Robust Inference for the Direct Average Treatment Effect with Treatment Assignment Interference

Matias D. Cattaneo* Yihan He† Ruiqi (Rae) Yu‡

cattaneo@princeton.edu yihan.he@princeton.edu rae.yu@princeton.edu

February 18, 2025

Abstract

Uncertainty quantification in causal inference settings with random network interference is a challenging open problem. We study the large sample distributional properties of the classical difference-in-means Hajek treatment effect estimator, and propose a robust inference procedure for the (conditional) direct average treatment effect, allowing for cross-unit interference in both the outcome and treatment equations. Leveraging ideas from statistical physics, we introduce a novel Ising model capturing interference in the treatment assignment, and then obtain three main results. First, we establish a Berry-Esseen distributional approximation pointwise in the degree of interference generated by the Ising model. Our distributional approximation recovers known results in the literature under no-interference in treatment assignment, and also highlights a fundamental fragility of inference procedures developed using such a pointwise approximation. Second, we establish a uniform distributional approximation for the Hajek estimator, and develop robust inference procedures that remain valid regardless of the unknown degree of interference in the Ising model. Third, we propose a novel resampling method for implementation of robust inference procedure. A key technical innovation underlying our work is a new *De-Finetti Machine* that facilitates conditional i.i.d. Gaussianization, a technique that may be of independent interest in other settings.

Keywords: causal inference under interferece; Ising Model; Distribution Theory; Robust Inference

*Department of Operations Research and Financial Engineering, Princeton University.

†Department of Operations Research and Financial Engineering, Princeton University.

‡Department of Operations Research and Financial Engineering, Princeton University.

1 Introduction

We study the large sample distributional properties of the classical Hajek average treatment effect estimator, and propose a robust inference procedure for the (conditional) direct average treatment effect, in the presence of cross-unit interference in both the outcome and treatment equations. This causal inference problem arises in a variety of contexts such as (online) social networks, medical trials, and socio-spatial studies, and has received renewed attention in recent years. Recent contributions include [1], [11], [10], [12], [18], [20], and references therein. See [8] for a modern textbook introduction to causal inference.

The key challenge in causal inference settings with interference is that units can affect each other in arbitrary ways, making statistical inference difficult without disciplining the degree of cross-unit interference: it is common to assume that units correspond to vertices in a network, typically represented as a graph, such that only when units are connected by an edge, they may influence each other. Early literature assumed that the underlying network was fixed, or otherwise known, but more recent advances have considered estimation and inference methods allowing the network to be a random (unobserved) graph (see Assumption 1 below). Furthermore, due to the challenges introduced by the presence of the latent random graph structure, it is common in the literature to restrict the degree of interference entering the outcome and treatment equations: prior work has focused on the special case where the potential outcomes exhibit restricted interference in the form of *annonymity* or *exchangability* (see Assumption 2 below), but the treatment assigment mechanism does not exhibit interference. We contribute to this emerging causal inference literature by allowing for the treatment assignment mechanism to also exhibit restricted cross-unit interference, while retaining the other semiparametric modelling assumptions imposed in previous work.

Leveraging ideas from statistical physics [7], we introduce a class of Ising equiprobable treatment assignment mechanisms described by

$$\mathbb{P}_\beta(\mathbf{T} = \mathbf{t}) \propto \exp\Big(\frac{\beta}{n}\sum_{i\neq j}(2t_i - 1)(2t_j - 1)\Big), \tag{1}$$

where $\mathbf{T} = (T_1, \dots, T_n)^\top \in \{0,1\}^n$ denote the vector of binary treatment assignments for n units, $\mathbf{t} = (t_1 \dots, t_n)^\top$, and the unknown parameter $\beta \geq 0$ controls the degree of cross-unit interference in their treatment assignments (see Assumption 3 below). This model explicitly accounts for the stochastic nature of network formation in the treatment equation, and reduces to the classical independent equiprobable treatment assignment rule when $\beta = 0$ (i.e., random assignment with equal probability). Thus, the Ising equiprobable treatment assignment model allow us to investigate how prior conclusions in the literature change as a function of the degree of cross-unit interference in treatment assignment as controlled by the unknown parameter β.

To streamline the presentation, and due to some technical issues, we focus on the moderate cross-unit interference regime $\beta \in [0, 1]$. See Section 8 for more discussion. Our first contribution concerns the large sample distributional properties of the classical difference-in-means Hajek estimator (see (3) below). Theorem 3.1 establishes a Berry-Esseen bound for the estimator, that is, a distributional approximation in Kolgomorov distance with explicit convergence rates. The closest antecedent is [12], who considered the same causal model with network interference but under the assumption $\beta = 0$ (random treatment assignment), and established a Gaussian distributional approximation for the Hajek estimator. Theorem 3.1 establishes a precise distributional approximation with explicit convergence rates and, more importantly, shows that: (i) for $\beta \in [0, 1)$, the limiting distribution continues to be Gaussian, but the asymptotic variance exhibits an additional term that captures the cross-unit interference in the treatment equation; (ii) for $\beta \in [0, 1)$, the new asymptotic variance coincides with the one obtained in [12] when $\beta = 0$ but is increasing and unbounded as a function of β; and (iii) for $\beta = 1$, the limiting distribution is non-Gaussian. These findings have an important implication for the robustness of inference procedures developed under the assumption of no-interference in the treatment assignment ($\beta = 0$): the distribution approximation changes as a function of $\beta \geq 0$, exhibiting a discontinuity at $\beta = 1$, thereby invalidating inference procedures obtained from distributional approximations that only hold pointwise in β.

The lack of uniform validity demonstrated in Theorem 3.1 poses a major challenge for developing robust inference procedures in the presence of potential interference in the treatment assignment because β is unknown in practice. Moreover, [16] showed that no consistent estimator exists for $\beta \in [0, 1)$, making plug-in inference procedures infeasible, even pointwise in $\beta \in [0, 1)$. To address these challenges, Theorem 4.1 establishes a uniform in $\beta \in [0, 1]$ distributional approximation for the Hajek estimator and, as a necessary by-product, also establishes a uniform distributional approximation in $\beta \in [0, 1]$ for its Maximum Pseudo-Likelihood estimator (MPLE); see [17]. The resulting distributional approximations are indexed by a localization parameter offering a smooth transition between the discontinuous limit laws established in Theorem 3.1, as well as for those corresponding to the MPLE of β.

Building on Theorem 4.1, and employing a Bonferroni-correction procedure that works by creating hierarchical confidence intervals for different β-regimes, we present uniformly valid uncertainty quantification for the (conditional) direct average treatment effect τ_n (see : we develop infeasible (Theorem 4.2) and feasible (Theorem 5.1) prediction intervals $\mathcal{C}_n(\alpha)$ satisfying

$$\liminf_{n\to\infty} \inf_{\beta\in[0,1]} \mathbb{P}_\beta[\tau_n \in \mathcal{C}_n(\alpha)] \geq 1 - \alpha,$$

for $\alpha \in [0, 1]$, where $\mathcal{C}_n(\alpha)$ is based on the Hajek estimator and a novel resampling procedure aimed to capturing sampling uncertainty coming from the underlying network. To the best of our knowledge, our proposed feasible inference procedure is new for $\beta = 0$. More importantly, our proposed inference procedure is the first to offer robust (uniform) validity across all values of $\beta \in [0, 1]$. Section 6 presents a simulation study demonstrating the performance of our proposed methods.

1.1 Summary of Methodological and Technical Contributions

From a methodological perspective, our paper contributes to the literature on causal inference under cross-unit interference. Classical contributions include [9], [19], [15], and references therein. The closest antecedent to our work is [12], who studied distribution theory for the same casual inference model with network interference considered in this paper except for assuming random treatment assignment (i.e., without interference in the treatment assignment mechanism). Thus, our first methodological contribution is to propose a novel Ising equiprobable treatment assignment model to capture the possible interdependency between treatment assignments when units can interference with each other. The model covers the equiprobable experimental design, as well as a class of dependent treatment assignments as indexed by β in (1). Furthermore, our second main methodological contribution is to present a novel, feasible robust inference procedure for the (conditional) direct average treatment effect, which is uniformly valid for all $\beta \in [0, 1]$. This procedure relies on a Bonferroni correction together with a uniform distributional approximation for the Hajek estimator, taking into account the different β-regimes, and also leverages a new resampling-based variance estimator developed herein. Our proposed inference procedure appears to be new even in the special case of $\beta = 0$ (no-interference in treatment assignment).

From a technical perspective, our paper also offers a contribution to the applied probability literature, particularly in the context of statistical mechanics [7]. Allowing for interference in treatment assignment leads to major technical challenges for establishing distribution theory for the Hajek estimator, since the Ising equiprobable treatment model introduces new sources of dependence that need to be taken into account. For example, as shown in Theorem 3.1, the Hajek estimator exhibits different concentration rates around τ_n depending on whether $\beta = 1$ or not, in addition to having different limit laws. Our first technical contribution is to develop a new *De-Finetti Machine* that leverages the exchangeability structure in the treatment vector induced by Ising model, which we then use to establish a Berry-Esseen bound under the different β-regimes. This new technique is based on a carefully crafted conditioning argument that renders the elements of $\mathbf{T}$ conditionally i.i.d., thereby reducing the problem to establishing a Berry-Esseen bound for conditionally i.i.d. random variables. Our new technical approached generalizes [5] and [6] by considering a multiplier Curie-Weiss magnetization statistic, without relying on variants of Stein's method [5], and instead using a novel conditional i.i.d. Gaussianization approach. Our new technique may be of independent interest in other settings considering establishing a Berry-Esseen bound for sum of exchangeable random variables. To address the uniform inference problem, we further establish uniform in $\beta \in [0, 1]$ distributional approximations: our results cover both the Hajek estimator and the MPLE for β. Thus, a second technical contribution of our work is to the literature on distributional properties of the Ising model.

1.2 Organization

Section 2 formalizes the setup. Section 3 presents pointwise in $\beta \in [0, 1]$ distribution theory for the Hajek estimator. Section 4 presents uniform in $\beta \in [0, 1]$ distribution theory, and discusses an infeasible uniformly valid inference procedure. Section 5 proposes a feasible inference procedure based on resampling methods. Section 6 presents simulation evidence. Section 7 overviews our technical contributions, including Berry-Esseen bounds for Curie-Weiss magnetization with independent multipliers, and Section 8 concludes with open questions and future research directions.

2 Setup

We consider a random potential outcome framework under network interference. For each unit $i \in [n] = \{1, 2, \cdots, n\}$, let $Y_i(t; \mathbf{t}_{-i})$ denote its random potential outcome when assigned to treatment level $t \in \{0, 1\}$ while the other units are assigned to treatment levels $\mathbf{t}_{-i} \in \{0, 1\}^{n-1}$. The vector of observed random treatment assignments for the n units is $\mathbf{T} = (T_i : i \in [n])$, and $\mathbf{T}_{-i}$ denotes the associated random treatment assignment vector excluding T_i. Thus, the observed data is $(Y_i, T_i : i \in [n])$ with $Y_i = (1 - T_i)Y_i(0; \mathbf{T}_{-i}) + T_i Y_i(1; \mathbf{T}_{-i})$ for each $i \in [n]$.

Interference among the n units is modelled via a latent network characterized by an undirected random graph $G(\mathbf{V}, \mathbf{E})$ with vertex set $\mathbf{V} = [n]$ and (random) adjacency matrix $\mathbf{E} = (E_{ij} : (i, j) \in [n] \times [n]) \in \{0, 1\}^{n \times n}$. The following assumption restricts this random graph structure.

Assumption 1 (Network Structure)**.** *The random network* $\mathbf{E}$ *satisfies: For all* $1 \leq i \leq j \leq n$ *and* $\rho_n \in (0, 1]$*,* $E_{ii} = 0$*,* $E_{ij} = E_{ji}$*, and* $E_{ij} = \mathbb{1}(\xi_{ij} \leq \min\{1, \rho_n G(U_i, U_j)\})$*, where* $G : [0, 1]^2 \mapsto \mathbb{R}_+$ *is symmetric, continuous and positive on* $[0, 1]^2$*,* $\mathbf{U} = (U_i : i \in [n])$ *are i.i.d.* $\mathsf{Uniform}[0, 1]$ *random variables,* $\mathbf{\Xi} = (\xi_{ij} : (i, j) \in [n] \times [n], i < j)$ *are i.i.d* $\mathsf{Uniform}[0, 1]$ *random variables. Finally,* $\mathbf{U}$ *and* $\mathbf{\Xi}$ *are independent.*

This assumption corresponds to the *sparse graphon model* of [3]. The parameter ρ_n controls the sparsity of the network, and will play an important role in our theoretical results. The variable U_i is a *latent* heterogenous property of the ith unit, and $G(U_i, U_j)$ measures similarity between traits of U_i and U_j. This allows for a stochastic model for the edge formation.

Building on the underlying random graph structure, the following assumption imposes discipline on the interference entering the outcome equation.

Assumption 2 (Exchangable Smooth Potential Outcomes Model)**.** *For all* $i \in [n]$*,* $Y_i(T_i; \mathbf{T}_{-i}) = f_i(T_i; \frac{M_i}{N_i})$ *where* $M_i = \sum_{j \neq i} E_{ij} T_j$*,* $N_i = \sum_{j \neq i} T_j$*, and* $\mathbf{f} = (f_i : i \in [n])$ *are i.i.d random functions. In addition, for all* $i \in [n]$ *and some integer* $p \geq 4$*,* $\max_{1 \leq i \leq n} \max_{t \in \{0,1\}} |\partial_2^{(p)} f_i(t, \cdot)| < C$ *for some* C *not depending on* n *and* β*. Finally,* $\mathbf{f}$ *is independent to* $\mathbf{\Xi}$*.*

This second assumption imposes two main restrictions on the potential outcomes. First, a dimension reduction is assumed via the underlying network structure (Assumption 1), making the potential outcomes for each unit $i \in [n]$ a function of only their own treatment assignment and the fraction of other treated units among their (connected) peers. Second, the potential outcomes are assumed to be smooth as a function of the fraction of treated peers, thereby ruling out certain types of outcome variables (e.g., binary or similarly limited dependent variable models). Assumption 2 explicitly parametrizes the smoothness level p because, together with the the sparsity parameter ρ_n in Assumption 1, it will play an important role in our theoretical results.

To close the causal inference model, the following assumption restricts the treatment assignment distribution. We propose an Ising model from statistical physics [7].

Assumption 3 (Ising Equiprobable Treatment Assignment)**.** *The treatment assignment mechanism follows a Curie-Weiss distribution:*

$$\mathbb{P}_\beta(\mathbf{T} = \mathbf{t}) = \frac{1}{C_\beta} \exp\left(\frac{\beta}{n}\sum_{i\neq j}(2t_i - 1)(2t_j - 1)\right), \tag{2}$$

where $\mathbf{t} \in \{0,1\}^n$*,* $\beta \in [0,1]$*, and* C_β *is determined by the condition* $\sum_{\mathbf{t}} \mathbb{P}_\beta(\mathbf{T} = \mathbf{t}) = 1$*.*

This model naturally encodes a class of equiprobable, possibly dependent treatment assignment mechanisms. Assumption 3 implies $\mathbb{P}_\beta(T_i = 1) = \frac{1}{2}$ for $i \in [n]$ and all $\beta \geq 0$, but allows for correlation in treatment assignment as controlled by β. When $\beta = 0$, treatment assignment becomes independent across units, and thus the assignment mechanism reduces to the canonical (equiprobable) randomized allocation. For $\beta \in [0, 1]$, the Ising mechanism induces positive pairwise correlations, capturing social interdependence phenomena like peer influence [14] characteristic of observational settings.

We propose a robust inference procedure based on the popular Hajek estimator

$$\widehat{\tau}_n = \frac{\sum_{i=1}^n T_i Y_i}{\sum_{i=1}^n T_i} - \frac{\sum_{i=1}^n (1 - T_i) Y_i}{\sum_{i=1}^n (1 - T_i)}. \tag{3}$$

This classical estimator is commonly used in causal inference, both with and without interference. In particular, [12] studied the asymptotic properties of $\widehat{\tau}_n$ when $\beta = 0$, under Assumptions 1–3, and showed that

$$\sqrt{n}(\widehat{\tau}_n - \tau_n) \rightsquigarrow \mathsf{N}(0, \kappa_2), \qquad \kappa_s = \mathbb{E}[(R_i - \mathbb{E}[R_i] + Q_i)^s], \tag{4}$$

where $\rightsquigarrow$ denotes weak convergence as $n \to \infty$, the standard target is the (conditional) direct average treatment effect given by

$$\tau_n = \frac{1}{n}\sum_{i=1}^n \mathbb{E}[Y_i(1; \mathbf{T}_{-i}) - Y_i(0; \mathbf{T}_{-i}) | f_i(\cdot), \mathbf{E}], \tag{5}$$

and $R_i = f_i(1, \frac{1}{2}) + f_i(0, \frac{1}{2})$ and $Q_i = \mathbb{E}[\frac{G(U_i,U_j)}{\mathbb{E}[G(U_i,U_j)|U_j]}(f_j'(1, \frac{1}{2}) - f_j'(0, \frac{1}{2}))|U_i]$. The (conditional) direct average treatment effect in (5) is a *predictand*, not an *estimand*, in the sense that it is a random variable that needs not to settle to a non-random probability limit under the assumptions imposed. Consequently, our uncertainty quantification methods can be regarded as prediction intervals for the classical target predictand τ_n in the causal inference literature.

3 Pointwise Distribution Theory

Our first main result is a Berry-Esseen bound for the Hajek estimator, pointwise in $\beta \in [0, 1]$, that is, the degree of treatment assignment interference. We provide a proof sketch in Section 7, with full technical details deferred to the supplementary material.

Theorem 3.1 (Pointwise Distribution Theory). *Suppose Assumptions 1, 2, and 3 hold. Then,*

$$\sup_{t\in\mathbb{R}} |\mathbb{P}[\widehat{\tau}_n - \tau_n \le t] - L_n(t;\beta,\kappa_1,\kappa_2)| = O\Big(\frac{\log n}{\sqrt{n\rho_n}} + \mathtt{r}_{n,\beta}\Big),$$

where $L_n(\cdot;\beta,\kappa_1,\kappa_2)$ and $\mathtt{r}_{n,\beta}$ are as follows. Then: (1) High temperature: if $\beta \in [0,1)$,

$$L_n(t;\beta,\kappa_1,\kappa_2) = \mathbb{P}_\beta\Big[n^{-1/2}\Big(\kappa_2 + \kappa_1^2\frac{\beta}{1-\beta}\Big)^{1/2}\mathsf{Z} \le t\Big] \tag{6}$$

with $\mathsf{Z} \sim \mathsf{N}(0,1)$, and $\mathtt{r}_{n,\beta} = \sqrt{n\log n}(n\rho_n)^{-\frac{p+1}{2}}$.
(2) Critical temperature: if $\beta = 1$,

$$L_n(t;\beta,\kappa_1,\kappa_2) = \mathbb{P}_\beta[n^{-1/4}\kappa_1\mathsf{W}_0 \le t] \tag{7}$$

with $\mathbb{P}[\mathsf{W}_0 \le w] = \frac{\int_{-\infty}^{w}\exp(-z^4/12)dz}{\int_{-\infty}^{\infty}\exp(-z^4/12)dz}$, $w \in \mathbb{R}$, and $\mathtt{r}_{n,\beta} = (\log n)^3 n^{-\frac{1}{4}} + \sqrt[4]{n}\sqrt{\log n}(n\rho_n)^{-\frac{p+1}{2}}$.

In the high temperature regime ($\beta \in [0,1)$), $\sqrt{n}(\widehat{\tau}_n - \tau_n)$ is asymptotically normal with variance $\kappa_2 + \kappa_1^2\frac{\beta}{1-\beta}$. Thus, when $\beta = 0$, our result recovers (4), but for $\beta \in (0,1)$, the asymptotic variance is strictly increasing unless $\kappa_1 = 0$ (i.e., no randomness from the underlying network). In the Critical temperature regime ($\beta = 1$), the limiting distribution is non-Gaussian. The distinct asymptotic behaviors of $\widehat{\tau}_n$ across these regimes mirror the phase transition phenomena observed in the Ising model's magnetization $m = \frac{1}{n}\sum_{i=1}^n(2T_i - 1)$. The first term in the Berry-Esseen bound, $\log n(n\rho_n)^{-1/2}$, is not improvable beyond the extra logarithmic factor because $\log(n)n^{-1/2}$ when $\rho_n \asymp 1$. For the second term, $\mathtt{r}_{n,\beta}$, the bound depends on the smoothness p of the potential outcome function and the temperature regime.

Theorem 3.1 highlights key challenges in uncertainty quantification, with unknown quantities κ_1 and κ_2, and the unknown regime parameter $\beta \in [0,1]$. Furthermore, [2] established an impossibility result showing that no consistent estimator for β exists in the high-temperature regime. In the following section, we address the estimation of β and the complications arising from the discontinuous transition between Gaussian and non-Gaussian laws.

4 Infeasible Robust Inference

This section addresses inference on the treatment effect when the regime parameter β is unknown, but assuming that κ_1 and κ_2 are known.

4.1 Maximum Pseudo-Likelihood Estimator (MPLE) for Temperature

Due to the existence of the normalizing constant C_β in (2), maximum likelihood estimation is not computationally efficient. However, the conditional distribution of T_i given the rest of treatments adopts a closed form solution and can be optimized efficiently [17]. Define $W_i = 2T_i - 1$, $\mathbf{W}_{-i} = \{W_j : j \in [n], j \neq i\}$, and $m_i = \frac{1}{n}\sum_{j\neq i} W_j$. The MPLE for β is

$$\widehat{\beta}_n = \arg\max_{\beta\in[0,1]} \sum_{i\in[n]} \log \mathbb{P}_\beta[W_i|\mathbf{W}_{-i}] = \arg\max_{\beta\in[0,1]} \sum_{i\in[n]} -\log\left(\frac{1}{2}W_i \tanh(\beta m_i) + \frac{1}{2}\right).$$

We show in Lemma 6 in the supplementary appendix that the limiting distribution of $\widehat{\beta}_n$ also depends on the regime $\beta \in [0,1]$. For $\beta \in [0,1)$, $1 - \widehat{\beta}_n \rightsquigarrow (1-\beta)\max\{(\chi_1^2)^{-1}, 0\}$, thereby ruling out consistent estimation. For $\beta = 1$, $\sqrt{n}(\widehat{\beta}_n - 1) \rightsquigarrow \min\{\mathsf{W}_0^2/3 - 1/\mathsf{W}_0^2, 1\}$, where W_0 is given in Theorem 3.1. For fixed n, the distribution of $\widehat{\beta}_n - 1$ exhibits the same discontinuity at $\beta = 1$ as $\widehat{\tau}_n - \tau_n$, highlighting the need for a distributional approximation that is uniform in β for valid inference across all regimes.

4.2 Robust Distribution Theory

We develop valid large sample inference for all values of $\beta \in [0,1]$. From Theorem 3.1, for all $\beta \in [0,1)$, the limiting variance of $\sqrt{n}(\widehat{\tau}_n - \tau_n)$ is $\kappa_2 + \kappa_1^2\frac{\beta}{1-\beta}$. Thus, when $\kappa_1 \neq 0$, the asymptotic variance diverges as β approaches the critical value $\beta = 1$. In contrast, Theorem 3.1 shows that when $\beta = 1$ the limiting variance of $n^{1/4}(\widehat{\tau}_n - \tau_n)$ is finite. This discrepancy indicates a lack of uniform validity in the distributional approximations in Theorem 3.1. To address this issue, we establish a uniform distributional approximation based on the drifting sequence $\beta_n = 1 + \frac{c}{\sqrt{n}}$. This sequence follows the *knife-edge* rate, ensuring that the law of $\widehat{\tau}_n - \tau_n$ smoothly interpolates between the pointwise distributional approximations indexed by $\beta \in [0,1]$.

Theorem 4.1 (Robust Distribution Theory)**.** *Suppose Assumptions 1, 2 and 3 hold. Define* $c_{\beta,n} = \sqrt{n}(1-\beta)$. *Then,*

$$\lim_{n\to\infty} \sup_{0\leq\beta\leq 1} \sup_{t\in\mathbb{R}} \left| \mathbb{P}_\beta[\widehat{\tau}_n - \tau_n \leq t] - \mathbb{P}_\beta[n^{-\frac{1}{2}}\kappa_2^{\frac{1}{2}}\mathsf{Z} + \beta^{\frac{1}{2}} n^{-\frac{1}{4}}\kappa_1 \mathsf{W}_{c_{\beta,n}} \leq t] \right| = 0$$

with $\mathsf{Z} \sim \mathsf{N}(0,1)$ *independent of* W_c, *and* $\mathbb{P}[\mathsf{W}_c \leq w] = \frac{\int_{-\infty}^{w} \exp(-\frac{x^4}{12} - \frac{cx^2}{2})dx}{\int_{-\infty}^{\infty} \exp(-\frac{x^4}{12} - \frac{cx^2}{2})dx}$, $w \in \mathbb{R}$. *Furthermore,*

$$\lim_{n\to\infty} \sup_{0\leq\beta\leq 1} \sup_{t\in\mathbb{R}} \left| \mathbb{P}_\beta[1 - \widehat{\beta}_n \leq t] - \mathbb{P}_\beta[\min\{\max\{\mathsf{T}_{c_{\beta,n},n}^{-2} - \mathsf{T}_{c_{\beta,n},n}^2/(3n), 0\}, 1\} \leq t] \right| = 0$$

where $\mathsf{T}_{c,n} = \mathsf{Z} + n^{\frac{1}{4}}\mathsf{W}_c$.

Theorem 4.1 establishes that $H_n(t;\kappa_1,\kappa_2,c_{\beta,n}) = \mathbb{P}_\beta[n^{-\frac{1}{2}}\kappa_2^{\frac{1}{2}}\mathsf{Z} + \beta^{\frac{1}{2}}n^{-\frac{1}{4}}\kappa_1\mathsf{W}_{c_{\beta,n}} \leq t]$ uniformly approximates the distribution of $\widehat{\tau}_n - \tau_n$ in both the high-temperature and critical-temperature regimes. Under the *knife-edge* scaling, the leading term $n^{-1/2}\kappa_2^{1/2}\mathsf{Z}$ becomes negligible, and the typical *knife-edge* representation retains only the second term $\beta^{1/2}n^{-1/4}\kappa_1\mathsf{W}_c$. However, when $\beta \in [0,1)$ is fixed and $c_{\beta,n} = \sqrt{n}(1-\beta) \to \infty$, $\mathsf{W}_{c_{\beta,n}}$ approximates $n^{-1/4}\mathsf{N}(0,(1-\beta)^{-1})$, making both terms comparable in order. Consequently, we retain both terms in the distributional approximation. In Lemma 4 in the supplementary , we show that when β is fixed and $c_{\beta,n} = \sqrt{n}(1-\beta)$, we have $\sup_{t\in\mathbb{R}}|H_n(t;\kappa_1,\kappa_2,c_{\beta,n}) - L_n(t;\kappa_1,\kappa_2,\beta)| \to 0$. The same ideas apply to the uniform approximation of $1 - \widehat{\beta}_n$.

4.3 Infeasible Uniform Inference

We can now propose a conservative prediction interval based on the following Bonferroni-correction procedure. In particular, in the first step, a uniform confidence interval for β is constructed under the *knife-edge* approximation, and in the second step, we choose the largest quantile for $\widehat{\tau}_n - \tau_n$ among all β's in the confidence interval. The quantile chosen is also based on the *knife-edge* approximation.

Algorithm 1: Infeasible Uniform Inference

Input: Treatments and outcomes $(T_i, Y_i)_{i\in[n]}$, MPLE-estimator $\widehat{\beta}_n$, an upper bound K_n such that $\kappa_2^{\frac{1}{2}} \leq K_n$, confidence level parameters $\alpha_1, \alpha_2 \in (0,1)$.

Output: An $(1-\alpha_1-\alpha_2)$ prediction interval $\mathcal{C}^\dagger(\alpha_1,\alpha_2)$ for τ_n.

Get the maximum pseudo-likelihood estimator $\widehat{\beta}_n$ of β;

Define the $(1-\alpha_1)$-confidence region given by $\mathtt{I}(\alpha_1) = \{\beta \in [0,1] : 1-\widehat{\beta}_n \in [\mathtt{q}, \infty)\}$, where $\mathtt{q} = \inf\{q : \mathbb{P}[\min\{\max\{\mathsf{T}_{c_{\beta,n},n}^{-2} - \mathsf{T}_{c_{\beta,n},n}^{2}/(3n), 0\}, 1\} \leq q] \geq \alpha_1\}$;

Take $\mathtt{U} = \sup_{\beta\in\mathcal{I}(\alpha_1)} H_n(1-\frac{\alpha_2}{2}; K_n, K_n, c_{\beta,n})$, $\mathtt{L} = \inf_{\beta\in\mathcal{I}(\alpha_1)} H_n(\frac{\alpha_2}{2}; K_n, K_n, c_{\beta,n})$.

return $\mathtt{C}^\dagger(\alpha_1,\alpha_2) = [\widehat{\tau}_n + \mathtt{L}, \widehat{\tau}_n + \mathtt{U}]$.

Theorem 4.2 (Infeasible Uniform Inference)**.** *Suppose Assumptions 1, 2 and 3 hold, and let K_n be a sequence such that $\kappa_2^{\frac{1}{2}} \leq K_n$. Then, the prediction interval given by Algorithm 1 satisfies* $\liminf_{n\to\infty} \inf_{\beta\in[0,1]} \mathbb{P}_\beta(\tau_n \in \mathcal{C}^\dagger(\alpha_1,\alpha_2)) \geq 1-\alpha_1-\alpha_2$.

Theorem 4.2 gives a lower bound on the coverage of the proposed confidence region. Algorithm 2 can be implemented without the knowledge of the parameter of the Ising treatment model, but requires knowledge of κ_1 and κ_2. A fully feasible implementation is discussed next.

5 Implementation

The unknown parameters κ_1 and κ_2 capture moments of the underlying random graph structure. Building on [13], we propose a resampling method for consistent estimation of those parameters under an additional nonparametric assumption on the outcome equation.

Assumption 4. *Suppose $f_i(\cdot,\cdot) = f(\cdot,\cdot) + \varepsilon_i$, where $f(t,\cdot)$ is 4-times continuously differentiable on $[0,1]$ for $t \in \{0,1\}$, and $(\varepsilon_i : 1 \leq i \leq n)$ are i.i.d and independent of $\mathbf{E}$ and $\mathbf{T}$, with $\mathbb{E}[\varepsilon_i] = 0$ and $\mathbb{E}[|\varepsilon_i|^{2+\nu}] < \infty$ for some $\nu > 0$.*

This assumption allows for nonparametric learning the regression function f. In Section 4 in the supplementary material, we provide one example of such learner, but here we remain agnostic and thus present high-level conditions. This step aims to find a consistent estimate for both the function f and its derivative $\frac{\partial f(\cdot,x)}{\partial x}$, which can be achieved through the introduction of Assumption 4. We propose the following novel algorithm for estimating κ_2 based on resampling methods.

Algorithm 2: Estimation of κ_2

Input: Treatments and outcomes $(T_i, Y_i)_{i\in[n]}$, realized graph $\mathbf{E}$, non-parametric learner $\widehat{f}$ of f.

Output: An upper bound $\widehat{K}_n$ for κ_2.

Generate a new sample $(T_i^* : 1 \leq i \leq n)$ with $\beta = 0$;
Take $M_j^* = \sum_{l\neq j} E_{jl}T_l^*$, $N_j^* = \sum_{l\neq j} E_{jl}$, $M_{j,(i)}^* = \sum_{l\neq i,j} E_{jl}T_l^*$, $N_{j,(i)}^* = \sum_{l\neq i,j} E_{jl}$;
Take $\widehat{\varepsilon}_i = Y_i - T_i\widehat{f}(1, \frac{\sum_{j\neq i} E_{ij}T_j}{\sum_{j\neq i} T_j}) - (1-T_i)\widehat{f}(0, \frac{\sum_{j\neq i} E_{ij}T_j}{\sum_{j\neq i} T_j})$;
Take $\tau_{(i)}^a = n^{-1}\sum_{j\neq i} 2T_j^*(\widehat{f}(1, \frac{M_j^*}{N_j^*}) + \widehat{\varepsilon}_j) - 2(1-T_j^*)(\widehat{f}(0, \frac{M_j^*}{N_j^*}) + \widehat{\varepsilon}_j)$, and
$\tau_{(i)}^b = n^{-1}\sum_{j\in[n]} 2T_j^*(\widehat{f}(1, \frac{M_{j,(i)}^*}{N_{j,(i)}^*}) + \widehat{\varepsilon}_j) - 2(1-T_i^*)(\widehat{f}(0, \frac{M_{j,(i)}^*}{N_{j,(i)}^*}) + \widehat{\varepsilon}_j)$;
Take $\overline{\tau}^a = n^{-1}\sum_{i\in[n]} \tau_{(i)}^a$, $\overline{\tau}^b = n^{-1}\sum_{i\in[n]} \tau_{(i)}^b$, and
$\widehat{K}_n = n\sum_{i\in[n]}(\tau_{(i)}^a - \overline{\tau}^a + \tau_{(i)}^b - \overline{\tau}^b)^2$.
return $\widehat{K}_n$.

Our procedure consists of three steps. In step 1, we estimate f non-parametrically by $\widehat{f}$. In step 2, we construct *two* types of plug-in and leave-one-out estimator, denoted by $\{\tau_{(i)}^a\}_{i\in[n]}$ and $\{\tau_{(i)}^b\}_{i\in[n]}$ respectively. $\tau_{(i)}^a$ accounts for the randomness from flipping i-th unit's own treatment. $\tau_{(i)}^b$ accounts for randomness from flipping j-th unit's treatment, where j is a neighbor of i. In Step 3, we form our final variance estimator using the resampling based treatment effect estimators similar to the i.i.d. case. Formal results on the guarantees given in Lemma 16 in the supplementary material.

Algorithm 3: Feasible Uniform Inference

Input: Treatments and outcomes $(T_i, Y_i)_{i\in[n]}$, realized graph $\mathbf{E}$, non-parametric learner $\widehat{f}$ of f.

Output: A fully data-driven $(1-\alpha_1-\alpha_2)$ prediction interval $\widehat{\mathcal{C}}(\alpha_1, \alpha_2)$ for τ_n.

Get $\widehat{K}_n$ from Algorithm 2 using the treatments and outcomes $(T_i, Y_i)_{i\in[n]}$, the realized random graph $\mathbf{E}$, a non-parametric learner $\widehat{f}$ for f;
Get $\widehat{\mathcal{C}}(\alpha_1, \alpha_2)$ from Algorithm 1 given $(T_i, Y_i)_{i\in[n]}$ and $\widehat{K}_n$.
return $\widehat{\mathcal{C}}(\alpha_1, \alpha_2)$.

Theorem 5.1 (Feasible Robust Confidence Interval)**.** *Suppose Assumptions 1, 2, 3, and 4 hold. Suppose the non-parametric learner $\widehat{f}$ satisfies $\widehat{f}(\ell,\cdot) \in C_2([0,1])$, and $|\widehat{f}(\ell,\pi_*) - f(\ell,\pi_*)|= o_{\mathbb{P}}(1)$, $|\partial_2\widehat{f}(\ell,\pi_*) - \partial_2 f(\ell,\pi_*)|= o_{\mathbb{P}}(1)$, for $\ell \in \{0,1\}$. If $n\rho_n^3 \to \infty$, then the prediction interval given by Algorithm 3 satisfies*

$$\liminf_{n\to\infty} \sup_{\beta\in[0,1]} \mathbb{P}_\beta[\tau_n \in \widehat{\mathcal{C}}(\alpha_1, \alpha_2)] \geq 1 - \alpha_1 - \alpha_2.$$

6 Simulations

We study the finite sample performance of our robust inference procedure. Take $(U_i : 1 \leq i \leq n)$ i.i.d Uniform$([0,1])$-distributed, graph function $G(\cdot,\cdot) \equiv 0.5$ and density $\rho_n = 0.5$. The

Ising-treatments satisfy Assumption 3 with various n and β. Y_i has data generating process $Y_i = \mathbb{1}(T_i = 1)f(1, \frac{M_i}{N_i}) + \mathbb{1}(T_i = 0)f(0, \frac{M_i}{N_i}) + \varepsilon_i$, with $f(x_1, x_2) = x_1^2 + x_1(x_2+1)^2$, $(x_1, x_2) \in \mathbb{R}^2$ and $(\varepsilon_i : 1 \le i \le n)$ are i.i.d $N(0, 0.05)$ noise terms independent to $((U_i, T_i) : 1 \le i \le n)$. The Monte-Carlo simulations are repeated with 5000 iterations and look at the $1 - \alpha$ confidence interval with $\alpha = 0.1$.

Figure 1 (a) and (b) demonstrate the empirical coverage and interval length against β, while fixing $n = 500$. To compare multiple methods, `conserv` stands for Algorithm 3, "$\beta = 0$" stands for using the formula from Theorem 3.1, `Oracle` stands for using the law $n^{-1/2}\widehat{\kappa}_2^{1/2}\mathsf{Z} + n^{-1/4}\widehat{\kappa}_1\mathsf{W}_{c_{\beta,n}}$ from Theorem 4.1 with $c_{\beta,n} = \sqrt{n}(1-\beta)$ assumed to be known, and `Onestep` stands for Algorithm 1 but taking the first step confidence interval $\mathtt{I}(\alpha_1)$ to be the full range $[0, 1]$ instead. For interval length, `Simulated` stands for the true interval length from Monte-Carlo simulations. `Conservative` and `Onestep` remain conservative except when β is close to 1, due to the second step in Algorithm 1 taking maximum quantile from $\beta \in \mathtt{I}(\alpha_1)$; `Oracle` has empirical coverage close to $1 - \alpha$ and interval length close to the true interval length from Monte-Carlo simulation; the approach of plugging in $\beta = 0$ becomes invalid as β deviates from zero. Figure 1 (c) and (d) demonstrate log-log plots of interval length against sample size, fixing $\beta = 0$. While the Monte-Carlo interval length `Simulated` interval length $\propto n^{-0.52}$, consistent with the $\sqrt{n}$-convergence with $\beta = 0$, `Conserv` has interval length $\propto n^{-0.34}$, an effect of taking the maximum quantile among $\beta \in \mathtt{I}(\alpha_1)$.

7 Main Technical Contribution

This section reports the main novel technical result in our paper: a Berry-Esseen distributional approximation for Curie-Weiss magnetization with independent multipliers. This section is self-contained, but omitted details are given in the supplemental appendix.

Lemma 7.1 (Ising Berry-Esseen Bound). *For $\beta \ge 0$, suppose $\mathbb{P}[\mathbf{W} = \mathbf{w}] \propto \exp(\frac{\beta}{n}\sum_{i \ne j} w_i w_j)$, where $\mathbf{W} = (W_1, \cdots, W_n)^\top$, $\mathbf{w} = (w_1, \cdots, w_n)^\top \in \{-1, 1\}^n$, and $(X_1, \cdots, X_n)$ are i.i.d. with $\mathbb{E}[|X_i|^3] < \infty$, and independent of $\mathbf{W}$. Then:*
(1) Fix $\beta \in [0, 1]$, then $\sup_{t \in \mathbb{R}} |\mathbb{P}(\frac{1}{n}\sum_{i=1}^n X_i W_i \le t) - L_n(t; (\mathbb{E}[X_i], \mathbb{E}[X_i^2]), \beta)| = O(\mathtt{r}_{n,\beta})$, where $\mathtt{r}_{n,\beta} = n^{-1/2}$ for $\beta \in [0, 1]$, $\mathtt{r}_{n,\beta} = n^{-1/2}(\log n)^3$ for $\beta = 1$, where L_n is given in Theorem 3.1.
(2) $\sup_{\beta \in [0,1]} \sup_{t \in \mathbb{R}} |\mathbb{P}(\frac{1}{n}\sum_{i=1}^n X_i W_i \le t) - H_n(t; \mathbb{E}[X_i], \mathbb{E}[X_i^2], c_{\beta,n})| = O(n^{-1/2}(\log n)^3)$, where $c_{\beta,n} = \sqrt{n}(\beta - 1)$, and H_n is given immediately after Theorem 4.1.

These result generalize the Berry-Esseen bounds for Curie-Weiss magnetization $\frac{1}{n}\sum_{i=1}^n W_i$ with multiplers set to $X_i = 1$ for $i \in [n]$ obtained by [5] and [6]. Our generalized result differs from theirs only in a logarithmic term, allowing for fairly general weights with third moment bounded.

7.1 Proof Sketch of Lemma 7.1

The magnetization $n^{-1}\sum_{i=1}^n W_i$ has been studied using Stein's method [6, 4]. Due to the multipliers, the Stein's method can not be directly applied for $n^{-1}\sum_{i=1}^n X_i W_i$. We use a novel strategy based on the following de Finetti's lemma.

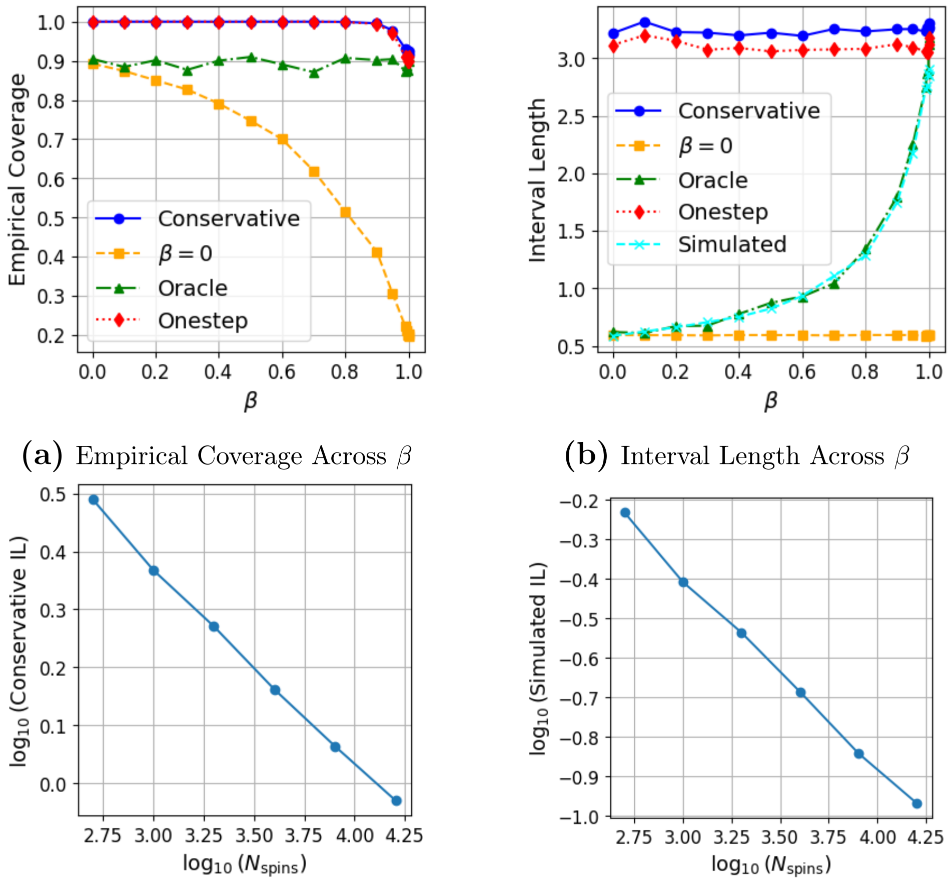

(a) Empirical Coverage Across β **(b)** Interval Length Across β

(c) `Conserv` Interval Length vs n **(d)** `Simulated` Interval Length vs n

Figure 1: (a) and (b) are empirical coverages and interval lengths of four methods across $\beta \in [0,1]$: `Conservative` and `Onestep` remain conservative except when β is close to 1; `Oracle` has empirical coverage close to $1-\alpha$ and interval length close to the true interval length from Monte-Carlo simulation; the approach of plugging in $\beta = 0$ becomes invalid as β deviates from zero. (c) shows `Conserv` interval length $\propto n^{-0.34}$. (d) shows `Simulated` interval length $\propto n^{-0.52}$.

de Finetti's Lemma. There exists a latent variable U_n such that $W_1, \cdots, W_n$ are i.i.d condition on U_n. Moreover, the density of U_n satisfies $f_{\mathsf{U}_n}(u) \propto \exp(-\frac{1}{2}u^2 + n \log\cosh(\sqrt{\beta/n}u)), u \in \mathbb{R}$.

We provide a proof sketch of Lemma 7.1 (2). Rigorous proofs for the other regimes given in Section 1 of the supplementary material. Denote by $\mathtt{C}$ an absolute constant, $\mathtt{K}$ a constant that only depends on the distribution of X_i, and $O(\cdot)$ is by an absolute constant. Throughout, take $c = \sqrt{n}(\beta - 1)$.

Step 1: Conditional Berry-Esseen. W_i's are i.i.d condition on U_n with $e(\mathsf{U}_n) = \mathbb{E}[X_i W_i | \mathsf{U}_n] = \mathbb{E}[X_i] \tanh(\sqrt{\beta/n}\mathsf{U}_n)$, and $v(\mathsf{U}_n) = \mathbb{V}[X_i W_i | \mathsf{U}_n] = \mathbb{E}[X_i^2] - \mathbb{E}[X_i]^2 \tanh^2(\sqrt{\beta/n}\mathsf{U}_n)$. Apply Berry-Esseen Theorem conditional on U_n, and take $\mathsf{Z} \sim \mathsf{N}(0,1)$ independent to U_n,

$$\sup_{t\in\mathbb{R}} |\mathbb{P}(\frac{1}{n}\sum_{i=1}^{n} X_i W_i \le t | \mathsf{U_n}) - \mathbb{P}(\sqrt{v(\mathsf{U}_n)}\mathsf{Z} + \sqrt{n} e(\mathsf{U}_n) \le t | \mathsf{U_n})| \le \mathtt{C}\mathbb{E}\left[|X_i|^3\right] v(\mathsf{U}_n) n^{-1/2}.$$

Lemma 2 in the supplementary material shows $\|\mathsf{U}_n\|_{\psi_2} \le \mathtt{C} n^{1/4}$, hence by concentration arguments, $\sup_{t\in\mathbb{R}}|\mathbb{P}(\frac{1}{n}\sum_{i=1}^n X_iW_i \le t) - \mathbb{P}(\sqrt{v(\mathsf{U}_n)}\mathsf{Z} + \sqrt{n}e(\mathsf{U}_n) \le t)| \le Kn^{-1/2}$.

Step 2: Non-Normal Approximation for $n^{-\frac{1}{4}}\mathsf{U}_n$. Consider $\mathsf{W}_n = n^{-1/4}\mathsf{U}_n$. By a change of variable from U_n and Taylor expand what is inside the exponent, we show W_n has density satisfying

$$f_{\mathsf{W}_n}(w) \propto \exp(-\frac{c}{2}w^2 - \frac{\beta_n^2}{12}w^4 + g(w)\beta_n^3 n^{-\frac{1}{2}}w^6),$$

where g is a bounded smooth function. We show based on sub-Gaussianity of W_n, with an upper bound of sub-Gaussian norm not depending on β, that the sixth order term is negligible and $\sup_{t\in\mathbb{R}}|\mathbb{P}(\mathsf{W}_n \le t) - \mathbb{P}(\mathsf{W} \le t)| = O((\log n)^3 n^{-1/2})$, where W has density proportional to $\exp(-\frac{c}{2}w^2 - \frac{\beta_n^2}{12}w^4)$.

Step 3: Concentration Arguments. Since Z is independent to $(\mathsf{U}_n, \mathsf{W}_n)$, we use data processing inequality and the previous two steps to show $\frac{1}{n}\sum_{i=1}^n X_iW_i$ is close to $n^{-\frac{1}{4}}v(n^{\frac{1}{4}}\mathsf{W}_c)^{\frac{1}{2}}\mathsf{Z} + n^{\frac{1}{4}}e(n^{\frac{1}{4}}\mathsf{W}_c))$. Lemma 2 in the supplementary appendix imply $\|\mathsf{W}\|_{\psi_2} \le \mathtt{K}$. By Taylor expanding $e(\cdot)$ and $v(\cdot)$ at 0, we show $n^{1/4}e(\mathsf{U}_n)$ is close to $\mathbb{E}[X_i]\mathsf{W}$ and $n^{-1/4}\sqrt{v(\mathsf{U}_n)}\mathsf{Z}$ is close to $n^{-\frac{1}{4}}v(n^{\frac{1}{4}}\mathsf{W})^{\frac{1}{2}}\mathsf{Z}$.

8 Discussion

This section discusses related results and future research directions.

8.1 Low Temperature Regime

The low temperature regime corresponds to $\beta > 1$, which was excluded from the main results presented. In this case the Hajek estimator converges to a different (conditional) direct treatment effect that also depends on which side of the half line $\operatorname{sgn}(m) = \operatorname{sgn}(\frac{2}{n}\sum_{i=1}^n T_i - 1)$ lies on, due to the convergence of $\frac{M_i}{N_i}$ to a two-point distribution depending on $\operatorname{sgn}(m)$. Define

$$\tau_{n,\ell} = \frac{1}{n}\sum_{i=1}^n \mathbb{E}[Y_i(1;\mathbf{T}_{-i}) - Y_i(0;\mathbf{T}_{-i})|f_i(\cdot),\mathbf{E},\operatorname{sgn}(m) = \ell], \qquad \ell \in \{-,+\},$$

which is a new causal predictand in the context of our causal inference model with interference. In the supplemental appendix, and under the assumptions imposed in the paper, we show that

$$\sup_{t\in\mathbb{R}}\max_{\ell\in\{-,+\}}|\mathbb{P}(\widehat{\tau}_n - \tau_{n,\ell} \le t|\operatorname{sgn}(m) = \ell) - L_n(t;\beta,\kappa_{1,\ell},\kappa_{2,\ell})| = O(\sqrt{\frac{n\log n}{(n\rho_n)^{p+1}}} + \frac{\log n}{\sqrt{n\rho_n}}),$$

where $\kappa_{s,\ell} = \mathbb{E}[(R_{i,\ell} + Q_{i,\ell})^s]$ with $R_{i,\ell} = f_i(1,\pi_\ell) - \mathbb{E}[f_i(1,\pi_\ell)] + f_i(0,\pi_\ell) - \mathbb{E}[f_i(0,\pi_\ell)]$ and $Q_{i,\ell} = \mathbb{E}[\frac{G(U_i,U_j)}{\mathbb{E}[G(U_i,U_j)|U_j]}(f_j'(1,\pi_\ell) - f_j'(0,\pi_\ell))|U_i]$, and

$$L_n(t;\beta,\kappa_1,\kappa_2) = \mathbb{P}\Big(n^{-1/2}\Big(\kappa_2(1-\pi_*^2) + \kappa_1^2\frac{\beta(1-\pi_*^2)}{1-\beta(1-\pi_*^2)}\Big)^{1/2}\mathsf{Z} \le t\Big)$$

with $\mathsf{Z} \sim \mathsf{N}(0,1)$ independent of m, π_* the positive root of $x = \tanh(\beta x)$, and $\pi_+ = \frac{1}{2} + \frac{1}{2}\pi_*$, $\pi_- = \frac{1}{2} - \frac{1}{2}\pi_*$. Inference for the conditional estimand is left for future works, with a challenge in a discontinuity in the estimand as we move from the critical regime to the low temperature regime.

8.2 Generalized Ising Model

In this work we assumed treatments are dependent through a fully connected graph. It is also of interest to study settings where the graph underlying treatment assignment has a block structure, or depends on unit-level properties. In the structured Ising setting, we might also consider estimation and inference for the block level or heterogenous direct average treatment effect.

Acknowledgements. Cattaneo gratefully acknowledges financial support from the National Science Foundation through DMS-2210561 and SES-2241575.

References

[1] Eric Auerbach and Max Tabord-Meehan. The local approach to causal inference under network interference. *arXiv preprint arXiv:2105.03810*, 2024.

[2] Bhaswar B Bhattacharya and Sumit Mukherjee. Inference in ising models. *Bernoulli*, 24(1):493–525, 2018.

[3] Peter J Bickel and Aiyou Chen. A nonparametric view of network models and newman–girvan and other modularities. *Proceedings of the National Academy of Sciences*, 106(50):21068–21073, 2009.

[4] Sourav Chatterjee. Spin glasses and stein's method. *Probability Theory and Related Fields*, 148(3):567–600, 2010.

[5] Sourav Chatterjee and Qi-Man Shao. Nonnormal approximation by stein's method of exchangeable pairs with application to the curie–weiss model. *Annals of Applied Probability*, 21(2):464–483, 2011.

[6] Peter Eichelsbacher and Matthias Loewe. Stein's Method for Dependent Random Variables Occuring in Statistical Mechanics. *Electronic Journal of Probability*, 15:962 – 988, 2010.

[7] Richard S Ellis. *Entropy, large deviations, and statistical mechanics*. Taylor & Francis, 2006.

[8] Miguel A. Hernán and James M. Robins. *Causal Inference: What If*. Boca Raton: Chapman & Hall/CRC, 2020.

[9] Michael G Hudgens and M Elizabeth Halloran. Toward causal inference with interference. *Journal of the American Statistical Association*, 103(482):832–842, 2008.

[10] Chanhwa Lee, Donglin Zeng, and Michael G Hudgens. Efficient nonparametric estimation of stochastic policy effects with clustered interference. *Journal of the American Statistical Association*, 2024.

[11] Michael P Leung. Causal inference under approximate neighborhood interference. *Econometrica*, 90(1):267–293, 2022.

[12] Shuangning Li and Stefan Wager. Random graph asymptotics for treatment effect estimation under network interference. *Annals of Statistics*, 50(4):2334–2358, 2022.

[13] Qiaohui Lin, Robert Lunde, and Purnamrita Sarkar. On the theoretical properties of the network jackknife. In *International Conference on Machine Learning*, pages 6105–6115. PMLR, 2020.

[14] Adam Lipowski, Dorota Lipowska, and António Luis Ferreira. Phase transition and power-law coarsening in an ising-doped voter model. *Physical Review E*, 96(3):032145, 2017.

[15] Charles F Manski. Identification of treatment response with social interactions. *The Econometrics Journal*, 16(1):S1–S23, 2013.

[16] Rajarshi Mukherjee, Sumit Mukherjee, and Ming Yuan. Global testing against sparse alternatives under ising models. *Annals of Statistics*, 46(5):2062–2093, 2018.

[17] Somabha Mukherjee, Jaesung Son, Swarnadip Ghosh, and Sourav Mukherjee. Efficient estimation in tensor curie-weiss and erdős-rényi ising models. *Electronic Journal of Statistics*, 18(1):2405–2449, 2024.

[18] Elizabeth L Ogburn, Oleg Sofrygin, Ivan Diaz, and Mark J Van der Laan. Causal inference for social network data. *Journal of the American Statistical Association*, 119(545):597–611, 2024.

[19] Eric J Tchetgen Tchetgen and Tyler J VanderWeele. On causal inference in the presence of interference. *Statistical methods in medical research*, 21(1):55–75, 2012.

[20] Gonzalo Vazquez-Bare. Identification and estimation of spillover effects in randomized experiments. *Journal of Econometrics*, 237(1):105237, 2023.