independently across clusters) based on the cluster-specific assignment probabilities. For example, in our case study, villages correspond to the clusters, and each household within a village corresponds to an experimental unit. In this manner, we can analyze the effectiveness of an insurance plan implemented in India under the two-stage randomized design. Hudgens and Halloran (2008) demonstrated how both direct treatment effects and indirect treatment effects (i.e., those effects that can be attributed to the treatments received by other units) can be inferred under this design in the presence of interference. This design and the corresponding causal inference methods that can be performed under it have been further studied by VanderWeele and Tchetgen (2011), Tchetgen and Vanderweele (2012), Liu and Hudgens (2014), and Basse and Feller (2018). For experiments that have treatment nonadherence but not interference, one standard methodology that has been considered for their analyses is the intention-to-treat (ITT) method. Under this approach, the treatment received by an experimental unit is ignored, and instead only the treatment assigned is considered. This method follows the traditional principle of analyzing an experiment according to its physical randomization (Cox and Reid, 2000, p. 14) mechanism, and can yield valid causal inferences on the effects of treatment *assigned* in certain situations. However, the ITT method will generally yield biased inferences on the effects of treatment *received* because it does not account for the latent stratification of experimental units defined according to their adherence behaviors to the different treatment assignments, and consequently does not provide inferences for the target stratum of compliers. Angrist

et al. (1996) and Imbens and Rubin (1997) developed a framework to provide a more principled approach for the analyses of experiments with nonadherence, and Frangakis and Rubin (2002) extended those frameworks to develop the more general principal stratification framework. This framework has been applied to a wide variety of real-life problems involving complications such as censoring or truncation due to death (Zhang and Rubin, 2003) and the occurrence of intermediate variables that are thought to mediate the effects of treatments on the outcome (Gallop et al., 2009). VanderWeele (2011) provides a detailed review of principal stratification.

Recently, methodologies have been developed to address combinations of interference, nonadherence, and missing outcomes, but not all of these issues simultaneously. Under the assumption of fully observed outcomes, Imai et al. (2021) presented a non-parametric identification of the complier average direct and indirect effects, and proposed consistent estimators for them under the two-stage randomized design in the presence of interference and nonadherence. They derived large-sample nonparametric bounds for the causal effects, but their results are not necessarily valid for the finite-sample regime. In addition, the estimators of Imai et al. (2021) are sensitive to outliers. Vazquez-Bare (2022) built on the work of Imai et al. (2021) and analyzed spillover effects using instrumental variables in the two-stage randomized experiment with fully observed outcomes. They considered the identification of causal direct and spillover effects under one-sided noncompliance and demonstrated that these effects can be estimated using two-stage least squares (2SLS). However, their methods do not work in general under two-sided noncompliance or when units have multiple peers without a strong structural assumption about peers' compliance types. Kang and Imbens (2016) considered the peer encouragement design to study network treatment effects when treatment randomization cannot feasibly be forced on experimental units, and presented identification results only for the case of one-sided noncompliance. Forastiere et al. (2016) developed a Bayesian principal stratification method for causal inference in clustered encouragement designs (CEDs), where the assignment of treatment encouragement is performed at the cluster level. The CED is effectively a special case of the two-stage randomized experiment with no randomization within clusters, i.e., with treatment assignment probability for units within clusters being either zero or one. As there is no randomization performed at the unit level within clusters, their methodology cannot capture the complexity of adherence behaviors in the two-stage experiments, where units could possibly change their adherence behaviors depending on what assignment probabilities their clusters are assigned to. It is important to note that none of the above methods can easily accommodate missing outcomes in the presence of interference and nonadherence. Experiments with treatment nonadherence and missing outcomes have been analyzed by Frangakis and Rubin (1999), Mattei and Mealli (2007), Frumento et al. (2012), and Mattei et al. (2014), but none of these studies involved interference.

We develop a new Bayesian causal inferential methodology for two-stage randomized experiments with interference, noncompliance, and missing outcomes. Our methodology utilizes the principal stratification framework to address the identification issues arising from treatment nonadherence and missing outcomes in the presence of interference. To the best of our knowledge, none of the existing causal inference methods have applied Bayesian principal stratification to two-stage randomized designs and addressed these

issues simultaneously as our method does. Not only does our Bayesian approach provide a principled framework to analyze two-stage randomized experiments in the presence of all of these complications, but it clarifies what can be learned when causal estimands are not identifiable but are instead weakly identifiable (i.e., when the likelihood functions of parameters and causal estimands have substantial regions of flatness). It is important to recognize that issues of identifiability under the Bayesian paradigm are distinct from those under the frequentist paradigm because the specification of proper prior distributions always yields proper posterior distributions (Imbens and Rubin, 1997). In particular, our Bayesian method enables us to infer principal causal effects under two-sided nonadherence without the need for strong structural assumptions, and define new types of interpretable and informative causal estimands (such as the complier spillover and overall treatment effects) that would be of great interest to policy makers. Furthermore, our methodology highlights new assumptions on compliance types and missingness mechanisms for making causal inferences more efficient and stable in such complicated experiments. The use of Bayesian models in our methodology enables us to naturally accommodate complicated outcome generation mechanisms that frequently occur in modern experiments, such as heavy-tailed, skewed, zero-inflated, and/or multi-modal outcome distributions.

We proceed in Section 2 to review the Rubin Causal Model (Rubin, 1974; Holland, 1986), the principal stratification framework, and relevant assumptions and causal estimands for the two-stage randomized experiment. Section 3 introduces the models and computational algorithms involved in our Bayesian methodology. In Section 4 we perform extensive simulation studies to investigate the frequentist performance of our Bayesian method for a heavy-tailed distribution with an excess of zero outcomes. These simulation studies effectively validate our methodology in a situation where the existing frequentist approach (Imai et al., 2021) performs poorly in terms of bias and mean squared error (MSE). Finally, in Section 5 we apply our methodology to the real-life data from the evaluation of India’s National Health Insurance Program (RSBY) (Nandi et al., 2015). In our analysis of this case study we are able to uncover more definitive evidence of causal effects that were previously found to be insignificant in past analyses. Our concluding remarks are in Section 6.

2 Background

2.1 Two-Stage Randomized Experiments with Interference and Noncompliance

Throughout this manuscript we consider two-stage randomized experiments involving two treatments and J clusters, with N_j experimental units in cluster $j = 1, \dots, J$ (each experimental unit belongs to only one cluster). We let $N = \sum_{j=1}^J N_j$ denote the total number of experimental units. The assignment mechanism in the two-stage randomized experiment is performed sequentially, with each stage involving a type of completely randomized design. In the first stage, J_1 clusters are randomly chosen to have a treatment assignment probability of $a_1 \in (0, 1)$ for their constituent experimental units, and

the remaining $J - J_1$ clusters have a treatment assignment probability of $a_0 \in (0, 1)$ for their units. For each $j = 1, \dots, J$ we let A_j be the indicator for whether cluster j was assigned a_1 ($A_j = 1$) or a_0 ($A_j = 0$). We let $\mathbf{A} = (A_1, \dots, A_J)^\top$, and without loss of generality we let $a_1 > a_0$. In the second stage, the experimental units within the clusters are randomly assigned treatment and control according to the treatment assignment probabilities assigned to their clusters, with the treatment assignment performed independently across clusters. For each cluster j with $A_j = 1$, $N_j a_1$ of their experimental units are randomly assigned treatment and the remaining $N_j(1 - a_1)$ are assigned control. Similarly, for each cluster j' with $A_{j'} = 0$, $N_{j'} a_0$ of their experimental units are randomly assigned treatment and the remaining $N_{j'}(1 - a_0)$ are assigned control. We assume that $N_j a_1$ and $N_j a_0$ are integers for all $j = 1, \dots, J$. We let $Z_{i,j}$ denote the treatment assignment indicator for unit i in cluster j , with $Z_{i,j} = 1$ if it is assigned treatment and $Z_{i,j} = 0$ otherwise. We let $\mathbf{Z}_j = (Z_{1,j}, \dots, Z_{N_j,j})^\top$ denote the vector of treatment assignment indicators for all units in cluster j , and $\mathbf{Z}_{-i,j}$ denote the subvector of \mathbf{Z}_j with the i th entry removed. Other assignment mechanisms for two-stage randomized experiments are provided by VanderWeele and Tchetgen (2011), but we do not consider them here.

As we consider two-stage randomized experiments with interference, nonadherence and missingness under the Rubin Causal Model, we must introduce two types of potential outcomes for the treatment received by an experimental unit and the final outcome of interest that are functions of the experimental units' treatment assignments. We let $D_{i,j}(\mathbf{z})$ denote the treatment received for unit i in cluster j under treatment assignment $\mathbf{z} \in \{0, 1\}^N$, $\mathbf{D}_j(\mathbf{z}) = (D_{1,j}(\mathbf{z}), \dots, D_{N_j,j}(\mathbf{z}))^\top$ be the vector of treatments received by the units in cluster j , and $\mathbf{D}(\mathbf{z}) = (\mathbf{D}_1(\mathbf{z}), \dots, \mathbf{D}_J(\mathbf{z}))^\top$. Furthermore, we let $Y_{i,j}(\mathbf{z}, \mathbf{D}(\mathbf{z}))$ denote the potential outcome for unit i in cluster j under treatment assignment vector \mathbf{z} and treatment received vector $\mathbf{D}(\mathbf{z})$. Although the potential outcomes can be written solely as a function of \mathbf{z} (because $\mathbf{D}(\mathbf{z})$ is a function of \mathbf{z}) we include $\mathbf{D}(\mathbf{z})$ in the notation for $Y_{i,j}(\mathbf{z}, \mathbf{D}(\mathbf{z}))$ to emphasize the existence of nonadherence. We let \mathbf{D} and \mathbf{Y} denote the matrices containing all the potential values of treatment receipts and outcomes, respectively, for all experimental units across all treatment assignments.

Finally, we let $M_{i,j}(\mathbf{z})$ denote the missingness indicator for the realized outcome of unit i in cluster j under treatment assignment $\mathbf{z} \in \{0, 1\}^N$, $\mathbf{M}_j(\mathbf{z}) = (M_{1,j}(\mathbf{z}), \dots, M_{N_j,j}(\mathbf{z}))^\top$ be the vector of missingness indicators of the units in cluster j , and $\mathbf{M}(\mathbf{z}) = (\mathbf{M}_1(\mathbf{z}), \dots, \mathbf{M}_J(\mathbf{z}))^\top$. We note that $Y_{i,j}(\mathbf{z}, \mathbf{D}(\mathbf{z}))$ is observed when $M_{i,j}(\mathbf{z}) = 1$, and is missing otherwise. It is important to distinguish between missing potential outcomes and unrealized potential outcomes. For a given treatment assignment vector \mathbf{z} that was realized in an experiment, the potential outcomes $Y_{i,j}(\mathbf{z}', \mathbf{D}(\mathbf{z}'))$ for any other distinct treatment assignment vector \mathbf{z}' are referred to as unrealized potential outcomes. Also, for a given treatment assignment vector \mathbf{z} that was realized in an experiment, if $M_{i,j}(\mathbf{z}) = 1$ then $Y_{i,j}(\mathbf{z}, \mathbf{D}(\mathbf{z}))$ is not observed, although the potential outcome was realized in the experiment.

Besides the treatments, potential outcomes, and missingness indicators, we assume covariates are measured for the experimental units. These covariates are either measured prior to treatment assignment, or are measured afterwards but are not affected by treatment assignment. We denote the vector of covariates for unit i in cluster j by $\mathbf{X}_{i,j}$.

2.2 Assumptions on the Structure of Interference

We extend the assumptions proposed by (Imai et al., 2021) for the two-stage randomized experiment with nonadherence. We first invoke the partial interference assumption in which units in different clusters do not interact or affect one another. This assumption was first formulated by Hudgens and Halloran (2008), and was extended to the noncompliance setting by Imai et al. (2021). Partial interference facilitates causal inference in our setting of interest because an experimental unit i 's treatment received, missingness indicator will only be functions of treatment assignments for other units within the same cluster j as unit i .

Assumption 1. For all $\mathbf{z}, \mathbf{z}' \in \{0, 1\}^N$ such that $\mathbf{z}_j = \mathbf{z}'_j$ for a cluster j , then $D_{i,j}(\mathbf{z}) = D_{i,j}(\mathbf{z}')$ and $M_{i,j}(\mathbf{z}) = M_{i,j}(\mathbf{z}')$ for all experimental units i in cluster j .

The next assumption that we consider is the stratified interference assumption of Hudgens and Halloran (2008). This assumption imposes further structure on interference by having the treatment received, the missingness indicator, and the potential outcome for an experimental unit being a function of just the number of experimental units assigned treatment within the same cluster. This assumption was also considered by Forastiere et al. (2016) and Imai et al. (2021). It is important to recognize that a great deal of work has been conducted to move beyond this condition and consider more flexible structures of interference (Aronow, 2012; Manski, 2013; Basse and Airolidi, 2018b; Aronow and Samii, 2017; Baird et al., 2018; Basse and Airolidi, 2018a; Athey et al., 2018; Basse et al., 2019; Leung, 2020; Forastiere et al., 2021; Sävje et al., 2021). However, none of these works considered noncompliance and missing outcomes.

Assumption 2. For a cluster j and experimental unit i in cluster j , if $\mathbf{z}, \mathbf{z}' \in \{0, 1\}^N$ such that $z_{i,j} = z'_{i,j}$ and $\mathbf{z}_j^\top \mathbf{1} = (\mathbf{z}'_j)^\top \mathbf{1}$, then $D_{i,j}(\mathbf{z}) = D_{i,j}(\mathbf{z}')$ and $M_{i,j}(\mathbf{z}) = M_{i,j}(\mathbf{z}')$.

We invoke the same assumptions on outcomes.

Assumption 3. For all $\mathbf{z}, \mathbf{z}' \in \{0, 1\}^N$ such that $\mathbf{z}_j = \mathbf{z}'_j$ for a cluster j , then $Y_{i,j}(\mathbf{z}, \mathbf{D}(\mathbf{z})) = Y_{i,j}(\mathbf{z}', \mathbf{D}(\mathbf{z}'))$ for all experimental units i in cluster j .

Assumption 4. For a cluster j and experimental unit i in cluster j , if $\mathbf{z}, \mathbf{z}' \in \{0, 1\}^N$ such that $z_{i,j} = z'_{i,j}$ and $\mathbf{z}_j^\top \mathbf{1} = (\mathbf{z}'_j)^\top \mathbf{1}$, $Y_{i,j}(\mathbf{z}, \mathbf{D}(\mathbf{z})) = Y_{i,j}(\mathbf{z}', \mathbf{D}(\mathbf{z}'))$.

Finally, we assume the exclusion restriction with interference between units in the two-stage randomized experiment. In this case, the potential outcome of a unit i in cluster j only depends on the treatments received by units within cluster j .

Assumption 5. For any treatment assignment vectors $\mathbf{z}, \mathbf{z}' \in \{0, 1\}^N$ with $\mathbf{D}_j(\mathbf{z}) = \mathbf{D}_j(\mathbf{z}')$, $Y_{i,j}(\mathbf{z}, \mathbf{D}(\mathbf{z})) = Y_{i,j}(\mathbf{z}', \mathbf{D}(\mathbf{z}'))$ for all experimental units i in cluster j .

These five assumptions imply that for two-stage randomized experiments in which the number of treated units within any cluster is fixed by design, the potential outcomes and missingness indicator for an experimental unit are a function of its own treatment assignment and the treatment assignment probability for its cluster. We accordingly slightly abuse the notation to write $D_{i,j}(\mathbf{z})$, $M_{i,j}(\mathbf{z})$ and $Y_{i,j}(\mathbf{z}, \mathbf{D}(\mathbf{z}))$ as $D_{i,j}(z, a)$,

$M_{i,j}(z, a)$ and $Y_{i,j}(z, a)$, respectively, where z denotes the treatment assignment for the experimental unit and a denotes the treatment assignment probability for the unit's cluster.

2.3 Principal Strata, Monotonicity, and the Exclusion Restriction for Two-Stage Randomized Experiments

Under the principal stratification framework, we stratify the experimental units according to their values of $D_{i,j}(z, a)$ under the different possible treatment assignments $z \in \{0, 1\}$ and treatment assignment probabilities $a \in \{a_0, a_1\}$ for the clusters. There exist four such potential values: $D_{i,j}(0, a_0)$, $D_{i,j}(1, a_0)$, $D_{i,j}(0, a_1)$, and $D_{i,j}(1, a_1)$. A unique feature of our consideration of nonadherence for the two-stage randomized design is that, according to Assumption 2 and 4, units can have different compliance behaviors under different assignment probabilities for their clusters. We formally define the compliance behavior of unit i in cluster j under each treatment assignment probability $a \in \{a_0, a_1\}$ for the cluster as

$$G_{i,j}(a) = \begin{cases} n & \text{if } D_{i,j}(0, a) = D_{i,j}(1, a) = 0, \\ c & \text{if } D_{i,j}(0, a) = 0, D_{i,j}(1, a) = 1, \\ d & \text{if } D_{i,j}(0, a) = 1, D_{i,j}(1, a) = 0, \\ a & \text{if } D_{i,j}(0, a) = D_{i,j}(1, a) = 1, \end{cases}$$

where n, c, d , and a denote never-takers, compliers, defiers, and always-takers, respectively. Finally, the compliance behavior for unit i in cluster j is defined according to the pair of compliance indicators $G_{i,j} = (G_{i,j}(a_0), G_{i,j}(a_1))$.

A standard assumption for the compliance behavior is monotonicity.

Assumption 6. *For all units $i = 1, \dots, N_j$ in cluster $j = 1, \dots, J$ and any $a \in \{a_0, a_1\}$, $D_{i,j}(1, a) \geq D_{i,j}(0, a)$, and strict inequality exists for at least one experimental unit.*

This assumption was also considered by Imai et al. (2021) and Forastiere et al. (2016). Monotonicity eliminates the possibility of defiers under either treatment assignment probabilities a_0 and a_1 . It also reduces the number of principal strata from sixteen to nine.

In addition to this existing monotonicity assumption, we also consider the following new assumption on compliance behaviors across the different treatment assignment probabilities that can be assigned to the clusters.

Assumption 7. *The set $\{n, c, a\}$ is a partial order set. For $a_0 < a_1$, units whose clusters have been assigned a_1 would have a non-strictly lower compliance type of $\{n, c, a\}$ than the current compliance type if their clusters were assigned a_0 . Units whose clusters have been assigned a_0 would have a non-strictly greater compliance type of $\{n, c, a\}$ than the current compliance type if their clusters were assigned a_1 .*

This assumption further reduces the number of principal strata. To see this, we first recognize that there are only six possible definitions of partial orders for $\{n, c, a\}$:

$n \leq c \leq a$, $n \leq a \leq c$, $c \leq n \leq a$, $c \leq a \leq n$, $a \leq n \leq c$, and $a \leq c \leq n$. We adopt the partial orders $n \leq c \leq a$ as well as $a_0 \leq a_1$ throughout. These correspond to units being more likely to take treatment if a larger proportion of their neighbors take treatment. Alternatively, if a cluster is assigned a_1 then its units are more likely to receive treatment, no matter what the units are assigned, compared to the case if the cluster is assigned a_0 . The combination of Assumptions 6 and 7 thus reduces the number of principal strata to six: $\{(n, n), (c, c), (a, a), (n, c), (n, a), (c, a)\}$.

Our definition of principal strata based on Assumptions 6 and 7 is more general than existing definitions. For example, Imai et al. (2021) considered the monotonicity assumption with respect to the treatment assignment probability. Their assumption corresponds to the partial order $n \leq c \leq a$ but can not express other orderings of monotonicity, e.g., $a \leq c \leq n$. Forastiere et al. (2016) only defined three principal strata based on the treatment uptake status, excluding defiers. Also, Vazquez-Bare (2022) defined five principal strata by posing monotonicity directly on treatment received, which led to the removal of strata (n, a) from consideration. These distinctions exist because we consider the two-stage randomized design, and we define the monotonicity assumptions on the compliance behaviors with respect to the treatment probabilities for clusters. Forastiere et al. (2016) considered clustered encouragement designs (CEDs) in which encouragement is randomized at the level of clusters but with no randomization carried out within clusters. In contrast, the two-stage randomized design has encouragement randomized at the level of units within clusters, with clusters assigned a treatment probability that governs the proportion of treated units within the cluster. The latter design generates more complicated structures for the units' compliance behaviors because some units might behave differently based on the treatment probabilities assigned to their clusters. This distinction is critical because of our need to capture behavioral shifts of units between such treatment probabilities, and because of the direct and spillover causal estimands of interest that are defined in Section 2.4. Our definitions of the monotonicity assumptions provide more flexible orderings of compliance behaviors than that of Vazquez-Bare (2022) because we permit more relevant orderings.

In addition to monotonicity, we also consider the exclusion restriction for certain principal strata in the case of nonadherence.

Assumption 8. For any unit $i = 1, \dots, N_j$ in cluster $j = 1, \dots, J$:

- if $G_{i,j} \in \{(a, a), (n, n), (n, a)\}$, then $Y_{i,j}(0, a) = Y_{i,j}(1, a)$ for $a \in \{a_0, a_1\}$,
- if $G_{i,j} = (c, a)$, $Y_{i,j}(0, a_1) = Y_{i,j}(1, a_1)$, and
- if $G_{i,j} = (n, c)$, $Y_{i,j}(0, a_0) = Y_{i,j}(1, a_0)$.

This assumption captures the idea that treatment assignment has no effect on the outcome if the unit is either an always-taker or a never-taker under each treatment assignment probability. It also corresponds to Assumption 5, from which the outcome of an experimental unit is determined only through the treatment received by units within the same cluster. This assumption is equivalent to a modified form of the restricted interference assumption of Imai et al. (2021) that is stated in terms of principal strata.

2.4 Direct, Spillover, and Overall Causal Estimands

We consider the finite-population framework for causal estimands under the Rubin Causal Model. In this case, estimands are defined in terms of comparisons of potential outcomes for the N experimental units. Our approach for defining direct, spillover, and overall causal estimands in two-stage randomized experiments with interference and nonadherence follows that of Hudgens and Halloran (2008) and Imai et al. (2021).

To simplify our expressions of the causal estimands, we first write the treatment received and potential outcome for unit i in cluster j under treatment assignment \mathbf{z} as $D_{i,j}(\mathbf{z})$ and $Y_{i,j}(\mathbf{z}, \mathbf{D}(\mathbf{z}))$, respectively. We let $K_j(a)$ denote the number of treated units in cluster j under treatment assignment probability $a \in \{a_0, a_1\}$, and $\mathcal{Z}_{-i,j}$ denote the set of all subvectors of $\mathbf{z}_j \in \{0, 1\}^{N_j}$ with the i th element removed such that $\sum_{i=1}^{N_j} z_{i,j} = K_j(a)$. Under Assumptions 1 - 4, we recognize that $\bar{D}_{i,j}(z, a) = \sum_{\mathbf{z}_{-i,j} \in \mathcal{Z}_{-i,j}} D_{i,j}(\mathbf{z}) p(\mathbf{Z}_{-i,j} = \mathbf{z}_{-i,j} \mid Z_{i,j} = z, A_j = a) = D_{i,j}(z, a)$ and $\bar{Y}_{i,j}(z, a) = \sum_{\mathbf{z}_{-i,j} \in \mathcal{Z}_{-i,j}} Y_{i,j}(\mathbf{z}, \mathbf{D}(\mathbf{z})) p(\mathbf{Z}_{-i,j} = \mathbf{z}_{-i,j} \mid Z_{i,j} = z, A_j = a) = Y_{i,j}(z, a)$.

One set of causal estimands that we consider are the unit-level direct intention-to-treat (ITT) effects of treatment assignment on treatment received and the final outcome under treatment assignment probability $a \in \{a_0, a_1\}$, i.e., $\text{ITT}_{D,i,j}(a) = \bar{D}_{i,j}(1, a) - \bar{D}_{i,j}(0, a)$ and $\text{ITT}_{Y,i,j}(a) = \bar{Y}_{i,j}(1, a) - \bar{Y}_{i,j}(0, a)$, respectively. These estimands capture the adherence behavior and changes in the final outcome under the same treatment assignment probability a when unit i in cluster j is assigned treatment as opposed to control. We average the unit-level effects to define the cluster-level and finite population-level ITT effects as $\text{ITT}_{D,\cdot,j}(a) = \sum_{i=1}^{N_j} \text{ITT}_{D,i,j}(a)/N_j$, $\text{ITT}_{D,\cdot,\cdot}(a) = \sum_{j=1}^J N_j \text{ITT}_{D,\cdot,j}(a)/N$, $\text{ITT}_{Y,\cdot,j}(a) = \sum_{i=1}^{N_j} \text{ITT}_{Y,i,j}(a)/N_j$, $\text{ITT}_{Y,\cdot,\cdot}(a) = \sum_{j=1}^J N_j \text{ITT}_{Y,\cdot,j}(a)/N$.

Another causal estimand of interest in the presence of interference is the spillover (indirect) effect of treatments assigned to other units on the potential outcomes for a particular experimental unit. Following Imai et al. (2021) we define unit-level spillover effects on treatment receipt and outcome as $S_{D,i,j}(z) = \bar{D}_{i,j}(z, a_1) - \bar{D}_{i,j}(z, a_0)$ and $S_{Y,i,j}(z) = \bar{Y}_{i,j}(z, a_1) - \bar{Y}_{i,j}(z, a_0)$. In these two estimands, we consider the average potential outcomes corresponding to the two different treatment assignment probabilities a_0 and a_1 for cluster j but the same treatment assignment z for unit i in cluster j . These estimands quantify the intuition that, if differences exist in potential outcomes under the same treatment assignment z , then they can be attributed to the first stage of the experiment that governs the proportion of treated units in each cluster. This is because under Assumption 2 and 4 the treated units in a particular unit i 's cluster are the only factor besides the assigned treatment that can affect unit i 's outcomes, and so the $S_{D,i,j}(z)$ and $S_{Y,i,j}(z)$ can be reasonably regarded as spillover effects. We define cluster-level and population-level spillover effects by averaging the $S_{D,i,j}(z)$ and $S_{Y,i,j}(z)$ across i and j , respectively.

The final estimand that we consider is the overall effect of the first stage of the experiment, i.e., the effect of having treatment assignment probability a_1 versus a_0 for a cluster. This effect is usually of greatest interest for policy makers. For example, infectious disease experts may be interested in comparing infection rates under

two different vaccine allocation plans (e.g., 40% and 80%) within each cluster. We define unit-level overall effects of the first stage on treatment receipt and outcome as $O_{D,i,j} = \bar{D}_{i,j}(a_1) - \bar{D}_{i,j}(a_0)$ and $O_{Y,i,j} = \bar{Y}_{i,j}(a_1) - \bar{Y}_{i,j}(a_0)$, respectively, where $\bar{Y}_{i,j}(a) = \sum_{\mathbf{z}_j \in \mathcal{Z}_j} Y_{i,j}(\mathbf{z})p(\mathbf{Z}_j = \mathbf{z}_j \mid A_j = a)$ and \mathcal{Z}_j is a set of all possible assignment vectors \mathbf{z}_j . We have that $\bar{Y}_{i,j}(a)$ is effectively the average value of the individual's outcome under the treatment assignment probability $a \in \{a_0, a_1\}$. It is decomposed into

$$\bar{Y}_{i,j}(a) = \left(\frac{K_j(a)}{N_j} \right) \bar{Y}_{i,j}(1, a) + \left(\frac{N_j - K_j(a)}{N_j} \right) \bar{Y}_{i,j}(0, a). \quad (2.1)$$

The proof of equation (2.1) is provided in the supplementary material (Ohnishi and Sabbaghi, 2022). Plugging this into $O_{Y,i,j}$, we have $O_{Y,i,j} = \{K_j(a_1)/N_j\} \text{ITT}_{Y,i,j}(a_1) - \{K_j(a_0)/N_j\} \text{ITT}_{Y,i,j}(a_0) + S_{Y,i,j}(0)$. We average the $O_{Y,i,j}$ to define the cluster-level and population-level overall effects as $O_{Y,\cdot,j} = \sum_{i=1}^{N_j} O_{Y,i,j}/N_j$ and $O_{Y,\cdot,\cdot} = \sum_{j=1}^J N_j O_{Y,\cdot,j}/N$, respectively. VanderWeele and Tchetgen (2011) provided the same decomposition but with a slightly different proof. The overall effect is expressed as the sum of the spillover effect and the contrast of two ITT effects under the treatment assignment probability a_1 and under the treatment assignment probability a_0 , each of which are multiplied by the proportion of treated units under that assignment probability.

2.5 Principal Causal Estimands

In addition to defining direct, spillover, and overall causal estimands, we define new principal causal estimands. These estimands extend those in Section 2.4 to account for compliers under the different treatment probabilities a_0 and a_1 that can be assigned to clusters.

Imai et al. (2021) used Assumption 2 and 4 to define the complier average direct effect under treatment probability $a \in \{a_0, a_1\}$ as

$$\left[\sum_{j=1}^J \sum_{i=1}^{N_j} \{Y_{i,j}(1, a) - Y_{i,j}(0, a)\} \mathbb{1}(G_{i,j}(a) = c) \right] / \sum_{j=1}^J \sum_{i=1}^{N_j} \mathbb{1}(G_{i,j}(a) = c). \quad (2.2)$$

where $\mathbb{1}(\cdot)$ denotes the indicator for an event. We define a new complier average direct effect by generalizing the estimand in equation (2.2) to account for how experimental units comply under the different treatment probabilities assigned to clusters. This new definition is motivated by the practical interest in modifying the effect of a_1 to account for the units who would comply under a_0 , and vice versa. To define this new complier average direct effect, we introduce the notions of the “base cluster assignment” as the treatment probability for a cluster under which compliers are defined, and the “target cluster assignment” as the treatment probability for a cluster under which causal effects are considered. It is important to separate the base from the target cluster assignments because the treatments received by units are influenced by the treatments assigned to other units, and because compliance behaviors vary with the treatment probabilities a_0 and a_1 . For a target cluster assignment a and base cluster assignment a' , our new

population-level complier average direct effect is defined as

$$\text{CADE}(a, a') = \left[\sum_{j=1}^J \sum_{i=1}^{N_j} \{Y_{i,j}(1, a) - Y_{i,j}(0, a)\} \mathbb{1}(G_{i,j}(a') = c) \right] / \left[\sum_{j=1}^J \sum_{i=1}^{N_j} \mathbb{1}(G_{i,j}(a') = c) \right]. \quad (2.3)$$

Equation (2.2) is a special case of equation (2.3) with $a = a'$.

The complier average spillover effect for treatment z defined by Imai et al. (2021) is the ratio of $\left[\sum_{j=1}^J \sum_{i=1}^{N_j} \{Y_{i,j}(z, a_1) - Y_{i,j}(z, a_0)\} \mathbb{1}(D_{i,j}(z, a_1) = 1, D_{i,j}(z, a_0) = 0) \right]$ and $\sum_{j=1}^J \sum_{i=1}^{N_j} \mathbb{1}(D_{i,j}(z, a_1) = 1, D_{i,j}(z, a_0) = 0)$, which is equivalent to

$$\left[\sum_{j=1}^J \sum_{i=1}^{N_j} \{Y_{i,j}(z, a_1) - Y_{i,j}(z, a_0)\} \mathbb{1}(G_{i,j} \in \{(c, a), (n, a)\}) \right] / \left[\sum_{j=1}^J \sum_{i=1}^{N_j} \mathbb{1}(G_{i,j} \in \{(c, a), (n, a)\}) \right]. \quad (2.4)$$

The estimand in equation (2.4) is interpreted as a local average treatment effect for units who comply with the treatment probability assigned to the cluster. These estimands are not clearly interpretable as complier spillover effects because the compliance is defined for the compliance behaviors with respect to treatment probability, not the actual treatment assignment. We define a new complier average spillover effect using base and target cluster assignments. For treatment z and base cluster assignment a' , we define $\text{CASE}(z, a')$ as the ratio of $\left[\sum_{j=1}^J \sum_{i=1}^{N_j} \{Y_{i,j}(z, a_1) - Y_{i,j}(z, a_0)\} \mathbb{1}(G_{i,j}(a') = c) \right]$ and $\sum_{j=1}^J \sum_{i=1}^{N_j} \mathbb{1}(G_{i,j}(a') = c)$. Similar to the population-level spillover effect on outcome defined in Section 2.4, $\text{CASE}(z, a')$ captures the compliers' local average spillover effect by comparing the two sets of potential outcomes under the two cluster assignments a_0 and a_1 and treatment z . In this estimand, the compliers are defined under the base cluster assignment a' , and it is interpretable as an average effect for the compliers of the treatment received by others in the same cluster.

We finally define the finite-population complier average overall effect by combining equation (2.1) with the unit-level complier average overall effect under the base cluster assignment a' , $\{\bar{Y}_{i,j}(a_1) - \bar{Y}_{i,j}(a_0)\} \mathbb{1}(G_{i,j}(a') = c)$, and decomposing the effect according to the sum of the unit-level complier average direct and spillover effects. Specifically, we have

$$\begin{aligned} \text{CAOE}_{i,j}(a') &= \left\{ \frac{K_j(a_1)}{N_j} \right\} \{Y_{i,j}(1, a_1) - Y_{i,j}(0, a_1)\} \mathbb{1}(G_{i,j}(a') = c) \\ &\quad - \left\{ \frac{K_j(a_0)}{N_j} \right\} \{Y_{i,j}(1, a_0) - Y_{i,j}(0, a_0)\} \mathbb{1}(G_{i,j}(a') = c) \\ &\quad + \{Y_{i,j}(0, a_1) - Y_{i,j}(0, a_0)\} \mathbb{1}(G_{i,j}(a') = c). \end{aligned} \quad (2.5)$$

The population-level effect is then the average of the unit-level effects in equation (2.5) over compliers under the base cluster assignment a' , i.e.,

$$\text{CAOE}(a') = \sum_{j=1}^J \sum_{i=1}^{N_j} \text{CAOE}_{i,j}(a') / \sum_{j=1}^J \sum_{i=1}^{N_j} \mathbb{1}(G_{i,j}(a') = c) \quad (2.6)$$

ITT _Y	Direct effect on the outcome
ITT _D	Direct effect on the treatment receipt
S _Y	Spillover effect on the outcome
S _D	Spillover effect on the treatment receipt
O _Y	Overall effect on the outcome
O _D	Overall effect on the treatment receipt
CADE	Complier average direct effect on the outcome
CASE	Complier average spillover effect on the outcome
CAOE	Complier average overall effect on the outcome

Table 1: Causal estimands defined in Sections 2.4 and 2.5.

To facilitate references to causal estimands in this manuscript, Table 1 summarizes the estimands that we defined in Sections 2.4 and 2.5.

3 Bayesian Model and Inferential Approach

3.1 Overview of Methodology

Our Bayesian methodology for performing causal inferences on two-stage randomized experiments with interference, nonadherence, and missing outcomes is based on model-based imputation of missing and unrealized potential outcomes to derive the posterior distributions of the finite-population causal estimands. To describe this more formally, we let τ denote one of the causal estimands from Sections 2.4 and 2.5, \mathbf{X} the $N \times P$ matrix of covariates for all experimental units, \mathbf{A}° the $N \times 1$ vector of treatment probabilities that were assigned to the clusters in the experiment, \mathbf{Z}° the $N \times 1$ vector of treatments that were assigned to the units in the experiment, \mathbf{D}° the $N \times 1$ vector of treatments that were received by the units in the realized experiment, \mathbf{M}° the $N \times 1$ vector of missingness indicators for the outcomes in the realized experiment, and \mathbf{Y}^{obs} the vector of observed outcomes in the realized experiment. Then the Bayesian methodology calculates the distribution $p(\tau \mid \mathbf{X}, \mathbf{A}^\circ, \mathbf{Z}^\circ, \mathbf{D}^\circ, \mathbf{M}^\circ, \mathbf{Y}^{\text{obs}})$. As causal estimands are functions of observed, missing, and unrealized potential outcomes, we calculate the posterior distribution above by integrating over the unrealized treatments received, which we denote by \mathbf{D}^u , and both the missing and unrealized potential outcomes, which we denote by \mathbf{Y}^{mis} and \mathbf{Y}^u respectively, according to

$$\begin{aligned}
 p(\tau \mid \mathbf{X}, \mathbf{A}^\circ, \mathbf{Z}^\circ, \mathbf{M}^\circ, \mathbf{D}^\circ, \mathbf{Y}^{\text{obs}}) &= p(\tau \mid \mathbf{O}) \\
 &= \int p(\tau \mid \mathbf{O}, \mathbf{D}^u, \mathbf{Y}^{\text{mis}}, \mathbf{Y}^u) p(\mathbf{D}^u, \mathbf{Y}^{\text{mis}}, \mathbf{Y}^u \mid \mathbf{O}) d\mathbf{D}^u d\mathbf{Y}^{\text{mis}} d\mathbf{Y}^u \\
 &= \int p(\tau \mid \mathbf{O}, \mathbf{G}, \mathbf{Y}^{\text{mis}}, \mathbf{Y}^u) p(\mathbf{G}, \mathbf{Y}^{\text{mis}}, \mathbf{Y}^u \mid \mathbf{O}) d\mathbf{G} d\mathbf{Y}^{\text{mis}} d\mathbf{Y}^u,
 \end{aligned}$$

where $\mathbf{O} = (\mathbf{X}, \mathbf{A}^\circ, \mathbf{Z}^\circ, \mathbf{M}^\circ, \mathbf{D}^\circ, \mathbf{Y}^{\text{obs}})$ and \mathbf{G} is an $N \times 1$ vector of (latent) principal strata memberships for the experimental units. This motivates our Bayesian methodology of first deriving a Monte Carlo approximation for the posterior predictive distribution of the missing principal strata and outcomes conditional on the observed data, and

then using that posterior predictive distribution to derive a Monte Carlo approximation for the posterior distribution of τ . Our methodology is distinct from existing methods for two-stage randomized experiments in that we impute missing principal strata memberships and at most four missing/unrealized potential outcomes for each experimental unit according to the assumptions in Section 2.

Our imputation approach, and the latent ignorability assumption that underlies it, is in Section 3.2. In Section 3.3 we provide additional details on our Bayesian modeling approach and the unconfoundedness condition underlying it, which is justified by the design of the two-stage randomized experiment. The Gibbs sampling algorithm that we use to derive a Monte Carlo approximation to the posterior distribution of the causal estimand is in Section 3.4. We provide a justification of exchangeability in our Bayesian methodology in the appendix.

3.2 Imputation of Missing Values

By virtue of Assumptions 6 and 7, any experimental unit in the two-stage randomized experiment belongs to one of six principal strata. Table 2 illustrates the correspondence between the possible values of $A_j^\circ, Z_{i,j}^\circ$, and $D_{i,j}^\circ$ and the principal strata memberships, along with the missingness rates for the outcomes in the RSBY study that we analyze in Section 5. We let $\mathcal{G}(a, z, d)$ denote a set of possible principal strata for units with $(A_j^\circ, Z_{i,j}^\circ, D_{i,j}^\circ) = (a, z, d)$. Assumption 7 in particular helps to narrow down the possible strata that a unit could belong to based on the observed data. For example, if an experimental unit i in cluster j has $A_j^\circ = a_0, Z_{i,j}^\circ = 0$, and $D_{i,j}^\circ = 1$, then $\mathcal{G}(a_0, 0, 1) = \{(a, a)\}$, and so $G_{i,j} = (a, a)$. Once $G_{i,j}$ is fixed, we can immediately impute all unrealized $D_{i,j}^u$.

We also observe in Table 2 that different strata, e.g., (a, a) and (n, n) , exhibit different missingness rates. This corresponds to more general phenomenon in modern experiments in which different strata have different missing data mechanisms for the final outcomes. This phenomenon is intuitive for certain strata. For example, those units in the (n, n) strata may be more reluctant to take any actions compared to units in other strata, and thus their outcomes are more likely to be missing. As the missingness mechanism depends on the latent strata, the traditional missing completely at random (MCAR) and missing at random (MAR) assumptions may not be appropriate. In contrast to other methods, our Bayesian methodology accommodates a class of missing not at random (MNAR) mechanisms (Little and Rubin, 2002, p. 351) that correspond to the following latent ignorability (Frangakis and Rubin, 1999, LI) assumption.

Assumption 9 (Latent Ignorability of Missing Data). *For any unit $i = 1, \dots, N_j$ in cluster $j = 1, \dots, J$,*

$$\begin{aligned} p(\mathbf{M}_{i,j} \mid \mathbf{D}_{i,j}, \mathbf{Y}_{i,j}, X_{i,j}) \\ = p(\mathbf{M}_{i,j} \mid \mathbf{D}_{i,j}, X_{i,j}), \quad p(\mathbf{M}_{i,j} \mid G_{i,j}, \mathbf{Y}_{i,j}, X_{i,j}) = p(\mathbf{M}_{i,j} \mid G_{i,j}, X_{i,j}), \end{aligned}$$

where $\mathbf{D}_{i,j} = (D_{i,j}(0, a_0), D_{i,j}(0, a_1), D_{i,j}(1, a_0), D_{i,j}(1, a_1))$, $\mathbf{M}_{i,j} = (M_{i,j}(0, a_0), M_{i,j}(0, a_1), M_{i,j}(1, a_0), M_{i,j}(1, a_1))$, and $\mathbf{Y}_{i,j} = (Y_{i,j}(0, a_0), Y_{i,j}(0, a_1), Y_{i,j}(1, a_0), Y_{i,j}(1, a_1))$.

A_j^o	$Z_{i,j}^o$	$D_{i,j}^o$	$\mathcal{G}(A_j^o, Z_{i,j}^o, D_{i,j}^o)$	Number of units	Missing rates
a_0	0	0	$(c, c), (n, n), (c, a), (n, c), (n, a)$	2258	0.095
a_0	1	0	$(n, n), (n, c), (n, a)$	545	0.119
a_0	0	1	(a, a)	911	0.043
a_0	1	1	$(c, c), (a, a), (c, a)$	1554	0.062
a_1	0	0	$(c, c), (n, n), (n, c)$	779	0.105
a_1	1	0	(n, n)	1028	0.135
a_1	0	1	$(a, a), (c, a), (n, a)$	350	0.031
a_1	1	1	$(c, c), (a, a), (c, a), (n, c), (n, a)$	3454	0.059

Table 2: Illustration of the possible principal strata that an experimental unit can belong to under different combinations of A_j^o , $Z_{i,j}^o$, and $D_{i,j}^o$. We provide examples of the number of units and the missingness rates of the outcomes for these different sets of principal strata based on data from the RSBY study in Section 5.

We formulate LI to mean that if we knew the latent strata $G_{i,j}$ of each unit, the missingness mechanism would be ignorable. Assumptions 8 and 9 facilitate and unify our imputations of missing and unrealized potential outcomes. For those units with $M_{i,j}^o = 0$, as one of their potential outcomes is observed we need to impute at most three unrealized potential outcomes for each of them. In the case of a unit whose principal stratum is (c, c) , we need to impute three unrealized potential outcomes based on our Bayesian model. For another unit whose principal stratum is (a, a) , we only need to impute one unrealized potential outcome, as the potential outcomes are restricted according to Assumption 8 for this stratum. For a unit with $M_{i,j}^o = 1$, we impute at most four potential outcomes based on their imputed principal strata $G_{i,j}$.

3.3 Bayesian Model Building

Following the Bayesian paradigm of Imbens and Rubin (2015) and Gelman et al. (2013), five inputs are necessary for deriving the posterior distributions of the causal estimands. The first input is knowledge of the assignment mechanisms as encoded in the probability mass functions $p(\mathbf{A} | \mathbf{X})$ and $p(\mathbf{Z} | \mathbf{X}, \mathbf{A})$. These two probability mass functions are obtained immediately by virtue of the known design of the two-stage randomized experiment. In addition, these two probability mass functions are independent of the experimental units' covariates \mathbf{X} , so that $p(\mathbf{A} | \mathbf{X}) = p(\mathbf{A})$ and $p(\mathbf{Z} | \mathbf{X}, \mathbf{A}) = p(\mathbf{Z} | \mathbf{A})$. The second input is the model for treatment received conditional on $\mathbf{X}, \mathbf{A}, \mathbf{Z}$. We denote this input by the probability mass function $p(\mathbf{D} | \mathbf{X}, \mathbf{A}, \mathbf{Z}, \psi)$ with parameter vector ψ . An equivalent input is the model for principal strata memberships of all experimental units, \mathbf{G} , conditional on $\mathbf{X}, \mathbf{A}, \mathbf{Z}$, because there is a one-to-one mapping between \mathbf{G} and \mathbf{D} . To simplify our notation we denote this input as $p(\mathbf{G} | \mathbf{X}, \mathbf{A}, \mathbf{Z}, \psi)$ with the understanding that ψ is generic notation for the parameter vector associated with either of these two models. The third input is the model for the potential outcomes conditional on $\mathbf{X}, \mathbf{A}, \mathbf{Z}, \mathbf{D}$. We denote this model via a probability density/mass function $p(\mathbf{Y} | \mathbf{X}, \mathbf{A}, \mathbf{Z}, \mathbf{D}, \theta)$ with parameter vector θ . This model can also be equivalently formulated using \mathbf{G} as $p(\mathbf{Y} | \mathbf{X}, \mathbf{A}, \mathbf{Z}, \mathbf{G}, \theta)$. The fourth input is the model for the miss-

ingness indicators \mathbf{M} conditional on $\mathbf{X}, \mathbf{A}, \mathbf{Z}, \mathbf{D}, \mathbf{Y}$, denoted by $p(\mathbf{M} \mid \mathbf{X}, \mathbf{A}, \mathbf{Z}, \mathbf{D}, \mathbf{Y}, \phi)$ with parameter vector ϕ . Under the LI assumption, this model will not involve \mathbf{Y} conditional on the other variables. The final input is the joint prior distribution for ψ , ϕ , and θ , denoted by $p(\psi, \phi, \theta)$. We assume throughout that ψ , ϕ , and θ are distinct and do not share any common parameters.

Specifying the second and third inputs is facilitated by the fact that the two-stage randomized experiment has an unconfounded assignment mechanism at both cluster and individual experimental unit levels. Unconfoundedness is formally expressed via the following assumption.

Assumption 10. *For experimental unit $i = 1, \dots, N_j$ in cluster $j = 1, \dots, J$, $p(A_j, Z_{i,j} \mid X_{i,j}, \mathbf{D}_{i,j}) = p(A_j, Z_{i,j} \mid X_{i,j})$, $p(A_j, Z_{i,j} \mid X_{i,j}, \mathbf{Y}_{i,j}, \mathbf{D}_{i,j}) = p(A_j, Z_{i,j} \mid X_{i,j}, \mathbf{D}_{i,j})$, $p(A_j, Z_{i,j} \mid X_{i,j}, \mathbf{Y}_{i,j}, G_{i,j}) = p(A_j, Z_{i,j} \mid X_{i,j}, G_{i,j})$ and $p(A_j, Z_{i,j}, \mathbf{M}_{i,j} \mid X_{i,j}, \mathbf{D}_{i,j}, \mathbf{Y}_{i,j}) = p(A_j, Z_{i,j} \mid X_{i,j}, \mathbf{D}_{i,j}, \mathbf{Y}_{i,j})p(\mathbf{M}_{i,j} \mid X_{i,j}, \mathbf{D}_{i,j}, \mathbf{Y}_{i,j})$.*

Assumption 10 implies that the treatment assignment mechanism is ignorable. In addition, unit exchangeability implies by de Finetti's Theorem that we can specify parametric models on the level of the individual experimental units for the principal strata memberships $G_{i,j}$, the potential outcomes, and the missingness indicators to derive the joint model for \mathbf{G} , \mathbf{M} , and \mathbf{Y} . A justification of unit exchangeability in the two-stage randomized experiment is in the appendix. Hence we have that

$$\begin{aligned} p(\mathbf{Y}, \mathbf{D}, \mathbf{M} \mid \mathbf{X}, \psi, \theta, \phi) \\ &= \prod_{i,j} p(\mathbf{D}_{i,j} \mid X_{i,j}, \psi) p(\mathbf{Y}_{i,j} \mid X_{i,j}, \mathbf{D}_{i,j}, \theta) p(\mathbf{M}_{i,j} \mid X_{i,j}, \mathbf{D}_{i,j}, \mathbf{Y}_{i,j}, \phi) \\ &= \prod_{i,j} p(\mathbf{D}_{i,j} \mid X_{i,j}, \psi) p(\mathbf{Y}_{i,j} \mid X_{i,j}, \mathbf{D}_{i,j}, \theta) p(\mathbf{M}_{i,j} \mid X_{i,j}, \mathbf{D}_{i,j}, \phi). \end{aligned} \quad (3.1)$$

The last line in equation (3.1) follows from Assumption 9. The posterior distribution $p(\psi, \phi, \theta \mid \mathbf{X}, \mathbf{A}^\circ, \mathbf{Z}^\circ, \mathbf{D}^\circ, \mathbf{M}^\circ, \mathbf{Y}^{\text{obs}})$ is proportional to

$$\begin{aligned} p(\psi, \phi, \theta \mid \mathbf{X}, \mathbf{A}^\circ, \mathbf{Z}^\circ, \mathbf{D}^\circ, \mathbf{M}^\circ, \mathbf{Y}^{\text{obs}}) \\ &\propto p(\psi, \phi, \theta) \int \prod_{i,j} p(\mathbf{D}_{i,j} \mid X_{i,j}, \psi) p(\mathbf{Y}_{i,j} \mid X_{i,j}, \mathbf{D}_{i,j}, \theta) p(M_{i,j}^\circ \mid X_{i,j}, \mathbf{D}_{i,j}, \theta) d\mathbf{D}^u d\mathbf{Y}^{\text{mis}} d\mathbf{Y}^u \\ &= p(\psi, \phi, \theta) \int \prod_{i,j} p(\mathbf{G}_{i,j} \mid X_{i,j}, \psi) p(\mathbf{Y}_{i,j} \mid X_{i,j}, \mathbf{G}_{i,j}, \theta) p(M_{i,j}^\circ \mid X_{i,j}, \mathbf{G}_{i,j}, \theta) d\mathbf{G} d\mathbf{Y}^{\text{mis}} d\mathbf{Y}^u \end{aligned}$$

To express the posterior distribution of the model parameters, let $O(a, z, d, m)$ denote the set of experimental units whose realized values of $A_j^\circ, Z_{i,j}^\circ, D_{i,j}^\circ$, and $M_{i,j}^\circ$ are a, z, d , and m , respectively. For each unit i in cluster j and each principal strata g , define $w_{i,j}^{(g)} = \Pr(G_{i,j} = g \mid X_{i,j}, \psi)$, $\rho_{i,j}^{(g,a,z)} = \Pr(M_{i,j} = 1 \mid G_{i,j} = g, A_j = a, Z_{i,j} = z, X_{i,j}, \phi)$, and $f_{i,j}^{(g,a,z)}$ be the probability mass/density function of $Y_{i,j}(z, a)$ for $g \in \{(n, n), (c, c), (a, a), (n, c), (n, a), (c, a)\}$. Under Assumptions 6 - 8, the posterior

distribution of ψ , ϕ , and θ can be written as

$$\begin{aligned}
 p(\psi, \phi, \theta \mid \mathbf{X}, \mathbf{A}^o, \mathbf{Z}^o, \mathbf{D}^o, \mathbf{M}^o, \mathbf{Y}^{\text{obs}}) &\propto p(\psi, \phi, \theta) \\
 &\times \prod_{a \in \{a_0, a_1\}} \prod_{z \in \{0, 1\}} \prod_{d \in \{0, 1\}} \prod_{(i, j) \in O(a, z, d, 0)} \sum_{g \in \mathcal{G}(a, z, d)} (1 - \rho_{i, j}^{g, z, d}) w_{i, j}^g f_{i, j}^{g, z, d} \\
 &\times \prod_{a \in \{a_0, a_1\}} \prod_{z \in \{0, 1\}} \prod_{d \in \{0, 1\}} \prod_{(i, j) \in O(a, z, d, 1)} \sum_{g \in \mathcal{G}(a, z, d)} \rho_{i, j}^{g, z, d} w_{i, j}^g f_{i, j}^{g, z, d}.
 \end{aligned} \tag{3.2}$$

Our model-based Bayesian method ultimately utilizes the above mixture model that requires the specification of $w_{i, j}^{(g)}$, $\rho_{i, j}^{(g, a, z)}$ and $f_{i, j}^{(g, a, z)}$. Specific models and prior distributions are provided in our simulation studies in Section 4 and our case study in Section 5. Our Bayesian methodology also enables us to infer principal causal effects under two-sided noncompliance, whereas existing methods require additional strict assumptions to perform inferences for principal causal effects under two-sided noncompliance including always-takers and social-interaction compliers (i.e., the strata (a, a) and (c, a)) (Vazquez-Bare, 2022).

3.4 Gibbs Sampling Algorithm

We utilize the Gibbs sampling algorithm (Geman and Geman, 1984; Gelfand and Smith, 1990; Imbens and Rubin, 1997) to obtain the posterior predictive distributions of \mathbf{Y}^{mis} and \mathbf{Y}^u given the observed data, so as to impute all missing and unrealized outcomes and derive the posterior distributions for the causal estimands. Specifically, we iterate between drawing from the conditional distributions of (ψ, ϕ, θ) and \mathbf{G} given the other variables, respectively. Then for each iteration we impute \mathbf{Y}^{mis} and \mathbf{Y}^u , and calculate the causal estimands of interest to effectively get a draw from their respective posterior distributions. The essential algorithm is outlined below.

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

This Gibbs sampler enables us to obtain posterior draws of all the unobserved potential outcomes and principal strata for all units. We then perform causal inferences by means of the posterior distributions. For example, points estimates of the estimands

can be obtained via their posterior means or medians, and intervals for the estimands can be obtained via the central credible intervals. Detailed derivations for each step of the Gibbs sampler are provided in the supplementary material (Ohnishi and Sabbaghi, 2022).

4 Simulation Studies

4.1 Evaluation Metrics and Data Generating Mechanisms

We evaluate the frequentist properties of our Bayesian methodology with respect to those of the method of Imai et al. (2021), which we implement using its R package `experiment` (Imai et al., 2019). Imai et al. (2021) implicitly assumed MCAR data in their analyses, and so removed rows with missing outcomes from their analyses because their methodology does not accommodate missing outcomes. Therefore, to ensure fair comparisons, in our first simulation studies we have either no missing data or MCAR data. Our Bayesian methodology easily accommodates the MCAR mechanism without major modifications, as this would just entail removing units that have missing outcomes and ignoring the missing data mechanism model in the derivations of the likelihood functions and posteriors.

The evaluation metrics that we consider are bias and mean square error (MSE) in estimating a causal estimand, coverage of an interval estimator for a causal estimand, and the interval length. Bias and MSE are generally defined as $\sum_{m=1}^M (\tau - \hat{\tau}_m) / M$ and $\sum_{m=1}^M (\tau - \hat{\tau}_m)^2 / M$ respectively, where M denotes the number of simulated datasets, τ denotes the true causal estimand, and $\hat{\tau}_m$ denotes the estimate of the causal estimand in dataset $m = 1, \dots, M$. For our Bayesian method, the point estimator is the median of the posterior distribution of a causal estimand, and the interval estimator is the 95% central credible interval. Our summary of the interval length is the median of the lengths of the credible intervals computed from M simulated datasets. The median is desirable as it is more robust for the chosen data generating mechanism. The data generating mechanisms that we consider in our study are specified to simulate data resembling those in our case study in Section 5.

For the two-stage randomized experiments in our simulation study, the clusters have equal numbers of units (i.e., $N_j = N/J$), and we specify $p(A_j = a_0) = p(A_j = a_1) = 0.5$, $p(Z_{i,j} = 1 \mid A_j = a_0) = 0.4$, and $p(Z_{i,j} = 1 \mid A_j = a_1) = 0.8$ for any experimental unit $i = 1, \dots, N_j$ in cluster $j = 1, \dots, J$. Each experimental unit belongs to just one of the latent principal strata $\{(n, n), (c, c), (a, a), (n, c), (n, a), (c, a)\}$. The principal strata memberships are generated according to $G_{i,j} \sim \text{Multinomial}(\pi_{(n,n)}, \pi_{(c,c)}, \pi_{(a,a)}, \pi_{(n,c)}, \pi_{(n,a)}, \pi_{(c,a)})$. Evaluation for different values of J is presented in the appendix.

The potential outcomes in our simulation study are continuous and are generated to have a right-skewed distribution with an excess of zeros. The specific generation mechanism that we implement is a zero-inflated Log-Normal distribution, with the parameters of the underlying Bernoulli and Log-Normal random variables (representing the excess zeros and heavy tail of the outcomes, respectively) specified to be distinct for each strata and treatment. More formally, for $a \in \{a_0, a_1\}$ and $z \in \{0, 1\}$, the potential outcomes for

unit i in cluster j are generated by first sampling $W_{i,j}(z, a) \sim \text{Bernoulli}(p_{z,a,G_{i,j}})$, then sampling $\tilde{Y}_{i,j}(z, a) \sim \text{Log-Normal}(\mu_{z,a,G_{i,j}}, \sigma_{z,a,G_{i,j}}^2)$, and finally generating $Y_{i,j}(z, a) = \{1 - W_{i,j}(z, a)\} \tilde{Y}_{i,j}(z, a)$. For this simulation study, we assume that potential outcomes are generated independently of one another. In addition, Assumption 8 applies throughout the data generating mechanism, including for $W_{i,j}(z, a)$ and $\tilde{Y}_{i,j}(z, a)$. The generated $W_{i,j}(z, a)$ and $\tilde{Y}_{i,j}(z, a)$ are not recorded as data. Finally, we use conjugate prior distributions for all parameters, that is, $(\pi_{(n,n)}, \pi_{(c,c)}, \pi_{(a,a)}, \pi_{(n,c)}, \pi_{(n,a)}, \pi_{(c,a)}) \sim \text{Dirichlet}(2, 2, 2, 2, 2, 2)$, $p_{z,a,G_{i,j}} \sim \text{Beta}(1, 1)$, $\mu_{z,a,G_{i,j}} \sim \text{Normal}(0, 10^2)$ and $\sigma_{z,a,G_{i,j}}^2 \sim \text{InverseGamma}(1, 1)$. Detailed algorithms for the Gibbs sampler are provided in the appendix. We simulate 500 datasets for each $N = 5000, 10000, 50000$ with a fixed cluster size of $J = 100$. In the appendix we provide all parameter values that are utilized in our simulation study the frequentist evaluations for the super-population versions of the causal estimands, and an additional simulation study using the Gamma distribution for the outcome model.

4.2 Results

Table 3 summarizes the results of our simulation study in the case of the Log-Normal distribution. The method of Imai et al. (2021) is abbreviated as “IJM” in this table. We observe that both methods perform well with respect to coverage and bias, with the IJM method exhibiting less bias than our Bayesian method for small sample sizes. This difference in bias can be attributed to the effect of the prior distributions, which is not negligible for small N . In addition, the posterior median is not an unbiased estimator of the mean of the Log-Normal distribution, although it is arguably a desirable estimator as it should be more robust for the chosen data generating mechanism. This difference in bias should also be expected as the emphasis of the Bayesian approach is on the posterior distribution and whether it is well-calibrated. Finally, we recognize that our Bayesian methodology outperforms the IJM method with respect to MSE under all conditions. The differences in MSE imply that the Bayesian point estimator is less likely to deviate from the true causal effect over multiple experiments as compared to the IJM method, that the Bayesian estimator has less variability, and that the Bayesian intervals have smaller widths. Ultimately, the Bayesian methodology yields more precise causal inferences. The IJM method is sensitive to the shape of the distribution and exhibits greater variability and MSEs as it attempts to accommodate outliers.

4.3 Evaluation under Misspecification

We evaluate our Bayesian methodology under model misspecification. To generate data that resemble those in our case study, we consider heavy-tailed distributions for the data-generating process. We use the Half-Student-t distribution with different values of the degrees of freedom, $\nu = 1.5, 2.0, 3.0, 5.0$ to evaluate how the outliers and skewness of the distribution influence our results. We did not choose $\nu = 1$ because that corresponds to the Cauchy distribution, which does not have a finite first moment. The outcomes are generated by $\tilde{Y}_{i,j}(z, a) \sim C_{z,a,G_{i,j}} \times \text{Half-t}(\nu)$, where $C_{z,a,G_{i,j}}$ is a scaling constant. We fit the Log-Normal model described in the previous section to the generated data.

	N	Coverage		Bias		MSE		Interval Width	
		IJM	Bayes	IJM	Bayes	IJM	Bayes	IJM	Bayes
CADE(0, a_0)	5000	96%	99%	4.43E+02	5.02E+02	3.60E+07	4.57E+06	1.91E+04	9.11E+03
	10000	96%	98%	-2.52E+02	2.99E+02	1.39E+07	1.94E+06	1.36E+04	5.41E+03
	50000	94%	95%	-1.22E+02	1.59E+02	3.58E+06	2.71E+05	6.61E+04	2.09E+03
CADE(1, a_1)	5000	92%	97%	-1.05E+03	1.20E+03	2.28E+08	5.97E+06	4.52E+04	1.67E+04
	10000	93%	97%	2.12E+02	8.92E+02	1.29E+08	3.61E+06	3.38E+04	1.13E+04
	50000	94%	97%	4.81E+02	2.98E+02	2.18E+07	7.75E+05	1.68E+05	4.81E+03
ITT _{Y,·,·} (a_0)	5000	96%	98%	1.68E+02	2.46E+02	9.11E+06	1.20E+06	9.63E+03	4.68E+03
	10000	96%	98%	-1.21E+02	1.46E+02	3.49E+06	4.86E+05	6.86E+03	2.80E+03
	50000	93%	96%	-5.61E+01	8.06E+01	9.08E+05	7.39E+04	3.32E+04	1.08E+03
ITT _{Y,·,·} (a_1)	5000	92%	97%	-4.54E+02	6.00E+02	4.62E+07	1.42E+06	2.03E+04	7.76E+03
	10000	94%	97%	1.18E+02	4.41E+02	2.64E+07	8.34E+05	1.52E+04	5.30E+03
	50000	94%	97%	2.31E+02	1.42E+02	4.45E+06	1.69E+05	7.54E+04	2.23E+03
ITT _{D,·,·} (a_0)	5000	87%	95%	-1.56E-02	2.60E-03	5.63E-04	2.22E-04	6.96E-02	5.70E-02
	10000	93%	94%	-2.26E-04	-1.47E-03	1.74E-04	1.21E-04	4.92E-02	4.06E-02
	50000	94%	95%	-6.52E-04	-1.52E-04	3.31E-05	2.20E-05	2.20E-02	1.83E-02
ITT _{D,·,·} (a_1)	5000	94%	94%	-7.57E-03	4.87E-03	5.34E-04	4.10E-04	9.06E-02	7.64E-02
	10000	93%	93%	-2.34E-03	3.76E-03	2.85E-04	2.24E-04	6.40E-02	5.51E-02
	50000	94%	96%	-2.36E-03	4.44E-04	5.29E-05	3.70E-05	2.86E-02	2.53E-02

Table 3: Evaluation metrics for our Bayesian methodology versus the method of Imai et al. (2021) (abbreviated as “IJM”) under the Log-Normal data-generating process and the finite-population perspective.

We fix the number of units and the number of clusters at $N = 10000$ and $J = 100$ respectively. The additional parameters $C_{z,a,G_{i,j}}$ are provided in the supplementary material (Ohnishi and Sabbaghi, 2022).

Table 4 summarizes the results of the simulation study in the case of model misspecification. Our Bayesian model outperforms the IJM method for the skewed distribution, $\nu = 1.5$. It also exhibits smaller MSE and slightly larger bias for $\nu = 2.0, 3.0$ for the case of the CADE and ITT_Y estimands. The larger bias is attributable to the same reasons as those outlined in the explanation of the results for the first simulation study (Section 4.2). When the data distribution is less skewed and less heavy-tailed, i.e., when $\nu = 5.0$, the IJM method performs well. This results suggest that one should adopt more appropriate outcome models for less skewed data. On the other hand, our model is assumed to be more appropriate for the data in the case study that is considered in Section 5, as those data are heavily skewed and exhibit substantial outliers on the tail.

4.4 Evaluation under MNAR

Table 5 presents the simulation results under MNAR. The underlying parametrizations are the same as the Log-normal mechanism in Section 4.1, except that the outcome is missing with probabilities given in the supplementary material (Ohnishi and Sabbaghi, 2022). It is important to recognize that the IJM method is evaluated after applying listwise deletion of the rows with missing outcome. We see that for $N = 5000$ the interval lengths of the Bayesian method are a bit larger than those in Table 3. This is due to incorporating the missingness mechanism into our model. For sufficiently large data (i.e., $N = 10000, 50000$), the results are similar to those in Table 3.

	ν	Coverage		Bias		MSE		Interval Width	
		IJM	Bayes	IJM	Bayes	IJM	Bayes	IJM	Bayes
CADE($0, a_0$)	1.5	94%	92%	8.38E-01	4.43E-01	3.52E+02	2.78E+01	2.08E+01	2.19E+01
	2.0	92%	93%	-7.55E-02	-4.02E-01	7.21E+00	3.70E+00	7.77E+00	8.39E+00
	3.0	95%	96%	-5.92E-02	2.51E-02	1.11E+00	8.78E-01	4.01E+00	4.10E+00
	5.0	93%	84%	-4.01E-02	5.29E-01	5.65E-01	8.85E-01	2.84E+00	2.88E+00
CADE($1, a_1$)	1.5	97%	90%	2.01E+00	-1.64E+00	7.36E+02	1.43E+02	2.58E+01	2.61E+01
	2.0	98%	96%	-6.03E-02	-4.60E-01	1.24E+01	7.57E+00	1.07E+01	1.08E+01
	3.0	96%	98%	3.72E-02	-1.43E-01	2.10E+00	1.50E+00	5.42E+00	5.96E+00
	5.0	96%	93%	-1.41E-01	-6.70E-01	9.33E-01	1.62E+00	3.83E+00	4.39E+00
ITT _{Y,\cdot,\cdot} (a_0)	1.5	94%	92%	3.60E-01	1.61E-01	6.37E+01	5.02E+00	8.99E+00	9.50E+00
	2.0	93%	94%	-2.74E-02	-1.94E-01	1.34E+00	6.90E-01	3.37E+00	3.54E+00
	3.0	96%	97%	-1.92E-02	-1.10E-02	2.11E-01	1.61E-01	1.74E+00	1.75E+00
	5.0	93%	86%	-1.21E-02	2.02E-01	1.07E-01	1.50E-01	1.23E+00	1.21E+00
ITT _{Y,\cdot,\cdot} (a_1)	1.5	96%	89%	9.11E-01	-7.06E-01	1.14E+02	2.15E+01	1.03E+01	9.90E+00
	2.0	98%	96%	-3.04E-02	-2.30E-01	2.01E+00	1.16E+00	4.24E+00	4.28E+00
	3.0	95%	98%	1.23E-02	-1.10E-01	3.37E-01	2.36E-01	2.20E+00	2.30E+00
	5.0	96%	87%	-5.61E-02	-3.10E-01	1.58E-01	2.79E-01	1.55E+00	1.64E+00
ITT _{D,\cdot,\cdot} (a_0)	1.5	95%	93%	1.57E-03	-4.14E-03	1.51E-04	1.97E-04	5.09E-02	4.70E-02
	2.0	95%	85%	1.57E-03	-5.87E-03	1.51E-04	2.45E-04	5.09E-02	4.66E-02
	3.0	95%	77%	1.57E-03	-6.43E-03	1.51E-04	3.46E-04	5.09E-02	4.60E+00
	5.0	95%	67%	1.57E-03	-7.54E-03	1.51E-04	4.85E-04	5.09E-02	4.46E-02
ITT _{D,\cdot,\cdot} (a_1)	1.5	95%	76%	-1.24E-03	-1.27E-03	2.78E-04	6.22E-04	6.52E-02	5.90E-02
	2.0	95%	69%	-1.24E-03	-1.53E-02	2.78E-04	7.32E-04	6.52E-02	5.75E-02
	3.0	95%	59%	-1.24E-03	-1.77E-02	2.78E-04	1.00E-03	6.52E-02	5.55E-02
	5.0	95%	47%	-1.24E-03	2.43E-02	2.78E-04	1.53E-03	6.52E-02	5.29E-02

Table 4: Evaluation metrics for our Bayesian methodology versus the IJM method under model misspecification.

The performance of the IJM method in Table 5 is worse in all metrics compared to the IJM results in Table 3. This is due to the bias resulting from applying listwise deletion of rows under the assumption of an MCAR mechanism when the actual missing data mechanism is MNAR. The bias becomes more severe when the missingness probabilities are significantly large. We present additional simulations with larger missing probabilities in the appendix.

5 Case Study: The Rashtriya Swasthya Bima Yojana Health Insurance Dataset

5.1 Description of Experiment and Data

Approximately 63 million people in India are below the poverty line due to health care expenditures. In 2008 a large-scale national hospital insurance plan for the poor was launched. This plan is known as the Rashtriya Swasthya Bima Yojana (RSBY). It is a large-scale national hospital insurance plan that households below the poverty line can join with a nominal co-payment. Under RSBY, these households can be covered for up to five people and more than 700 medical treatments and procedures, with the price set by the government. Medical services are provided nationwide by government-contracted public and private hospitals. Beneficiaries use their RSBY biometric ID cards, elimi-

	N	Coverage		Bias		MSE		Interval Width	
		IJM	Bayes	IJM	Bayes	IJM	Bayes	IJM	Bayes
CADE(0, a_0)	5000	96%	98%	6.12E+02	4.65E+02	3.88E+07	4.77E+06	2.06E+04	3.24E+04
	10000	98%	98%	-2.52E+02	2.02E+02	1.58E+07	1.12E+06	1.48E+04	1.41E+04
	50000	94%	96%	-3.25E+02	1.22E+02	3.97E+06	5.21E+05	7.07E+03	2.87E+03
CADE(1, a_1)	5000	92%	96%	-1.27E+03	1.84E+03	2.82E+08	9.11E+06	4.91E+04	5.43E+04
	10000	93%	97%	6.10E+02	1.25E+03	1.47E+08	6.07E+06	3.71E+04	3.92E+04
	50000	96%	98%	3.65E+02	5.31E+03	2.56E+07	5.93E+06	1.84E+04	1.26E+04
ITT _{Y,·,·} (a_0)	5000	96%	98%	3.34E+02	2.22E+02	9.50E+06	1.26E+06	1.00E+04	1.60E+04
	10000	97%	98%	1.94E+02	9.60E+01	3.83E+06	3.02E+05	7.30E+03	7.06E+03
	50000	93%	96%	-2.14E+02	6.24E+01	9.86E+05	1.38E+05	3.47E+03	1.45E+03
ITT _{Y,·,·} (a_1)	5000	92%	95%	-4.79E+02	9.28E+02	5.51E+07	2.25E+06	2.20E+04	2.50E+04
	10000	93%	97%	2.66E+02	6.20E+02	2.83E+07	1.40E+06	1.63E+04	1.80E+04
	50000	96%	98%	1.30E+02	2.49E+02	5.03E+06	1.22E+06	8.17E+03	5.77E+03
ITT _{D,·,·} (a_0)	5000	94%	91%	1.89E-03	4.81E-03	3.43E-04	3.33E-04	7.10E-02	5.94E-02
	10000	72%	92%	-1.71E-02	-2.36E-03	4.72E-04	1.62E-04	5.04E-02	4.23E-02
	50000	47%	93%	-1.13E-02	-2.22E-04	1.62E-04	2.50E-05	2.24E-02	1.86E-02
ITT _{D,·,·} (a_1)	5000	96%	91%	8.62E-04	9.62E-03	5.47E-04	5.09E-04	9.83E-02	7.29E-02
	10000	89%	91%	-1.09E-02	5.62E-03	4.42E-04	2.50E-04	6.90E-02	5.43E-02
	50000	89%	95%	-6.75E-03	7.27E-04	1.00E-04	3.80E-05	3.08E-02	2.49E-02

Table 5: Evaluation for our Bayesian methodology under MNAR. The IJM method is evaluated after applying listwise deletion of rows with missingness.

nating the need for cash transactions and insurance claims. Additional information and references for RSBY are provided by Nandi et al. (2015).

Imai et al. (2021) conducted a two-stage randomized experiment to determine whether access to the national insurance plan provided by RSBY increases access to hospitals and reduces impoverishment due to high medical expenses. This experiment consisted of $N = 10,072$ households after Imai et al. (2021) applied listwise deletion for missing outcomes, with the households residing in $J = 435$ villages. Of these villages, $J_1 = 219$ were assigned treatment probability $a_1 = 0.8$ and the remaining 216 villages were assigned treatment probability $a_0 = 0.4$. One concern in this experiment was the spillover effects between households, because one household's enrollment in RSBY may depend on the treatments assigned to other households. Another concern is that some households assigned treatment may decide to not enroll in RSBY, and some households assigned control may ultimately manage to enroll in RSBY.

5.2 Bayesian Analyses of the Experiment

We utilize our Bayesian methodology to infer the direct and spillover effects accounting for interference, treatment nonadherence, and missing outcomes on the annual household hospital expenditure outcome (which ranges from 0 to INR 500,000). We first consider the case of Imai et al. (2021) in which the outcomes are assumed to be MCAR. In this case, listwise deletion of rows with missingness is performed. We use the same mixture model as in Section 4.1 to analyze the data. The model for principal strata memberships is the Multinomial distribution, $G_{i,j} \sim \text{Multinomial}(\pi_{(n,n)}, \pi_{(c,c)}, \pi_{(a,a)}, \pi_{(n,c)}, \pi_{(n,a)}, \pi_{(c,a)})$. For the potential outcomes we specify a zero-inflated Log-Normal distribution for each principal stratum, that is, for $a \in \{a_0, a_1\}$ and $z \in \{0, 1\}$, the potential outcomes

	Mean	SD	Median	95% interval	IJM Est.	IJM SD
CADE(a_1, a_1)	-2041	7247	-1813	(-3782, -256)	-1649	1061
CADE(a_0, a_0)	298	3912	158	(-1356, 1982)	1984	1215
ITT _{Y,.,.} (a_1)	-853	3040	-759	(-1586, -106)	-795	514
ITT _{Y,.,.} (a_0)	139	1811	74	(-632, 913)	875	530
ITT _{D,.,.} (a_1)	0.418	0.010	0.418	(0.397, 0.438)	0.482	0.023
ITT _{D,.,.} (a_0)	0.465	0.009	0.465	(0.446, 0.483)	0.441	0.021
S _{Y,.,.} (1)	-1129	1795	-1071	(-1741, -459)	-1374	823
S _{Y,.,.} (0)	-136	3044	-222	(-1003, 666)	297	858
S _{D,.,.} (1)	0.030	0.007	0.029	(0.018, 0.047)	0.086	0.053
S _{D,.,.} (0)	0.077	0.009	0.077	(0.060, 0.095)	0.045	0.028

Table 6: Comparisons of the causal inferences obtained from our Bayesian methodology with those obtained from the method of IJM.

for unit i in cluster j are determined by $Y_{i,j}(z, a) = \{1 - W_{i,j}(z, a)\} \tilde{Y}_{i,j}(z, a)$ where $W_{i,j}(z, a) \sim \text{Bernoulli}(p_{z,a,G_{i,j}})$ and $\tilde{Y}_{i,j}(z, a) \sim \text{Log-Normal}(\mu_{z,a,G_{i,j}}, \sigma_{z,a,G_{i,j}}^2)$. Finally, we use conjugate prior distributions for all parameters. $(\pi_{(n,n)}, \pi_{(c,c)}, \pi_{(a,a)}, \pi_{(n,c)}, \pi_{(n,a)}, \pi_{(c,a)}) \sim \text{Dirichlet}(\alpha_{(n,n)}, \alpha_{(c,c)}, \alpha_{(a,a)}, \alpha_{(n,c)}, \alpha_{(n,a)}, \alpha_{(c,a)})$, $p_{z,a,G_{i,j}} \sim \text{Beta}(a_0, b_0)$, $\mu_{z,a,G_{i,j}} \sim \text{Normal}(\mu_0, \sigma_0^2)$ and $\sigma_{z,a,G_{i,j}}^2 \sim \text{InverseGamma}(k_0, \theta_0)$ where θ_0 is a scale parameter. For the hyperparameters we choose $\alpha_{(n,n)} = \alpha_{(c,c)} = \alpha_{(a,a)} = \alpha_{(n,c)} = \alpha_{(n,a)} = \alpha_{(c,a)} = 1$, $a_0 = b_0 = 1$, $\mu_0 = 0$, $\sigma_0^2 = 5$, $k_0 = 0.1$ and $\theta_0 = 1$. Note that $\sigma_0^2 = 5$ is sufficiently large on a log scale. Our MCMC algorithm was performed for 100,000 iterations with a burn-in of 50,000 draws.

We consider the finite-population inference as was done by Imai et al. (2021). One advantage of the alternative, super-population inference is that the estimands are free of the parameters governing the associations between potential outcomes. This is an advantage because the data are not informative about the associations between potential outcomes as all potential outcomes can never be jointly observed simultaneously for an experimental unit. Ding and Li (2018) suggested isolating the parameters that govern the marginal distributions from the parameters that govern the associations between potential outcomes, and performing sensitivity analyses for the association parameters. We omit the sensitivity analysis for these parameters and instead perform a sensitivity analysis for the prior distributions.

Table 6 compares the results obtained from our Bayesian methodology with those obtained by the method of Imai et al. (2021) in the case of MCAR outcomes. We do not consider the complier average spillover effects because we defined these estimands differently from Imai et al. (2021), and because Imai et al. (2021) mentioned that these estimands were imprecisely estimated in their analyses. We observe in Table 6 that our results are generally consistent with those of Imai et al. (2021). However, differences exist because our Bayesian method is able to detect significance effects in CADE(a_1, a_1), ITT_{Y,.,.}(a_1), and S_{Y,.,.}(1). The negative spillover effects that our method detects indicate that treated households are more likely to be negatively affected by the shift from a_0 to a_1 . Alternatively, assigning a greater proportion of households to treatment

	Post. Mean	Post. Median	95% Credible Interval
$O_{Y,\cdot,\cdot}$	-731	-739	(-1280, -181)
CAOE(a_0)	-1242	-1154	(-2334, -30)
CAOE(a_1)	-1234	-1328	(-2600, -75)
CADE(a_1, a_0)	-1647	-1626	(-3329, -230)
CADE(a_0, a_1)	259	237	(-1444, 2162)
CASE(0, a_0)	-33	-64	(-1691, 1797)
CASE(1, a_0)	-2004	-1870	(-3484, -548)
CASE(0, a_1)	128	-76	(-1903, 1981)
CASE(1, a_1)	-2172	-2160	(-3831, -694)

Table 7: Additional causal estimands and their estimates for the RSBY data.

will cause another treated household in the same village to spend less. We also observe large posterior standard deviations, and that the posterior intervals always have smaller widths than IJM's confidence intervals. The new effects that our methodology detects can be attributed to the greater precision (hence power) that follows from the use of the Bayesian model. Furthermore, our Bayesian methodology's ability to consider a point estimator based on the median, and a model that accommodates both an abundance of zeros and heavy tails, is advantageous for analyzing the data as the inferences would be robust to outlying observations. In particular, there are 36 observations greater than INR 100,000, with the two largest ones being INR 403,000 and 500,000. For comparison, the median in the dataset is INR 1,000.

Table 7 summarizes the results for the other causal estimands in the case of MCAR outcomes. Interestingly, the credible intervals of the overall effects for all units, as well as for compliers, lie below zero. The overall effect is of greatest interest to policy makers as it captures a pure impact of the intervention on all units. Our inferences on compliers imply that the overall effects are negative for units who comply with the assignment regardless of which treatment probability a_0 or a_1 is assigned to their respective cluster. We can also conclude from both Tables 6 and 7 that the direct effect of the treatment assignment under treatment probability a_1 is negative for compliers, regardless of their base cluster assignment. Finally, we can conclude that the spillover effects for compliers are negative when all units are assigned to treatment, regardless of whether they are compliers under a_0 or a_1 . Combined with our inferences on $S_{Y,\cdot,\cdot}(1)$, the spillover effect of treatment assignment on units is negative, regardless of their principal strata.

Table 8 presents the results in the case of MNAR outcomes. As the overall missingness rate is small (7.8%), we do not observe significant changes from the results under the MCAR mechanism in Table 6 and 7. This observation implies that assuming the MCAR mechanism would not hurt the final estimates for our analyses. Our Bayesian methodology enables us to explicitly assess the plausibility of missingness assumptions. This is another advantage of our methodology over the previous works (Imai et al., 2021; Forastiere et al., 2016; Vazquez-Bare, 2022).

	Post. Mean	Post. Median	95% Credible Interval
CADE(a_1, a_1)	-1951	-1871	(-3808, -361)
CADE(a_0, a_0)	127	78	(-1473, 1854)
ITT _{Y,·,·} (a_1)	-814	-783	(-1597, -149)
ITT _{Y,·,·} (a_0)	58	34	(-674, 838)
ITT _{D,·,·} (a_1)	0.417	0.417	(0.395, 0.436)
ITT _{D,·,·} (a_0)	0.456	0.456	(0.437, 0.474)
S _{Y,·,·} (1)	-1069	-1038	(-1726, -465)
S _{Y,·,·} (0)	-196	-209	(-1011, 700)
S _{D,·,·} (1)	0.036	0.036	(0.024, 0.051)
S _{D,·,·} (0)	0.075	0.075	(0.059, 0.092)
O _{Y,·,·}	-729	-714	(-1284, -217)
CAOE(a_0)	-1106	-1067	(-2291, -18)
CAOE(a_1)	-1222	-1213	(-2455, -37)
CADE(a_1, a_0)	-1773	-1710	(-3456, -328)
CADE(a_0, a_1)	143	151	(-1565, 1965)
CASE(0, a_0)	54	37	(-1595, 1939)
CASE(1, a_0)	-1846	-1750	(-3390, -451)
CASE(0, a_1)	50	26	(-1778, 2097)
CASE(1, a_1)	-2044	-1990	(-3657, -596)

Table 8: Results in the case of MNAR outcomes.

5.3 Sensitivity Analysis to Prior Specifications

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. The results in this section are derived using proper prior distributions. In particular, we use $\mu_{z,a,G_{i,j}} \sim \text{Normal}(0, \sigma_0^2)$, $\sigma_{z,a,G_{i,j}}^2 \sim \text{InverseGamma}(a_0, b_0)$, $(\pi_{(n,n)}, \pi_{(c,c)}, \pi_{(a,a)}, \pi_{(n,c)}, \pi_{(n,a)}, \pi_{(c,a)}) \sim \text{Dirichlet}(\phi_{(c,c)}, \phi_{(a,a)}, \phi_{(n,n)}, \phi_{(n,c)}, \phi_{(n,a)}, \phi_{(c,a)})$, and $p_{z,a,G_{i,j}} \sim \text{Beta}(\alpha_0, \beta_0)$ for the prior specifications. Table 10 presents the specifications of hyperparameters. Note that $\sigma_0^2 = 3$ is substantially different from $\sigma_0^2 = 5$ and $\sigma_0^2 = 7$ on a log scale. Table 9 reports the results for CADE. we observe that the extreme cases for p , specifically, Cases 10 and 11, lead to slight changes in the intervals, and that the other cases only exhibit minor changes in the results. However, the extreme values of p indicate strong prior beliefs about whether the outcome is zero or not for each principle strata. In our case study, as we lack knowledge about the unobservable strata, we believe that such strong priors about these parameters are unreasonable.

6 Concluding Remarks

We presented a Bayesian model-based methodology to simultaneously address interference, treatment nonadherence, and missing outcomes in two-stage randomized experiments. These three complications are significant drivers of unstable and biased causal

Percentile	CADE(a_0, a_0)			CADE(a_1, a_0)			CADE(a_0, a_1)			CADE(a_1, a_1)		
	2.5%	50%	97.5%	2.5%	50%	97.5%	2.5%	50%	97.5%	2.5%	50%	97.5%
Case 1	-1346	125	1773	-3329	-1593	-161	-1423	210	2034	-3691	-1770	-181
Case 2	-1356	158	1982	-3328	-1626	-229	-1444	236	2162	-3782	-1813	-256
Case 3	-1265	165	2134	-3431	-1679	-170	-1357	212	2084	-3845	-1873	-241
Case 4	-1313	192	1944	-3322	-1579	-179	-1399	262	2094	-3742	-1764	-205
Case 5	-1351	143	2013	-3458	-1682	-308	-1458	190	2120	-3905	-1881	-353
Case 6	-1388	141	1783	-3326	-1610	-192	-1463	233	2055	-3718	-1782	-218
Case 7	-1194	259	1884	-3797	-2034	-576	-1238	261	1938	-3888	-2081	-596
Case 8	-1323	168	2004	-3447	-1652	-288	-1442	230	2115	-3874	-1853	-326
Case 9	-1335	176	2052	-3341	-1639	-195	-1387	255	2147	-3759	-1813	-239
Case 10	-902	484	2128	-3354	-1734	-326	-990	525	2326	-3627	-1881	-355
Case 11	-1411	-10	1629	-3447	-1762	-447	-1521	21	1770	-3796	-1928	-490
Case 12	-1384	136	1857	-3298	-1660	-219	-1445	223	2131	-3632	-1813	-239

Table 9: Sensitivity analysis to prior specifications. We present the posterior medians and 95% central credible intervals of the causal estimands.

	Log-Normal			Principal Strata						Excess Zeros	
	σ_0^2	a_0	b_0	$\phi_{(c,c)}$	$\phi_{(a,a)}$	$\phi_{(n,n)}$	$\phi_{(n,c)}$	$\phi_{(n,a)}$	$\phi_{(c,a)}$	α_0	β_0
Case 1	3.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
Case 2	5.0	0.1	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
Case 3	7.0	0.01	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
Case 4	5.0	0.1	1.0	2.0	2.0	2.0	2.0	2.0	2.0	1.0	1.0
Case 5	5.0	0.1	1.0	1.0	1.0	1.0	1.0	1.0	1.0	2.0	2.0
Case 6	10.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
Case 7	3.0	0.1	0.1	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
Case 8	3.0	1.0	1.0	0.1	0.1	0.1	0.1	0.1	0.1	1.0	1.0
Case 9	3.0	1.0	1.0	10.0	10.0	10.0	10.0	10.0	10.0	1.0	1.0
Case 10	3.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.1	0.1
Case 11	3.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	10.0	10.0
Case 12	3.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.5	0.5

Table 10: Parameter specifications for sensitivity analysis.

inferences in modern experiments. Existing methodologies are lacking in their abilities to address all these issues simultaneously. There are three novel contributions of our Bayesian methodology. First, it provides a set of assumptions for performing valid causal inference in two-stage randomized experiments in the presence of these complications. We extended existing assumptions on the structure of interference to handle missing outcomes and also clarified and formalized assumptions about adherence behaviors within and across clusters. Second, it defines new causal estimands, including the overall effects of intervention and interpretable spillover effects, that can be inferred by means of our flexible Bayesian models. Our Bayesian methodologies address the issues of identifiability under two-sided nonadherence. Finally, our methodology can enable more precise causal inferences compared to existing methods for complex, non-standard data-generating mechanisms. We illustrated the utility of our methodology in this respect via simulation studies and the RSBY case study. In the latter study our methodology was able to uncover more definitive evidence of spillover and overall effects of the intervention that were not recognized in the previous analyses by Imai et al. (2021). Our

results are further validated by sensitivity analyses to prior distribution specifications and consideration of the case of MNAR missing outcomes.

Of great interest for future research on our Bayesian methodology is the relaxation of the assumptions of the interference structure. In certain real-life contexts it may be too restrictive to employ the two-stage randomized design and assume stratified interference. A great deal of research has been conducted on inferring causal effects without using special designs such as the two-stage randomized design or clustered encouragement design. The work of Aronow and Samii (2017) and Sävje et al. (2021) provide one possible path for future research based on the network structure and exposure mapping.

Supplementary Material

Supplementary Material for “A Bayesian Analysis of Two-Stage Randomized Experiments in the Presence of Interference, Treatment Nonadherence, and Missing Outcomes” (DOI: [10.1214/22-BA1347SUPP](https://doi.org/10.1214/22-BA1347SUPP); .pdf).

- Appendix A: Derivation of $\overline{D}_{i,j}(z, a)$ and $\overline{Y}_{i,j}(a)$. Proofs of $\overline{D}_{i,j}(z, a)$ and $\overline{Y}_{i,j}(a)$.
- Appendix B: On Exchangeability. A justification of exchangeability in the two-stage randomized experiment.
- Appendix C: Derivations of Gibbs Samplers. Detailed steps of Gibbs samplers in Section 3.4.
- Appendix D: Simulation Studies for Different Cluster Sizes. Results of simulation studies for different cluster sizes J .
- Appendix E: Simulation Studies under the Gamma data-generating Process. Results of simulation studies under the Gamma data-generating process.
- Appendix F: Parameters for Simulation Studies and Super-population Evaluation. Parameters for simulation studies in Section 4.1 and super-population evaluation.
- Appendix G: Parameters for Misspecification and MNAR evaluation. Parameters for Misspecification and MNAR evaluation in Section 4.3 and 4.4.
- Appendix H: Additional Simulation Studies under MNAR with various missing rates. Additional simulation studies under MNAR with various missing rates.

References

Angrist, J. D., Imbens, G. W., and Rubin, D. B. (1996). “Identification of Causal Effects Using Instrumental Variables.” *Journal of the American Statistical Association*, 91(434): 444–455. URL <http://www.jstor.org/stable/2291629> 206

- Aronow, P. M. (2012). “A General Method for Detecting Interference Between Units in Randomized Experiments.” *Sociological Methods and Research*, 40: 3–16. MR3190698. doi: <https://doi.org/10.1177/0049124112437535>. 210
- Aronow, P. M. and Samii, C. (2017). “Estimating average causal effects under general interference, with application to a social network experiment.” *Annals of Applied Statistics*, 11: 1912–1947. MR3743283. doi: <https://doi.org/10.1214/16-AOAS1005>. 210, 230
- Athey, S., Eckles, D., and Imbens, G. W. (2018). “Exact p-Values for Network Interference.” *Journal of the American Statistical Association*, 113: 230–240. MR3803460. doi: <https://doi.org/10.1080/01621459.2016.1241178>. 210
- Baird, S., Bohren, J. A., McIntosh, C., and Özler, B. (2018). “Optimal Design of Experiments in the Presence of Interference.” *The Review of Economics and Statistics*, 100(5): 844–860. URL https://doi.org/10.1162/rest_a_00716 210
- Basse, G. and Feller, A. (2018). “Analyzing Two-Stage Experiments in the Presence of Interference.” *Journal of the American Statistical Association*, 113: 41–55. MR3803438. doi: <https://doi.org/10.1080/01621459.2017.1323641>. 206
- Basse, G. W. and Airolidi, E. M. (2018a). “Limitations of Design-based Causal Inference and A/B Testing under Arbitrary and Network Interference.” *Sociological Methodology*, 48: 136–151. 210
- Basse, G. W. and Airolidi, E. M. (2018b). “Model-assisted design of experiments in the presence of network-correlated outcomes.” *Biometrika*, 105(4): 849–858. URL <https://doi.org/10.1093/biomet/asy036> 210
- Basse, G. W., Feller, A., and Toulis, P. (2019). “Randomization tests of causal effects under interference.” *Biometrika*, 106: 487–494. MR3949317. doi: <https://doi.org/10.1093/biomet/asy072>. 210
- Cox, D. and Reid, N. (2000). *The Theory of Design of Experiments*, 14. Chapman and Hall/CRC. 206
- Ding, P. and Li, F. (2018). “Causal inference: A missing data perspective.” *Statistical Science*, 33: 214–237. MR3797711. doi: <https://doi.org/10.1214/18-STS645>. 226
- Forastiere, L., Airolidi, E. M., and Mealli, F. (2021). “Identification and Estimation of Treatment and Interference Effects in Observational Studies on Networks.” *Journal of the American Statistical Association*, 116: 901–918. MR4270033. doi: <https://doi.org/10.1080/01621459.2020.1768100>. 210
- Forastiere, L., Mealli, F., and VanderWeele, T. J. (2016). “Identification and Estimation of Causal Mechanisms in Clustered Encouragement Designs: Disentangling Bed Nets Using Bayesian Principal Stratification.” *Journal of the American Statistical Association*, 111: 510–525. MR3538683. doi: <https://doi.org/10.1080/01621459.2015.1125788>. 207, 210, 211, 212, 227
- Frangakis, C. and Rubin, D. (1999). “Addressing complications of intention-to-treat analysis in the combined presence of all-or-none treatment-noncompliance and subse-

- quent missing outcomes." *Biometrika*, 86(2): 365–379. URL <https://doi.org/10.1093/biomet/86.2.365> MR1705410. doi: <https://doi.org/10.1093/biomet/86.2.365>. 207, 217
- Frangakis, C. E. and Rubin, D. B. (2002). "Principal Stratification in Causal Inference." *Biometrics*, 58: 21–29. MR1891039. doi: <https://doi.org/10.1111/j.0006-341X.2002.00021.x>. 207
- Frumento, P., Mealli, F., Pacini, B., and Rubin, D. B. (2012). "Evaluating the Effect of Training on Wages in the Presence of Noncompliance, Nonemployment, and Missing Outcome Data." *Journal of the American Statistical Association*, 107(498): 450–466. URL <http://www.jstor.org/stable/23239583> MR2980057. doi: <https://doi.org/10.1080/01621459.2011.643719>. 207
- Gallop, R., Small, D. S., Lin, J. Y., Elliott, M. R., Joffe, M., and Have, T. R. T. (2009). "Mediation analysis with principal stratification." *Statistics in Medicine*, 28: 1108–1130. MR2662200. doi: <https://doi.org/10.1002/sim.3533>. 207
- Gelfand, A. E. and Smith, A. F. M. (1990). "Sampling-Based Approaches to Calculating Marginal Densities." *Journal of the American Statistical Association*, 85(410): 398–409. URL <http://www.jstor.org/stable/2289776> MR1141740. 220
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., and Rubin, D. B. (2013). *Bayesian Data Analysis*. Chapman and Hall/CRC, 3rd edition. MR3235677. 218
- Geman, S. and Geman, D. (1984). "Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images." *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-6(6): 721–741. 220
- Holland, P. W. (1986). "Statistics and Causal Inference." *Journal of the American Statistical Association*, 81(396): 945–960. URL <http://www.jstor.org/stable/2289064> MR0867618. 208
- Hudgens, M. G. and Halloran, M. E. (2008). "Toward causal inference with interference." *Journal of the American Statistical Association*, 103: 832–842. MR2435472. doi: <https://doi.org/10.1198/016214508000000292>. 206, 210, 213
- Imai, K., Jiang, Z., and Li, M. (2019). *experiment: R Package for Designing and Analyzing Randomized Experiments*. R package version 1.2.0. URL <https://CRAN.R-project.org/package=experiment> 221
- Imai, K., Jiang, Z., and Malani, A. (2021). "Causal Inference With Interference and Noncompliance in Two-Stage Randomized Experiments." *Journal of the American Statistical Association*, 116: 632–644. MR4270009. doi: <https://doi.org/10.1080/01621459.2020.1775612>. 207, 208, 210, 211, 212, 213, 214, 215, 221, 222, 223, 225, 226, 227, 229
- Imbens, G. W. and Rubin, D. B. (1997). "Bayesian Inference for Causal Effects in Randomized Experiments with Noncompliance." URL <http://www.jstor.org/stable/2242722> MR1429927. doi: <https://doi.org/10.1214/aos/1034276631>. 207, 208, 220

- Imbens, G. W. and Rubin, D. B. (2015). *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*, 3–5, 40. Cambridge University Press. 205, 218
- Kang, H. and Imbens, G. (2016). “Peer Encouragement Designs in Causal Inference with Partial Interference and Identification of Local Average Network Effects.” <https://arxiv.org/abs/1609.04464>. 207
- Leung, M. P. (2020). “Treatment and spillover effects under network interference.” *Review of Economics and Statistics*, 102: 368–380. 210
- Little, R. and Rubin, D. (2002). *Statistical analysis with missing data*. Wiley series in probability and mathematical statistics. Probability and mathematical statistics. Wiley. URL <http://books.google.com/books?id=aYPwAAAAMAAJ> MR0440638. 206, 217
- Little, R. J., D’Agostino, R., Cohen, M. L., Dickersin, K., Emerson, S. S., Farrar, J. T., Frangakis, C., Hogan, J. W., Molenberghs, G., Murphy, S. A., Neaton, J. D., Rotnitzky, A., Scharfstein, D., Shih, W. J., Siegel, J. P., and Stern, H. (2012). “The Prevention and Treatment of Missing Data in Clinical Trials.” *New England Journal of Medicine*, 367(14): 1355–1360. PMID: 23034025. URL <https://doi.org/10.1056/NEJMs1203730> 206
- Liu, L. and Hudgens, M. G. (2014). “Large sample randomization inference of causal effects in the presence of interference.” *Journal of the American Statistical Association*, 109: 288–301. MR3180564. doi: <https://doi.org/10.1080/01621459.2013.844698>. 206
- Manski, C. F. (2013). “Identification of treatment response with social interactions.” *Econometrics Journal*, 16. MR3030060. doi: <https://doi.org/10.1111/j.1368-423X.2012.00368.x>. 210
- Mattei, A. and Mealli, F. (2007). “Application of the Principal Stratification Approach to the Faenza Randomized Experiment on Breast Self-Examination.” *Biometrics*, 63(2): 437–446. URL <http://www.jstor.org/stable/4541355> MR2370802. doi: <https://doi.org/10.1111/j.1541-0420.2006.00684.x>. 207
- Mattei, A., Mealli, F., and Pacini, B. (2014). “Identification of causal effects in the presence of nonignorable missing outcome values.” *Biometrics*, 70(2): 278–288. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/biom.12136> MR3258033. doi: <https://doi.org/10.1111/biom.12136>. 207
- Nandi, A., Holtzman, E. P., Malani, A., and Laxminarayan, R. (2015). “The need for better evidence to evaluate the health & economic benefits of India’s Rashtriya Swasthya Bima Yojana.” *The Indian journal of medical research*, 142(4): 383–390. 26609029[pmid]. URL <https://pubmed.ncbi.nlm.nih.gov/26609029> 208, 225
- Ohnishi, Y., and Sabbaghi, A. (2022). “Supplementary Material for “A Bayesian Analysis of Two-Stage Randomized Experiments in the Presence of Interference, Treatment Nonadherence, and Missing Outcomes”.” *Bayesian Analysis*. doi: <https://doi.org/10.1214/22-BA1347SUPP>. 214, 221, 223

- Pearl, J. (2009). “Causal Inference in Statistics: An Overview.” *Statistics Surveys*, 3: 96–146. MR2545291. doi: <https://doi.org/10.1214/09-SS057>. 205
- Rubin, D. B. (1974). “Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies.” *Journal of Educational Psychology*, 66: 688–701. 208
- Rubin, D. B. (2008). “Comment: The Design and Analysis of Gold Standard Randomized Experiments.” *Journal of the American Statistical Association*, 103(484): 1350–1353. URL <https://doi.org/10.1198/016214508000001011> MR2655717. doi: <https://doi.org/10.1198/016214508000001011>. 205
- Sävje, F., Aronow, P. M., and Hudgens, M. G. (2021). “Average treatment effects in the presence of unknown interference.” *Annals of Statistics*, 49: 673–701. MR4255103. doi: <https://doi.org/10.1214/20-aos1973>. 210, 230
- Tchetgen, E. J. and Vanderweele, T. J. (2012). “On causal inference in the presence of interference.” *Statistical Methods in Medical Research*, 21: 55–75. MR2867538. doi: <https://doi.org/10.1177/0962280210386779>. 206
- VanderWeele, T. J. (2011). “Principal stratification – Uses and limitations.” *International Journal of Biostatistics*, 7. MR2843532. doi: <https://doi.org/10.2202/1557-4679.1329>. 207
- VanderWeele, T. J. and Tchetgen, E. J. T. (2011). “Effect partitioning under interference in two-stage randomized vaccine trials.” *Statistics and Probability Letters*, 81: 861–869. MR2793754. doi: <https://doi.org/10.1016/j.spl.2011.02.019>. 206, 209, 214
- Vazquez-Bare, G. (2022). “Causal Spillover Effects Using Instrumental Variables.” *Journal of the American Statistical Association*, 0(0): 1–12. URL <https://doi.org/10.1080/01621459.2021.2021920> 207, 212, 220, 227
- Zhang, J. L. and Rubin, D. B. (2003). “Estimation of Causal Effects via Principal Stratification When Some Outcomes are Truncated by “Death”.” *Journal of Educational and Behavioral Statistics*, 28(4): 353–368. URL <https://doi.org/10.3102/10769986028004353> 207