PrIsing: Privacy-Preserving Peer Effect Estimation via Ising Model

Abhinav Chakraborty[*]
University of Pennsylvania

Anirban Chatterjee[*]
University of Pennsylvania

Abhinandan Dalal[*]
University of Pennsylvania

Abstract

The Ising model, originally developed as a spin-glass model for ferromagnetic elements, has gained popularity as a network-based model for capturing dependencies in agents' outputs. Its increasing adoption in healthcare and the social sciences has raised privacy concerns regarding the confidentiality of agents' responses. In this paper, we present a novel (ε, δ)-differentially private algorithm specifically designed to protect the privacy of individual agents' outcomes. Our algorithm allows for precise estimation of the natural parameter using a single network through an objective perturbation technique. Furthermore, we establish error bounds for this algorithm and assess its performance on synthetic datasets and two real-world networks: one involving HIV status in a social network and the other concerning the political leaning of online blogs.

1 INTRODUCTION

The ubiquity of data available on interactions between agents in a system has led to several network models being developed to better understand and contemplate agents' responses in an interconnected environment. One such popular model is the Ising spin glass model, which was originally developed in statistical physics to model ferromagnetism. However, it has now gained popularity in applications in several social science domains, due to it's ease of interpretation and widespread applicability. Thomas Schelling's Ising-like model (Schelling (1971)) to explain racial segregation in US cities has become a standard practice in explaining urban dynamics (Fossett (2006)). Stauffer (2008) also discusses how the Ising model can be used to understand language dynamics and the adoption of linguistic features from different languages without external forces, a line of work pioneered in Nettle (1999), and later strongly reflected in future literature.

One of the interesting properties of Ising model (Equation (1), discussed in detail in Section 2), which is a joint distribution on the outcomes of the nodes $\boldsymbol{\sigma} = (\sigma_1, \cdots, \sigma_n) \in \{\pm 1\}^n$ given an arbitrarily encoded symmetric network information matrix $\boldsymbol{J}_n := ((\boldsymbol{J}_n(i,j)))$, is that it is easy to infer the influence of the neighboring nodes on the outcome of an individual node. This can be seen from the conditional probability of $\sigma_i = 1$ given the other node outcomes $\boldsymbol{\sigma}_{-i} := (\sigma_1, \ldots, \sigma_{i-1}, \sigma_{i+1}, \ldots, \sigma_n)$,

$$\mathbb{P}(\sigma_i = 1 | \boldsymbol{\sigma}_{-i}) = \frac{e^{\beta \sum_{j:j\neq i} \sigma_j \boldsymbol{J}_n(i,j)}}{e^{\beta \sum_{j:j\neq i} \sigma_j \boldsymbol{J}_n(i,j)} + e^{-\beta \sum_{j:j\neq i} \sigma_j \boldsymbol{J}_n(i,j)}},$$

which increases or decreases in β as $\sum_{j\neq i} \sigma_j \boldsymbol{J}_n(i,j)$, is positive or negative, respectively. This parameter β, often referred to as the inverse temperature in the physics literature, encapsulates the extent of influence of neighbors in the network (see for example, Daskalakis et al. (2020)).

However, the applicability of Ising model in social networks comes with its concern in privacy. In fact, Abawajy et al. (2016) and Zhou et al. (2008) discuss a multitude of privacy preservation techniques when presenting network data, particularly with the boom of current network data. The concerns of privacy in social network analysis has indeed been a concern echoed by many (Backstrom et al. (2007); Kleinberg (2007); Srivastava et al. (2008)), for instance, a powerful adversary with access to others' data might be able to conclude one's outcome from a non-private algorithm, particularly since the outcomes in a network are dependent (Liu et al. (2016)). In fact, Ising model in itself has been or can be used to study several sensitive or potentially sensitive data on:

- **Transmission of contagious diseases:** For instance Mello et al. (2021) study epidemic transmission concepts from Covid-19 using Ising model.

Proceedings of the 27[th] International Conference on Artificial Intelligence and Statistics (AISTATS) 2024, Valencia, Spain. PMLR: Volume 238. *The authors contributed equally to this work.

- **Tax evasion dynamics** and peers' influence on such behavior, as studied by Zaklan et al. (2009) using Ising model, and further enriched by Pickhardt and Seibold (2014), Giraldo-Barreto and Restrepo (2021).
- **Enforcing social behavior** as discussed by Cajueiro (2011) using Ising model for modeling harmful behaviors, like smoking decisions (Krauth (2006)), criminal behavior (Glaeser et al. (1996)), investment decisions (Duflo and Saez (2002)), etc.
- **Sexually transmitted diseases**, like that studied by Potterat et al. (2002), for HIV transmission in a network based study from Colorado Springs. They also incorporate several sensitive information like drug injection usage of agents involved.

Such applications motivate the need for privacy-preserving techniques for analysis. Indeed, quite a few of the papers cited study the model through simulations, as such data is hard to collect and are often unreliable due to the potential untruthful reporting for privacy concerns. From eavesdropping medical and financial agencies, to incriminating evidence, social taboos and voting freedoms; these examples show why privacy is of utmost importance in studying these behaviors, so that truthful data collection can be incentivized and valid inferences can be drawn while ensuring the individuals' privacy.

1.1 Related Works

There has been a growing literature for theoretical analysis of the Ising model, and advances have been made in understanding the non-standard estimation techniques in regard to the same. Chatterjee (2007) is one of the pioneers in this literature, where he shows the $\sqrt{n}$ consistency of the maximum pseudo-likelihood estimator. On the other hand, Bhattacharya and Mukherjee (2018) extends this result to $\sqrt{a_n}$-consistency based on conditions of the log partition function, thus completing the result of consistency for all the regimes. We build on these previous works to incorporate the non-statistical constraint of differential privacy, and quantify the loss of efficiency due to the privacy requirement. Mukherjee and Ray (2022) also analyze the difficulty in the estimation of the parameter of the Ising model in certain regimes, and draws parallels with the joint estimation strategies demonstrated in Ghosal and Mukherjee (2020). Theoretical explorations into distribution testing with Ising Models have been studied in Daskalakis et al. (2019).

In this work, we use techniques from Kifer et al. (2012). They however deal with independent data structures, which is in stark contrast with the dependent structure of that of the Ising model, thus requiring the necessity for developing new arguments and drawing insights from the Ising literature to prove error bounds on the differentially private estimator.

However, it must be noted that the notion of outcome-differential privacy is different from the usual edge-differential privacy (eg: Mohamed et al. (2022), Chen et al. (2023), etc.) or node-differential privacy (eg: Kasiviswanathan et al. (2013), Blocki et al. (2013), etc.) often considered in network privacy. In the Ising model, the network information incorporated into the $\boldsymbol{J}_n$ matrix is considered non-stochastic, and we are instead interested in the outcomes $\sigma_i \in \{\pm 1\}$, $i \in \{1, \cdots, n\}$ of the nodes. Taking up the tax-evasion example to elucidate, the choice to evade taxes, taken to be binary as ± 1 (which can be affected by peers' choices), are sensitive and hence require privacy guarantees. This would give the respondents plausible deniability against financially criminal behavior, while still allowing the researchers to study the influence of peers in such behavioral models.

To our knowledge, the work by Zhang et al. (2020) is the only one discussing differential privacy in Ising models. They focus on keeping the dataset private during both structure learning and parameter estimation from multiple realizations of the results.

However, their privacy concept differs from ours. They adopt a privacy model wherein the collection $\{\sigma_i\}_{i=1}^n$ treated as a singular unit of data. This approach is particularly designed for scenarios involving the observation of multiple independent replicates of datasets, each consisting of n sign flips. In contrast, our approach considers each node's outcome as an individual unit, observing only a single collection of n sign flips. This presents a more individualistic perspective on the preservation of privacy, making our applicability significantly different from theirs.

Our Contributions can be summarized as follows:

- Primarily we study the problem of preserving privacy for node outcomes of a network, in the context of parameter (β) estimation in an Ising model. This parameter is estimated with a single realization of the network, and can be used to infer about the extent of interference between node outcomes in a network. The problem of preserving node-outcome privacy in a single network data, as far as our knowledge is concerned, has not been studied before.
- We prove error bounds for our algorithm, quantify the cost of privacy and complement the theoretical results with extensive simulation study with Erdős-Rényi random graphs.
- Finally, we evaluate the performance on two real-

world networks-HIV status of individuals in a social network, and political leaning of online blogs that link to one another.

The article is organised as follows: Section 1 provides an introduction discussing the importance of privacy of node outcomes along-with the current state of the literature, Section 2 puts the problem formally in terms of the model and the privacy guarantee being provided, Section 3 discusses our algorithm and proves privacy and error guarantees, and Section 4 evaluates its performance through numerical experiments and real life data. Finally, Section 5 provides closing discussions. All the proofs and additional experiments are presented in the Supplementary Material. Code for all the experiments can be found in the Github repository `https://github.com/anirbanc96/PrIsing`.

2 PROBLEM FORMULATION

Two vectors $\boldsymbol{\tau}, \boldsymbol{\tau'} \in \{\pm 1\}^n$ are said to be adjacent if they differ in at most one coordinate. The notion of differential privacy tries to constraint an algorithm by limiting its output variability for adjacent training input $\boldsymbol{\tau}$ and $\boldsymbol{\tau'}$ (Dwork et al. (2014)).

Definition 2.1. (Dwork (2006); Dwork et al. (2006a)) A randomized algorithm $\mathcal{M}$ is said to be node outcome (ε, δ)-differentially private ($\varepsilon > 0$, $\delta \geqslant 0$) if

$$\mathbb{P}(\mathcal{M}(\boldsymbol{\sigma}) \in S) \leqslant e^{\varepsilon} \mathbb{P}(\mathcal{M}(\boldsymbol{\sigma'}) \in S) + \delta$$

for any adjacent vectors $\boldsymbol{\sigma}, \boldsymbol{\sigma'} \in \{\pm 1\}^n$ and all events S in the output space of $\mathcal{M}$. When $\delta = 0$, the algorithm is said to be ε-differentially private.

Note that in Definition 2.1, we have not specified anything about the graph information. Indeed, the privacy protection is for the outcomes on the nodes, even when the graph information is perfectly available to an adversary.

Given a non-negative symmetric matrix $\boldsymbol{J}_n \in \mathbb{R}^{n \times n}$ (encapsulating network information) with 0 on its diagonal, the Ising model constitutes assigning a probability distribution on a vector of dependent ± 1 random variables $\boldsymbol{\sigma} = (\sigma_1, \cdots, \sigma_n)$, given by a parametric distributions on $S_n := \{-1, 1\}^n$ given by

$$\mathbb{P}_\beta(\boldsymbol{\sigma} = \boldsymbol{\tau}) = \frac{1}{2^n} \exp\left(\frac{1}{2}\beta H_n(\boldsymbol{\tau}) - F_n(\beta)\right); \quad (1)$$

with $\beta \geqslant 0$, where

$$H_n(\boldsymbol{\tau}) = \boldsymbol{\tau}^T \boldsymbol{J}_n \boldsymbol{\tau} = \sum_{1 \leqslant i,j \leqslant n} \boldsymbol{J}_n(i,j)\tau_i\tau_j; \quad \boldsymbol{\tau} \in S_n \quad (2)$$

and $F_n(\beta)$ is the log-partition function determined by the normalizing constraint $\sum_{\boldsymbol{\tau} \in S_n} \mathbb{P}_\beta(\boldsymbol{\sigma} = \boldsymbol{\tau}) = 1$ resulting in the formulation

$$\begin{aligned} F_n(\beta) &:= \log\left[\frac{1}{2^n} \sum_{\tau \in S_n} \exp\left(\frac{1}{2}\beta H_n(\boldsymbol{\tau})\right)\right] \\ &= \log \mathbb{E}_0 \exp\left(\frac{1}{2}\beta H_n(\boldsymbol{\sigma})\right) \end{aligned}$$

where $\mathbb{E}_0$ denotes the expectation over $\boldsymbol{\sigma}$ distributed as $\mathbb{P}_0$ (the uniform measure on S_n). The parameter β , in parallel with the physics literature, is often known as the inverse temperature and captures the strength of dependence in the various entries of $\boldsymbol{\sigma}$.

A very popular way of estimating β is obtaining the maximum pseudo-likelihood estimator (MPLE) $\hat{\beta}_n(\boldsymbol{\sigma})$ (Bhattacharya and Mukherjee (2018); Chatterjee (2007)), given by

$$\hat{\beta}_n(\boldsymbol{\sigma}) = \arg\max_{\beta} \prod_{i=1}^n f_i(\beta, \sigma_i) \quad (3)$$

where $f_i(\beta, \sigma_i)$ is the conditional probability density of σ_i given $\boldsymbol{\sigma}_{-i}$, under parameter β.

For any $\boldsymbol{\tau} \in S_n$, defining the function $L_{\boldsymbol{\tau}} : [0, \infty) \to \mathbb{R}$ as

$$L_{\boldsymbol{\tau}}(x) := -\frac{1}{n} \sum_{i=1}^n m_i(\boldsymbol{\tau})(\tau_i - \tanh(x m_i(\boldsymbol{\tau}))), \quad (4)$$

where

$$m_i(\boldsymbol{\tau}) := \sum_{j=1}^n \boldsymbol{J}_n(i,j)\tau_j, \quad (5)$$

it can be verified (see for example, Chatterjee (2007); Bhattacharya and Mukherjee (2018)) that

$$\hat{\beta}_n(\boldsymbol{\sigma}) := \inf\{x \geqslant 0 : L_{\boldsymbol{\sigma}}(x) = 0\}, \quad (6)$$

interpreting the infimum of an empty set as ∞ as usual, where $\boldsymbol{\sigma} \sim \mathbb{P}_\beta$. Henceforth the dependence on σ is suppressed with $\hat{\beta}_n := \hat{\beta}_n(\boldsymbol{\sigma})$ denoting the MPLE of β, and the function defined in Equation (4) is referred to as the pseudo log-likelihood function. Furthermore, in the following we use the notation $t_n = \Theta(s_n)$ to denote $t_n = O(s_n)$ and $s_n = O(t_n)$. Also we say a random variable $X_n = O_p(t_n)$ to imply that for any $\varepsilon > 0$ there exists $M_\varepsilon > 0$ such that,

$$\mathbb{P}\left[|X_n/t_n| > M_\varepsilon\right] \leqslant \varepsilon \text{ for all large enough } n.$$

3 OUR METHOD

Our algorithm for private parameter estimation in one-parameter Ising model is given in Algorithm 1.

Algorithm 1 Private Estimation in One-parameter Ising Model (`PrIsing`)

Require: $\boldsymbol{\sigma} = (\sigma_1, \cdots, \sigma_n)$, privacy parameters $\varepsilon > 0, \delta \geqslant 0$, symmetric matrix $\boldsymbol{J}_n \in \mathbb{R}^{n \times n}$ with non-negative entries such that $\boldsymbol{J}_n(i,i) = 0\ \forall 1 \leqslant i \leqslant n$.

Set $m_i(\boldsymbol{\sigma}) = \sum_{j=1}^n \boldsymbol{J}_n(i,j)\sigma_j$; $i = 1, \cdots, n$; $L_{\boldsymbol{\sigma}}(\beta) = -\frac{1}{n}\sum_{i=1}^n m_i(\boldsymbol{\sigma})(\sigma_i - \tanh(\beta m_i(\boldsymbol{\sigma}))$.

Set $d_i = n\sum_{j=1}^n \boldsymbol{J}_n(i,j)$ for all $i = 1(1)n$.

Set $\zeta = \max_j \left\{8\dfrac{d_j}{n}\right\}$.

if $\delta > 0$ **then**

Sample $b \in \mathbb{R}$ from $\nu(b; \varepsilon, \delta) = \mathcal{N}(0, \gamma^2)$ where

$$\gamma = \frac{\zeta\sqrt{8\log(2/\delta) + 4\varepsilon}}{\varepsilon}$$

else if $\delta = 0$ **then**

Sample $b \in \mathbb{R}$ from $\nu(b; \varepsilon, 0) = \mathrm{Lap}\left(0, 2\zeta/\varepsilon\right)$

end if

Set $\Delta \geqslant \max_j \left\{\frac{24}{\varepsilon n}\sum_{i=1}^n d_i \boldsymbol{J}_n(i,j)\right\}$

return $\hat{\beta}^{\mathtt{priv}} = \inf\{\beta \geqslant 0 : L_{\boldsymbol{\sigma}}(\beta) + \Delta\beta/n + b/n = 0\}$

Although the non-private estimate is given by the MPLE obtained through equation (6), our algorithm builds on the MPLE method by equating the pseudo-likelihood equation not to 0, but to a random noise perturbation, calibrated according to the privacy requirement. The algorithm builds on Kifer et al. (2012) and uses similar proof ideas by bounding the ratio of the gradients of the MPLE equation, and the density of the noise. However, in the former ratio, they could use an identical bound as their data points were i.i.d., whereas due to the dependent structure of the MPLE equation ($m_i(\boldsymbol{\sigma})$ depends on $\boldsymbol{\sigma}_{-i}$'s), we need to obtain the noise variance calibrated to the global-sensitivity (Dwork et al. (2006b)) of the pseudo-loglikelihood function, demonstrated in proof of Theorem 3.1 in Section A.

Theorem 3.1. *Given any $\epsilon > 0$ and $\delta \geqslant 0$, Algorithm 1 is (ε, δ)-differentially private on node-outcome $\boldsymbol{\sigma}$.*

Next, we quantify the error bound of our Algorithm 1, and quantify the cost of privacy in contrast with the non-private error bound. Under regularity conditions, Bhattacharya and Mukherjee (2018) shows $\sqrt{a_n}$ consistency of the non-private estimators, where a_n is determined by conditions on the log-partition function. In the following result we adopt a conditions similar to those required for consistency of the non-private estimator and provide the error bounds attained by $\hat{\beta}^{\mathtt{priv}}$ from Algorithm 1.

Theorem 3.2 (Simpler Version of Theorem B.1)**.** *Let $\sup_{n\geqslant 1}\|\boldsymbol{J}_n\| < \infty$, and let $\beta_0 > 0$ be fixed. Suppose $\{a_n : n \geqslant 1\}$ is a sequence such that, $a_n \to \infty$ as $n \to \infty$ and,*

(i) $F_n(\beta) = \Theta(a_n)$ for all β in a neighbourhood of β_0,

(ii) $\mathbb{E}_{\beta_0}\left[\sum_{i=1}^n m_i(\boldsymbol{\sigma})^2\right] = o(a_n)$,

(iii) $\sum_{i=1}^n\sum_{j=1}^n \boldsymbol{J}_n(i,j)^2 = O(a_n)$.

Then, whenever Δ is chosen to be the smallest permitted by Algorithm 1, the estimator $\hat{\beta}^{priv}$ satisfies,

$$|\hat{\beta}^{priv} - \beta_0| = O_p\left(\frac{1}{\sqrt{a_n}} + \frac{\lambda_n\sqrt{\log(2/\delta)}}{a_n\varepsilon}\right) \text{ if } \delta > 0,$$

and,

$$|\hat{\beta}^{priv} - \beta_0| = O_p\left(\frac{1}{\sqrt{a_n}} + \frac{\lambda_n}{a_n\varepsilon}\right) \text{ if } \delta = 0.$$

where $\lambda_n := 1 \vee \|\boldsymbol{J}_n\|^2_{1\to\infty}$, with $\|\cdot\|_{1\to\infty}$ denoting the vector induced 1-norm of a matrix.

The first term in the rate, denoted as $1/\sqrt{a_n}$, represents the non-private rate inherent in the problem. The additional term $\frac{\lambda_n}{a_n\varepsilon}$ accounts for the privacy cost introduced by differential privacy (DP) constraints. This is in line with the privacy literature, where it is widely observed that the privacy cost manifests as a higher-order term, and its dependence on n (in our case, a_n) is quadratic in nature (as discussed in, for example, Acharya et al. (2021)).

However, It is important to note that Acharya et al. (2021) deal with problems exhibiting an independent and identically distributed (i.i.d.) structure. Therefore, while drawing analogies, one should consider this context carefully. We anticipate that in the presence of dependence structures among observations, privacy is compromised to a greater extent (Liu et al. (2016)). This compromise is quantified by the parameter λ_n, where larger values of λ_n signify an exacerbation of the privacy cost.

3.1 Examples

In this section we consider specific examples of the underlying network to quantify the error bound obtained by the private estimator $\hat{\beta}^{\mathtt{priv}}$ and contrast with the corresponding non-private estimator.

3.1.1 Degree Regular Graphs

For Ising Models on degree regular graphs G_n, the matrix $\boldsymbol{J}_n$ in (1) becomes $\boldsymbol{J}_n = \boldsymbol{A}_n/D_n$ where $\boldsymbol{A}_n$ is the adjacency matrix of the graph G_n and D_n is the degree. This encompasses Ising models on complete graphs, random regular graphs and lattices which have been comprehensively investigated in probability and statistical physics (see Dembo and Montanari (2010); Levin et al. (2010)). For such $\boldsymbol{J}_n$ the parameter $\lambda_n = 1$, and hence by Theorem 3.2 and Corollary 3.1 from Bhattacharya and Mukherjee (2018) we have the following result.

Corollary 3.1. *Fix $\beta_0 > 0$. Then for any sequence of D_n regular graphs G_n,*

$$\left|\hat{\beta}^{priv} - \beta_0\right| = \begin{cases} O_p\left(\sqrt{\frac{D_n}{n}} + \frac{D_n}{n\varepsilon}\eta_\delta\right) & 0 < \beta_0 < 1 \\ O_p\left(\frac{1}{\sqrt{n}} + \frac{1}{n\varepsilon}\eta_\delta\right) & \beta_0 > 1 \end{cases}$$

where $\eta_\delta = \log(2/\delta)$ if $\delta > 0$, and $\eta_\delta = 1$ if $\delta = 0$.

3.1.2 Erdős-Rényi Graphs

Consider G_n to be a sequence of Erdős-Rényi random graphs on n vertices with edge probability p_n. For Ising Models with such an underlying network structure the matrix $\boldsymbol{J}_n$ from (1) is taken as $\boldsymbol{A}_n/np_n$, where $\boldsymbol{A}_n$ is the adjacency matrix of the network G_n. It is easy to infer that for large enough n the parameter $\lambda_n \leqslant 2$ with high probability. Combining Theorem 3.2 and Corollary 3.2 from Bhattacharya and Mukherjee (2018) we have the following result.

Corollary 3.2. *Fix $\beta_0 > 0$. Then for a sequence of Erdős-Rényi random graphs G_n with edge probability $\frac{\log n}{n} \ll p_n \leqslant 1$,*

$$\left|\hat{\beta}^{priv} - \beta_0\right| = \begin{cases} O_p\left(\sqrt{p_n} + \frac{p_n}{\varepsilon}\eta_\delta\right) & 0 < \beta_0 < 1 \\ O_p\left(\frac{1}{\sqrt{n}} + \frac{1}{n\varepsilon}\eta_\delta\right) & \beta_0 > 1 \end{cases}$$

where $\eta_\delta = \log(2/\delta)$ if $\delta > 0$, and $\eta_\delta = 1$ if $\delta = 0$.

In Corollary 3.1 and 3.2, we observe similar phase transition behavior as in non-private scenarios, with a distinct change in the rate of convergence occurring at $\beta_0 = 1$. Specifically, when $\beta_0 < 1$, we witness a phenomenon reminiscent of mean estimation problems, where the observed rate follows $O_p(\sqrt{\frac{d}{n}} + \frac{d}{n\varepsilon})$ (as discussed in Cai et al. (2021)), where d represents the dimensionality of the problem. In these regimes, D_n and np_n signifies the intrinsic dimensionality of our problem. Notably, the cost of privacy becomes more pronounced for larger values of D_n and np_n, indicating a higher degree of dependence in our model. However, intriguingly, this intensification of the privacy cost diminishes in the high-dependence regime $\beta_0 > 1$. Here, the cost of privacy no longer exhibits dependency on the graph's intrinsic dimensionality, paralleling the non-private rate.

4 NUMERICAL EXPERIMENTS

To complement the error guarantees in Section 3, we perform numerical experiments to evaluate the performance of our method. We conduct a set of simulations on Erdős-Rényi graphs and on two real datasets- HIV status of individuals in a social network, and political leaning of online blogs, repeating the simulation 500 times to plot the results. The Erdős-Rényi simulations show a regime change in the estimation rate of β, a phenomenon also seen in the nonprivate setting Bhattacharya and Mukherjee (2018). On the other hand, the real network experiments show that the validity of the method is upheld in realistic networks as well.

4.1 Experiments on Erdős-Rényi graphs

In this section, we rigorously validate our theory through a comprehensive simulation study. We begin by providing a detailed description of the simulation setup.

Simulation Setup

We aim to estimate the parameter β based on a dataset consisting of n observations. These observations are generated from an Ising model, where the underlying dependency graph G_n is created using the Erdos-Renyi model with a designated parameter p_n. $\boldsymbol{J}_n$ is taken to be $\boldsymbol{A}(G_n)/np_n$, where $\boldsymbol{A}(G)$ is taken to be the adjacency matrix of a network G, ie, the corresponding H from (2) becomes

$$H(\boldsymbol{\tau}) = \frac{1}{np_n}\boldsymbol{\tau}^T \boldsymbol{A}(G_n)\boldsymbol{\tau}.$$

Our investigation into the performance of our proposed estimator encompasses three primary simulation studies, as discussed below.

1. Impact of β on Our Estimator We want to compare the performance of our $\hat{\beta}^{\texttt{priv}}$ with $\hat{\beta}_n$ over a range of true β. We set $n = 2000$, $\delta = 1/n$ and $\varepsilon = 5$ for the comparison, and p_n is chosen to be $n^{-\frac{1}{3}}$.

Figure 1 shows the performance of the estimator over $\beta \in [0, 2]$ alongwith 1 standard deviation errorbars. Notice the phase transition at 1 in the error-bars produced in both the non-private and private estimators. This is in line with what we expect in theory, as $\hat{\beta}_n$ is $\frac{1}{\sqrt{p_n}} = n^{\frac{1}{6}}$ consistent for $\beta < 1$ and $\sqrt{n}$ consistent for

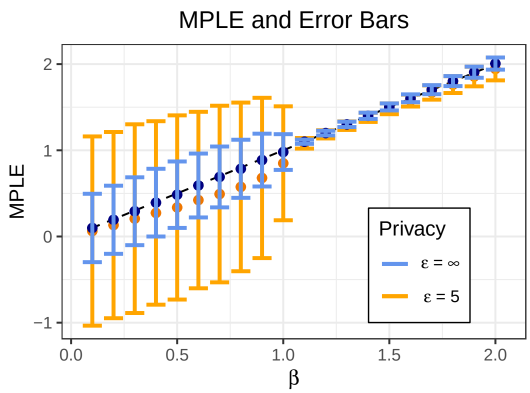

Figure 1: Private and non-private MPLE in an Ising model on and Erdős-Rényi random graph.

$\beta > 1$ (see Corollary 3.2), and the private estimator follows the same trend.

2. Effect of the Number of Observations n on Mean Squared Error (MSE) Next we focus on how the MSE of the estimators vary with the number of nodes n in G_n. p_n δ are still taken to be as before, and we use a range of ε for comparison. Since the rate of consistency is different in the two regimes of β, we plot the MSE vs n at $\beta = 0.5$ (Figure 2) and at $\beta = 1.5$ (Figure 3).

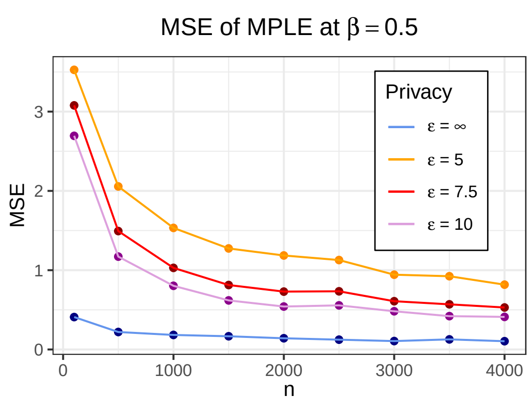

Figure 2: Effect of n on MSE of MPLE in an Ising model on and Erdős-Rényi random graph with $\beta = 0.5$

Note that both Figure 2 and 3 show a downward trend in MSE as expected, but the speed at which the trend dips down varies over ε. This effect can be attributed to the cost of privacy in Corollary 3.2, in particular for $\beta < 1$ the cost is $\frac{p_n}{\varepsilon} = \frac{1}{n^{1/3}\varepsilon}$ while for $\beta > 1$ the cost is $\frac{1}{n\varepsilon}$, which complements the observation that for $\beta > 1$ the reduction in MSE is much faster than $\beta < 1$.

3. Effect of Edge Density of G_n on MSE Here, we investigate the relationship between the edge density of the underlying graph G_n and the resulting Mean Squared Error (MSE) of our estimator. Since the number of edges is expected to be around $n^2 p_n$, we take $p_n = n^{-\alpha}$, and vary α to contrast the edge densities.

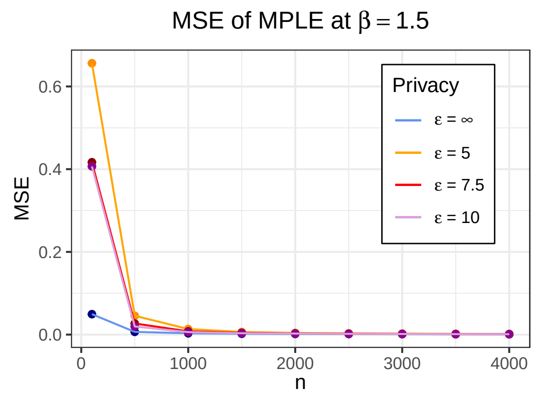

Figure 3: Effect of n on MSE of MPLE in an Ising model on and Erdős-Rényi random graph with $\beta = 1.5$

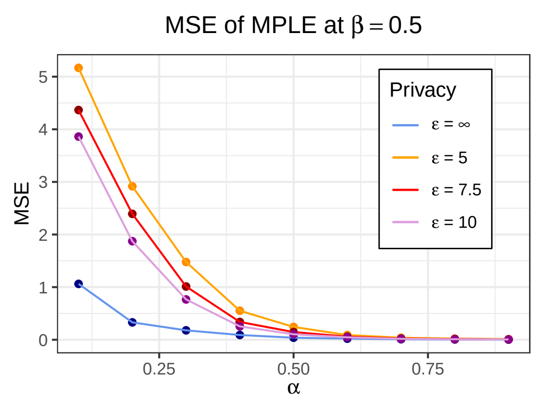

Figure 4: Effect of p_n on MSE of MPLE in an Ising model on and Erdős-Rényi random graph with $\beta = 0.5$

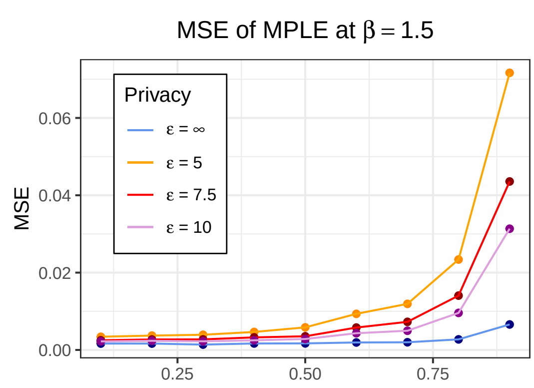

Figure 5: Effect of p_n on MSE of MPLE in an Ising model on and Erdős-Rényi random graph with $\beta = 1.5$

Figure 4 and 5 show the MSE of MPLE for a range of α. Recall from Corollary 3.2 that rate of convergence in the high temperature regime ($\beta < 1$) is inversely proportional to α, which explains the relation betweeen MSE and α in Figure 4. On the other hand in the low temperature regime ($\beta > 1$) the rate of convergence is $\sqrt{n}$, independent of the choice of α, reflected in non-private curve in Figure 5. However, in the private case, following Theorem 3.2, the increment in error with an increasing α can be attributed to the parameter λ_n, which is approximately 1 with additional error proportional to α with high probability.

4.2 Real world networks: Experiments & Real data

In the second set of simulations, we adopt real networks from two datasets, HIV transmission in social networks, and political affiliations of online blogs. We conduct synthetic experiments adopting the corresponding networks as fixed, and perform simulations of Ising model realizations on these networks with $\boldsymbol{J}_n = \boldsymbol{D}(G_n)^{-1/2}\boldsymbol{A}(G_n)\boldsymbol{D}(G_n)^{-1/2}$, where $\boldsymbol{D}(G) = \text{diag}(D_1(G), \cdots, D_n(G))$ is a diagonal matrix of the degrees of the nodes in a graph G. This leads to

$$H(\boldsymbol{\tau}) = \boldsymbol{\tau}^T \boldsymbol{D}(G_n)^{-\frac{1}{2}} \boldsymbol{A}(G_n) \boldsymbol{D}(G_n)^{-\frac{1}{2}} \boldsymbol{\tau} \qquad (7)$$

which can thus handle moderate degree heterogeneity in the network G_n, and is a generalization of the scaling used for regular or Erdős-Rényi graphs. In fact this choice of $\boldsymbol{J}_n$ can be linked to the normalized graph Laplacian $\mathcal{L}_n$ as $\boldsymbol{J}_n = \boldsymbol{I}_n - \mathcal{L}_n$ (Chung (1997)), and have been extensively used in node label classification problems (Li et al. (2018), Zhou et al. (2020); Wu et al. (2020)).

4.2.1 HIV status of individuals in a social network

We consider the network of HIV status of individuals in Colorado Springs with 403 individuals in the years 1988-1993 pooled together (Morris and Rothenberg (2011)), of which 23 have HIV status positive. Clearly this is a network where the privacy of the node outcomes is of great importance, and the outcomes are heavily imbalanced in the network, as given by the numbers as well as the network plot in Figure 6.

We conduct synthetic experiments simulating Ising model realizations from this network under the Laplacian scaling as in Equation 7, and plot the results in Figure 7. $\varepsilon = 5$ and $\delta = 1/n$ are chosen for the plot. Both the private as well as non-private estimate appears to be consistent around the true β.

Next we perform the analysis of the real data. The non-private $\hat{\beta}_n = 1.8$, and we produce $\hat{\beta}^{\texttt{priv}}$ and take

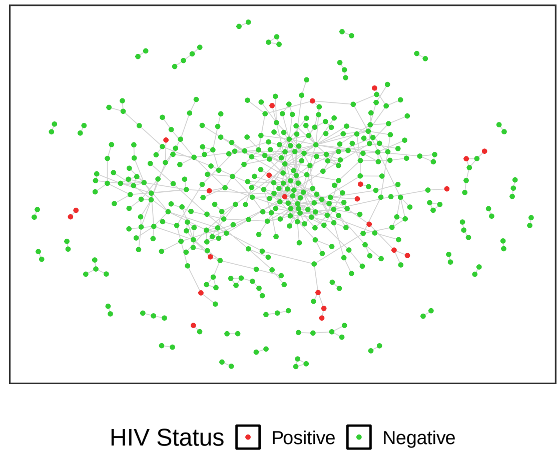

Figure 6: Social Network with HIV Status

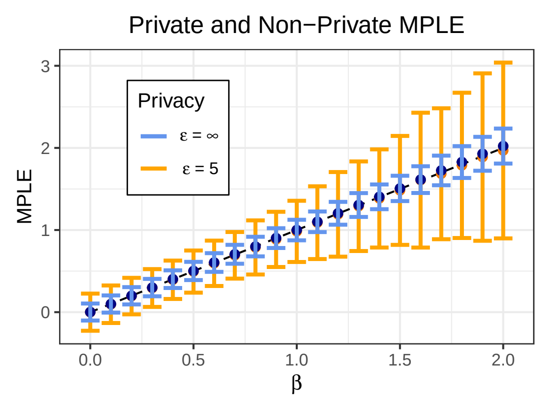

Figure 7: Performance of private and non-private estimators based on Ising model synthetic data on real HIV status network.

the Monte-Carlo conditional expectation of $\mathbb{E}[(\hat{\beta}^{\texttt{priv}} - \hat{\beta}_n)^2 | \boldsymbol{\sigma}, \boldsymbol{J}_n]$ to quantify the cost of privacy. It can be seen from Figure 8 that the cost of privacy has a decreasing trend over ε, which is as expected.

4.2.2 Political leaning of online blogs

Next we consider the network of popular political blogs over the period of two months preceding the U.S. Presidential Election of 2004 (Adamic and Glance (2005)) and their political leaning. The graph, plotted in 9, have nodes representing the blogs, color coded by their political leaning, and edges between two nodes if and only if at least one of them link to the other. We have removed nodes with very high degrees ($\geqslant 50$) as they are very popular blogs anyway (like `blogforamerica.com`, `churchofcriticalthinking.com`, `brilliantatbreakfast.blogspot.com`, `busybusybusy.com`,etc.) and are outliers in mea-

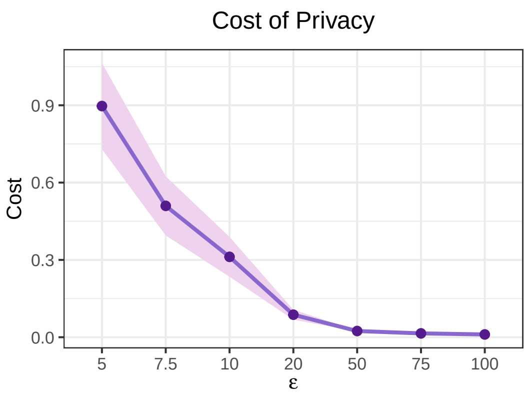

Figure 8: Cost of privacy across ε on the HIV status network and real data

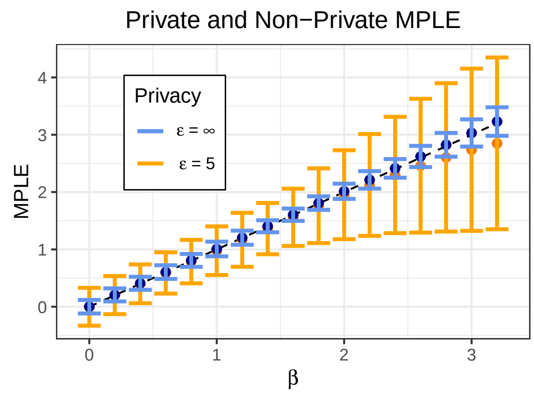

Figure 10: Performance of private and non-private estimators based on Ising model synthetic data on real political blogs network.

suring the influence of the linking network on the political leaning. Any isolated node have also been removed to create a connected graph. The removed nodes are relatively balanced on both sides of the political spectrum. The resulting network, of $n = 815$ nodes is relatively balanced in the two outcomes, liberal (382) and conservative (433), and we want to maintain the privacy of the political leanings of the blogs while measuring the influence of the links on a blog being red or blue.

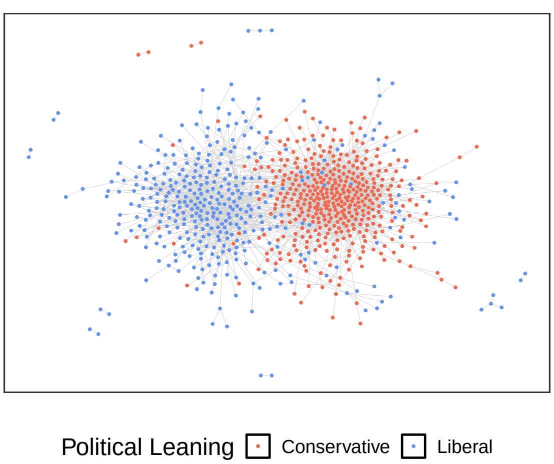

Figure 9: Link Network between Political Blogs

As before we conduct synthetic experiments simulating Ising model on this network, and the results in Figure 10 show how the estimates, both private and non-private concentrate around the true β.

In the real data $\hat{\beta}_n = 2.85$ here, and as before we conduct the cost of privacy analysis as in Section 4.2.1. The results plotted in Figure 11 show the expected downward trend of MSE for rising ε.

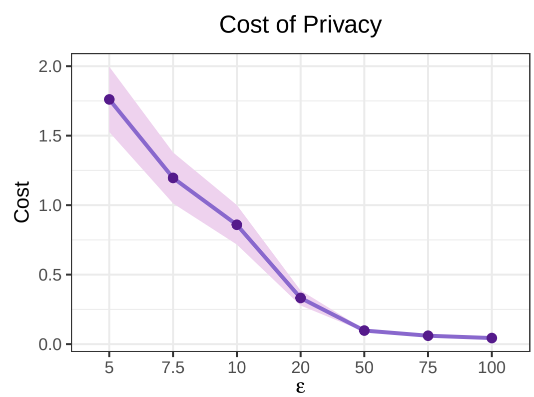

Figure 11: Cost of privacy across ε on the political blogs network and real data.

5 DISCUSSION

The Ising model, initially developed for ferromagnetism, has wide applications in various fields, including social sciences and healthcare, for modeling outcome dependencies in networked systems. However, its use in contexts like disease transmission, tax evasion, and social behavior raises privacy concerns.

Current privacy research primarily focuses on independent models, whereas network analysis mainly employs edge and node differential privacy. However, the Ising model presents a unique challenge of protecting interdependent node outcomes alongside network structure, which current literature does not adequately address.

To address this gap, we introduced an (ε, δ)-differentially private algorithm for Ising models, ensuring node outcome privacy with a single network realization. Our work contributes theoretical insights, including the consistency of the maximum pseudo-likelihood estimator and quantifying privacy cost as

$O\left(\frac{\lambda_n}{a_n \varepsilon}\right)$.

Our experiments demonstrate the algorithm's practicality, preserving privacy and utility in synthetic and real-world networks, like HIV status in a social network and political leaning of online blogs.

Despite our contributions, limitations remain, particularly assumptions related to log-partition functions (see Dagan et al. (2021)).Possible avenues for exploration may involve integrating network privacy and node outcome protection while evaluating privacy guarantees for different privacy notions, such as Renyi Differential Privacy or f-DP (see Dong et al. (2022)).

In summary, our research emphasizes the importance of privacy in Ising models. Our (ε, δ)-differentially private algorithm addresses this concern effectively, offering valuable contributions to this critical field. We hope this work encourages further exploration of privacy preservation techniques in Ising models, promoting a more secure approach to network analysis.

Acknowledgements

The authors would like to thank Bhaswar B. Bhattacharya for helpful discussions and Samuel R Friedman for pointing us to HIV Transmission Network Metastudy Project. They also convey their gratitude to the anonymous referees for their valuable and useful comments.

References

Abawajy, J. H., Ninggal, M. I. H., and Herawan, T. (2016). Privacy preserving social network data publication. *IEEE communications surveys & tutorials*, 18(3):1974–1997.

Acharya, J., Sun, Z., and Zhang, H. (2021). Differentially private assouad, fano, and le cam. In *Algorithmic Learning Theory*, pages 48–78. PMLR.

Adamic, L. A. and Glance, N. (2005). The political blogosphere and the 2004 us election: divided they blog. In *Proceedings of the 3rd international workshop on Link discovery*, pages 36–43.

Backstrom, L., Dwork, C., and Kleinberg, J. (2007). Wherefore art thou r3579x? anonymized social networks, hidden patterns, and structural steganography. In *Proceedings of the 16th international conference on World Wide Web*, pages 181–190.

Bhattacharya, B. B. and Mukherjee, S. (2018). Inference in ising models. *Bernoulli*, 24(1):493–525.

Blocki, J., Blum, A., Datta, A., and Sheffet, O. (2013). Differentially private data analysis of social networks via restricted sensitivity. In *Proceedings of the 4th conference on Innovations in Theoretical Computer Science*, pages 87–96.

Cai, T. T., Wang, Y., and Zhang, L. (2021). The cost of privacy: Optimal rates of convergence for parameter estimation with differential privacy. *The Annals of Statistics*, 49(5):2825–2850.

Cajueiro, D. (2011). Enforcing social behavior in an ising model with complex neighborhoods. *Physica A: Statistical Mechanics and its Applications*, 390(9):1695–1703.

Chatterjee, S. (2007). Estimation in spin glasses: A first step. *The Annals of Statistics*, 35(5):1931–1946.

Chen, H., Cohen-Addad, V., d'Orsi, T., Epasto, A., Imola, J., Steurer, D., and Tiegel, S. (2023). Private estimation algorithms for stochastic block models and mixture models. *arXiv preprint arXiv:2301.04822*.

Chung, F. R. (1997). *Spectral graph theory*, volume 92. American Mathematical Soc.

Dagan, Y., Daskalakis, C., Dikkala, N., and Kandiros, A. V. (2021). Learning ising models from one or multiple samples. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 161–168.

Daskalakis, C., Dikkala, N., and Kamath, G. (2019). Testing ising models. *IEEE Transactions on Information Theory*, 65(11):6829–6852.

Daskalakis, C., Dikkala, N., and Panageas, I. (2020). Logistic regression with peer-group effects via inference in higher-order ising models. In *International Conference on Artificial Intelligence and Statistics*, pages 3653–3663. PMLR.

Dembo, A. and Montanari, A. (2010). Ising models on locally tree-like graphs.

Dong, J., Roth, A., and Su, W. J. (2022). Gaussian differential privacy. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 84(1):3–37.

Duflo, E. and Saez, E. (2002). Participation and investment decisions in a retirement plan: The influence of colleagues' choices. *Journal of public Economics*, 85(1):121–148.

Dwork, C. (2006). Differential privacy. In *International colloquium on automata, languages, and programming*, pages 1–12. Springer.

Dwork, C., Kenthapadi, K., McSherry, F., Mironov, I., and Naor, M. (2006a). Our data, ourselves: Privacy via distributed noise generation. In *Advances in Cryptology-EUROCRYPT 2006: 24th Annual International Conference on the Theory and Applications of Cryptographic Techniques, St. Petersburg,* *Russia, May 28-June 1, 2006. Proceedings 25*, pages 486–503. Springer.

Dwork, C., McSherry, F., Nissim, K., and Smith, A. (2006b). Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography: Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006. Proceedings 3*, pages 265–284. Springer.

Dwork, C., Roth, A., et al. (2014). The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407.

Fossett, M. (2006). Ethnic preferences, social distance dynamics, and residential segregation: Theoretical explorations using simulation analysis*. *Journal of Mathematical Sociology*, 30(3-4):185–273.

Ghosal, P. and Mukherjee, S. (2020). Joint estimation of parameters in ising model. *The Annals of Statistics*, 48(2):785–810.

Giraldo-Barreto, J. and Restrepo, J. (2021). Tax evasion study in a society realized as a diluted ising model with competing interactions. *Physica A: Statistical Mechanics and its Applications*, 582:126264.

Glaeser, E. L., Sacerdote, B., and Scheinkman, J. A. (1996). Crime and social interactions. *The Quarterly journal of economics*, 111(2):507–548.

Kasiviswanathan, S. P., Nissim, K., Raskhodnikova, S., and Smith, A. (2013). Analyzing graphs with node differential privacy. In *Theory of Cryptography Conference*, pages 457–476. Springer.

Kifer, D., Smith, A., and Thakurta, A. (2012). Private convex empirical risk minimization and high-dimensional regression. In *Conference on Learning Theory*, pages 25–1. JMLR Workshop and Conference Proceedings.

Kleinberg, J. M. (2007). Challenges in mining social network data: processes, privacy, and paradoxes. In *Proceedings of the 13th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 4–5.

Krauth, B. (2006). Social interactions in small groups. *Canadian journal of Economics*, pages 414–433.

Levin, D. A., Luczak, M. J., and Peres, Y. (2010). Glauber dynamics for the mean-field ising model: cut-off, critical power law, and metastability. *Probability Theory and Related Fields*, 146:223–265.

Li, Q., Han, Z., and Wu, X.-M. (2018). Deeper insights into graph convolutional networks for semi-supervised learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32.

Liu, C., Chakraborty, S., and Mittal, P. (2016). Dependence makes you vulnberable: Differential privacy under dependent tuples. In *NDSS*, volume 16, pages 21–24.

Mello, I. F., Squillante, L., Gomes, G. O., Seridonio, A. C., and de Souza, M. (2021). Epidemics, the ising-model and percolation theory: a comprehensive review focused on covid-19. *Physica A: Statistical Mechanics and its Applications*, 573:125963.

Mohamed, M. S., Nguyen, D., Vullikanti, A., and Tandon, R. (2022). Differentially private community detection for stochastic block models. In *International Conference on Machine Learning*, pages 15858–15894. PMLR.

Morris, M. and Rothenberg, R. (2011). Hiv transmission network metastudy project: An archive of data from eight network studies, 1988–2001. *ICPSR Data Holdings*.

Mukherjee, R. and Ray, G. (2022). On testing for parameters in Ising models. *Annales de l'Institut Henri Poincaré, Probabilités et Statistiques*, 58(1):164 – 187.

Nettle, D. (1999). Is the rate of linguistic change constant? *Lingua*, 108(2-3):119–136.

Pickhardt, M. and Seibold, G. (2014). Income tax evasion dynamics: Evidence from an agent-based econophysics model. *Journal of Economic Psychology*, 40:147–160.

Potterat, J. J., Phillips-Plummer, L., Muth, S. Q., Rothenberg, R., Woodhouse, D., Maldonado-Long, T., Zimmerman, H., and Muth, J. (2002). Risk network structure in the early epidemic phase of hiv transmission in colorado springs. *Sexually transmitted infections*, 78(suppl 1):i159–i163.

Schelling, T. C. (1971). Dynamic models of segregation. *Journal of mathematical sociology*, 1(2):143–186.

Srivastava, J., Ahmad, M. A., Pathak, N., and Hsu, D. K.-W. (2008). Data mining based social network analysis from online behavior. In *Tutorial at the 8th SIAM international conference on data mining (SDM'08)*.

Stauffer, D. (2008). Social applications of two-dimensional ising models. *American Journal of Physics*, 76(4):470–473.

Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., and Philip, S. Y. (2020). A comprehensive survey on graph neural networks. *IEEE transactions on neural networks and learning systems*, 32(1):4–24.

Zaklan, G., Westerhoff, F., and Stauffer, D. (2009). Analysing tax evasion dynamics via the ising model.

Journal of Economic Interaction and Coordination, 4(1):1–14.

Zhang, H., Kamath, G., Kulkarni, J., and Wu, S. (2020). Privately learning markov random fields. In *International conference on machine learning*, pages 11129–11140. PMLR.

Zhou, B., Pei, J., and Luk, W. (2008). A brief survey on anonymization techniques for privacy preserving publishing of social network data. *ACM Sigkdd Explorations Newsletter*, 10(2):12–22.

Zhou, J., Cui, G., Hu, S., Zhang, Z., Yang, C., Liu, Z., Wang, L., Li, C., and Sun, M. (2020). Graph neural networks: A review of methods and applications. *AI open*, 1:57–81.

Checklist

1. For all models and algorithms presented, check if you include:
   (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. [Yes]
   (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. [Yes]
   (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [Yes]

2. For any theoretical claim, check if you include:
   (a) Statements of the full set of assumptions of all theoretical results. [Yes]
   (b) Complete proofs of all theoretical results. [Yes]
   (c) Clear explanations of any assumptions. [Yes]

3. For all figures and tables that present empirical results, check if you include:
   (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). [Yes]
   (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [Yes]
   (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Yes]
   (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [Not Applicable]

4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
   (a) Citations of the creator If your work uses existing assets. [Yes]
   (b) The license information of the assets, if applicable. [Not Applicable]
   (c) New assets either in the supplemental material or as a URL, if applicable. [Yes]
   (d) Information about consent from data providers/curators. [Not Applicable]
   (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [Yes]

5. If you used crowdsourcing or conducted research with human subjects, check if you include:
   (a) The full text of instructions given to participants and screenshots. [Not Applicable]
   (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [Not Applicable]
   (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [Not Applicable]

PrIsing: Supplementary Materials

A Proof of Theorem 3.1

In this section we prove that Algorithm 1 satisfies (ε, δ) privacy for any $\varepsilon > 0$ and $\delta \geqslant 0$. The proof is organised as follows. First, in the following lemma, we bound the amount of change in $L_{\boldsymbol{\sigma}}(\beta)$ induced by flipping a coordinate in $\boldsymbol{\sigma}$.

Lemma A.1. *Fix $1 \leqslant j \leqslant n$. If $\boldsymbol{\sigma}$ and $\boldsymbol{\sigma'}$ differ in only the j-th entry, then for any $\beta > 0$,*

$$|L_{\boldsymbol{\sigma}}(\beta) - L_{\boldsymbol{\sigma'}}(\beta)| \leqslant 8\frac{d_j}{n^2}.$$

The proof of Lemma A.1 is provided in Section A.1. Now consider $\boldsymbol{\sigma}, \boldsymbol{\sigma'} \in \{-1, 1\}^n$ such that,

$$\sum_{i=1}^{n} \mathbb{1}\left\{\sigma_i \neq \sigma_i'\right\} = 1. \tag{8}$$

Fix $\alpha > 0$. If $\hat{\beta}^{\texttt{priv}} = \alpha$, then α must satisfy,

$$L_{\boldsymbol{\sigma}}(\alpha) + \Delta\alpha/n + b/n = 0$$

or equivalently, define

$$b(\alpha; \boldsymbol{\sigma}) := -(nL_{\boldsymbol{\sigma}}(\alpha) + \Delta\alpha). \tag{9}$$

Using a change of variable approach the ratio of densities can be written as,

$$\frac{f_{\widehat{\beta}^{\texttt{priv}}, \boldsymbol{\sigma}}(\alpha)}{f_{\widehat{\beta}^{\texttt{priv}}, \boldsymbol{\sigma'}}(\alpha)} = \frac{\nu(b(\alpha, \boldsymbol{\sigma}); \varepsilon, \delta)}{\nu(b(\alpha, \boldsymbol{\sigma'}); \varepsilon, \delta)} \cdot \frac{|\nabla b(\alpha; \boldsymbol{\sigma'})|}{|\nabla b(\alpha; \boldsymbol{\sigma})|}, \tag{10}$$

where ∇ denotes the partial derivative with respect to α and $f_{\widehat{\beta}^{\texttt{priv}}, \boldsymbol{\tau}}$ denotes the density of $\widehat{\beta}^{\texttt{priv}}$ given the data $\boldsymbol{\tau} = \boldsymbol{\sigma}, \boldsymbol{\sigma'}$. Now in the subsequent lemmas we bound the two ratios appearing in R.H.S of (10) separately. First, in the following result, we bound the second term.

Lemma A.2. *For any $\boldsymbol{\sigma}, \boldsymbol{\sigma'} \in \{-1, 1\}^n$ satisfying* (8),

$$\left|\frac{\nabla b(\alpha; \boldsymbol{\sigma})}{\nabla b(\alpha; \boldsymbol{\sigma'})}\right| \leqslant e^{\frac{\varepsilon}{2}}$$

where $b(\alpha, \cdot)$ is defined in (9).

Next, we provide a bound on the ratio of densities in the following lemma.

Lemma A.3. *Consider any $\boldsymbol{\sigma}, \boldsymbol{\sigma'} \in \{-1,1\}^n$ satisfying (8). Then using Algorithm 1 for (ε, δ) privacy with $0 < \delta < \frac{2}{\sqrt{e}}$, we get,*

$$\frac{\nu(b(\alpha, \boldsymbol{\sigma}); \varepsilon, \delta)}{\nu(b(\alpha, \boldsymbol{\sigma'}); \varepsilon, \delta)} \leqslant e^{\varepsilon/2}$$

on a set $S \subseteq \mathbb{R}$ such that $\mathbb{P}\left(b(\alpha, \boldsymbol{\sigma}) \in S\right) \geqslant 1 - \delta$, and for $(\varepsilon, 0)$ privacy we get,

$$\frac{\nu(b(\alpha, \boldsymbol{\sigma}); \varepsilon, 0)}{\nu(b(\alpha, \boldsymbol{\sigma'}); \varepsilon, 0)} \leqslant e^{\varepsilon/2}.$$

The proofs of Lemma A.2 and Lemma A.3 are given in Sections A.2 and A.3 respectively. We now proceed to show that Algorithm 1 preserves the notion of (ε, δ) differential privacy as defined in Definition 2.1. The proof of $(\varepsilon, 0)$ differential privacy is now immidiate by combining the bounds from Lemma A.2, Lemma A.3 and (10). For $\delta > 0$, recalling S from Lemma A.3 observe that,

$$\begin{aligned} f_{\widehat{\beta}^{\mathtt{priv}}, \boldsymbol{\sigma}}(\alpha) &\leqslant e^{\varepsilon/2} f_{\widehat{\beta}^{\mathtt{priv}}, \boldsymbol{\sigma'}}(\alpha) \mathbb{1}\left\{b(\alpha, \boldsymbol{\sigma}) \in S\right\} + f_{\widehat{\beta}^{\mathtt{priv}}, \boldsymbol{\sigma}}(\alpha) \mathbb{1}\left\{b(\alpha, \boldsymbol{\sigma}) \in S^c\right\} \\ &\leqslant e^{\varepsilon/2} f_{\widehat{\beta}^{\mathtt{priv}}, \boldsymbol{\sigma'}}(\alpha) + f_{\widehat{\beta}^{\mathtt{priv}}, \boldsymbol{\sigma}}(\alpha) \mathbb{1}\left\{b(\alpha, \boldsymbol{\sigma}) \in S^c\right\} \end{aligned}$$

Then for any borel set $A \subseteq \mathbb{R}$ and using a change of variable we get,

$$\begin{aligned} \int_A f_{\widehat{\beta}^{\mathtt{priv}}, \boldsymbol{\sigma}}(\alpha) \mathrm{d}\alpha &\leqslant e^{\varepsilon/2} \int_A f_{\widehat{\beta}^{\mathtt{priv}}, \boldsymbol{\sigma'}}(\alpha) \mathrm{d}\alpha + \int f_{\widehat{\beta}^{\mathtt{priv}}, \boldsymbol{\sigma}}(\alpha) \mathbb{1}\left\{b(\alpha, \boldsymbol{\sigma}) \in S^c\right\} \mathrm{d}\alpha \\ &\leqslant e^{\varepsilon/2} \int_A f_{\widehat{\beta}^{\mathtt{priv}}, \boldsymbol{\sigma'}}(\alpha) \mathrm{d}\alpha + \int \nu\left(b(\alpha, \boldsymbol{\sigma}); \varepsilon, \delta\right) \mathbb{1}\left\{b(\alpha, \boldsymbol{\sigma}) \in S^c\right\} \mathrm{d}b(\alpha, \boldsymbol{\sigma}) \\ &\leqslant e^{\varepsilon/2} \int_A f_{\widehat{\beta}^{\mathtt{priv}}, \boldsymbol{\sigma'}}(\alpha) \mathrm{d}\alpha + \delta \end{aligned}$$

where the last bound follows by definiton of S from Lemma A.3. The proof is now completed by recalling the choice of $\boldsymbol{\sigma}$ and $\boldsymbol{\sigma'}$ from (8).

A.1 Proof of Lemma A.1

Recalling (5), note that $m_i(\boldsymbol{\sigma})$ does not depend on σ_i for all $1 \leqslant i \leqslant n$. Using (4), we have

$$L_{\boldsymbol{\tau}}(\beta) = -\frac{1}{n} \sum_{i=1}^{n} m_i(\boldsymbol{\tau}) \tau_i + \frac{1}{n} \sum_{i=1}^{n} m_i(\boldsymbol{\tau}) \tanh(\beta m_i(\boldsymbol{\tau}))$$

for $\boldsymbol{\tau} = \boldsymbol{\sigma}, \boldsymbol{\sigma'}$. Then,

$$\begin{aligned} L_{\boldsymbol{\sigma}}(\beta) - L_{\boldsymbol{\sigma'}}(\beta) = &-\frac{1}{n} \sum_{i: i \neq j} [m_i(\boldsymbol{\sigma}) - m_i(\boldsymbol{\sigma'})] \sigma_i - \frac{1}{n} m_j(\boldsymbol{\sigma})(\sigma_j - \sigma'_j) \\ &- \frac{1}{n} \sum_{i=1}^{n} [m_i(\boldsymbol{\sigma'}) \tanh(\beta m_i(\boldsymbol{\sigma'})) - m_i(\boldsymbol{\sigma}) \tanh(\beta m_i(\boldsymbol{\sigma}))] \end{aligned} \tag{11}$$

Now recalling that all entries of $\boldsymbol{J}_n$ are non-negative,

$$|m_i(\boldsymbol{\sigma}) - m_i(\boldsymbol{\sigma'})| = |\sum_{k=1}^{n} (\sigma_k - \sigma'_k) \boldsymbol{J}_n(i,k)| = |(\sigma_j - \sigma'_j) \boldsymbol{J}_n(i,j)| \leqslant 2 \boldsymbol{J}_n(i,j). \tag{12}$$

Consider,

$$\kappa(x) := x \tanh(\beta x), \ \forall x \in \mathbb{R}.$$

It is easy to see that,

$$\kappa^{'}(x) = \tanh(\beta x) + x \beta \mathrm{sech}^2(\beta x), \ \forall x \in \mathbb{R},$$

and by definition $|\kappa'| \leqslant 2$. Then using the Mean Value Theorem we conclude,

$$|\kappa(x) - \kappa(y)| \leqslant 2|x-y| \text{ for all } x, y \in \mathbb{R}.$$

Now recalling the definition of κ and (12) we get,

$$|m_i(\boldsymbol{\sigma'}) \tanh(\beta m_i(\boldsymbol{\sigma'})) - m_i(\boldsymbol{\sigma}) \tanh(\beta m_i(\boldsymbol{\sigma}))| = |\kappa(m_i(\boldsymbol{\sigma})) - \kappa(m_i(\boldsymbol{\sigma'}))| \leqslant 4\boldsymbol{J}_n(i,j) \tag{13}$$

Thus combining (11),(13) and noticing that $|m_j(\boldsymbol{\sigma})| \leqslant \sum_{k=1}^{n} \boldsymbol{J}_n(j,k) \leqslant d_j/n$, we have,

$$|L_{\boldsymbol{\sigma}}(\beta) - L_{\boldsymbol{\sigma'}}(\beta)| \leqslant \frac{2}{n} \sum_{i=1}^{n} \boldsymbol{J}_n(i,j) + \frac{2}{n}|m_j(\boldsymbol{\sigma})| + \frac{4}{n} \sum_{i=1}^{n} \boldsymbol{J}_n(i,j) \leqslant 2\frac{d_j}{n^2} + 2\frac{d_j}{n^2} + 4\frac{d_j}{n^2} = 8\frac{d_j}{n^2},$$

completing the proof of the lemma.

A.2 Proof of Lemma A.2

Since $\boldsymbol{\sigma}$ and $\boldsymbol{\sigma'}$ satisfy (8), then there exists $1 \leqslant j \leqslant n$ such that $\sigma_j \neq \sigma'_j$. Note that,

$$\left|\frac{\nabla b(\alpha, \boldsymbol{\sigma})}{\nabla b(\alpha, \boldsymbol{\sigma'})}\right| \leqslant 1 + \left|\frac{\nabla b(\alpha, \boldsymbol{\sigma}) - \nabla b(\alpha, \boldsymbol{\sigma'})}{\nabla b(\alpha, \boldsymbol{\sigma'})}\right|. \tag{14}$$

Once again by (5), note that $m_i(\boldsymbol{\sigma})$ does not depend on $\boldsymbol{\sigma}_i$ for all $1 \leqslant i \leqslant n$. Now recalling (9) and taking derivative on both sides of (11) shows,

$$\begin{aligned}
|\nabla b(\alpha, \boldsymbol{\sigma}) - \nabla b(\alpha, \boldsymbol{\sigma'})| &= \left|\sum_{i:i\neq j} m_i(\boldsymbol{\sigma})^2 \text{sech}^2\left(\alpha m_i(\boldsymbol{\sigma})\right) - m_i(\boldsymbol{\sigma'})^2 \text{sech}^2\left(\alpha m_i(\boldsymbol{\sigma'})\right)\right| \\
&\leqslant \sum_{i:i\neq j} \left|m_i(\boldsymbol{\sigma})^2 \text{sech}^2\left(\alpha m_i(\boldsymbol{\sigma})\right) - m_i(\boldsymbol{\sigma'})^2 \text{sech}^2\left(\alpha m_i(\boldsymbol{\sigma'})\right)\right| \\
&\leqslant \frac{2}{n} \sum_{i:i\neq j} d_i \left|m_i(\boldsymbol{\sigma}) \text{sech}(\alpha m_i(\boldsymbol{\sigma})) - m_i(\boldsymbol{\sigma'}) \text{sech}(\alpha m_i(\boldsymbol{\sigma'}))\right| \qquad (15)
\end{aligned}$$

where the inequality in (15) follows from the bounds $|m_i(\boldsymbol{\sigma})| \leqslant d_i/n$ and $|\text{sech}(\cdot)| \leqslant 1$. Define,

$$\kappa_0(x) := x\text{sech}(\alpha x) \ \forall x \in \mathbb{R}.$$

Observe that,

$$\kappa_0'(x) = \text{sech}(\alpha x) - x\alpha\text{sech}(\alpha x)\tanh(\alpha x) \ \forall x \in \mathbb{R}.$$

Now it is easy to infer that $|\kappa_0'(\cdot)| \leqslant 3$. Using Mean value theorem we get,

$$|\kappa_0(x) - \kappa_0(y)| \leqslant 3|x-y| \ \forall x, y \in \mathbb{R}. \tag{16}$$

Finally by the definition of κ_0, (16), (12) and recalling that entries of $\boldsymbol{J}_n$ are non-negative we conclude,

$$\left|m_i(\boldsymbol{\sigma}) \text{sech}(\alpha m_i(\boldsymbol{\sigma})) - m_i(\boldsymbol{\sigma'}) \text{sech}(\alpha m_i(\boldsymbol{\sigma'}))\right| \leqslant 6\boldsymbol{J}_n(i,j), \ \forall i \neq j. \tag{17}$$

Next, note that

$$|\nabla b(\alpha, \boldsymbol{\sigma'})| = |\Delta + \sum_{i=1}^{n} m_i(\boldsymbol{\sigma'})^2 \text{sech}^2(\alpha m_i(\boldsymbol{\sigma'}))| \geqslant \Delta$$

Thus recalling (14), (15) and (17) shows,

$$\left|\frac{\nabla b(\alpha, \boldsymbol{\sigma})}{\nabla b(\alpha, \boldsymbol{\sigma'})}\right| \leqslant 1 + \frac{\frac{12}{n}\sum_{i:i\neq j} d_i \boldsymbol{J}_n(i,j)}{\Delta} \leqslant 1 + \frac{\varepsilon}{2} \leqslant e^{\varepsilon/2}$$

completing the proof.

A.3 Proof of Lemma A.3

First suppose that we are using Algorithm 1 for (ε, δ) privacy. Let $\Gamma = b(\alpha, \boldsymbol{\sigma}) - b(\alpha, \boldsymbol{\sigma'})$. By Lemma A.1, $|\Gamma| \leqslant \zeta = 8 \max_j \frac{d_j}{n}$. Then,

$$\begin{aligned}\frac{\nu(b(\alpha,\boldsymbol{\sigma});\varepsilon,\delta)}{\nu(b(\alpha,\boldsymbol{\sigma'});\varepsilon,\delta)} &= \exp\left(\frac{1}{2\gamma^2}(b(\alpha,\boldsymbol{\sigma'})^2 - b(\alpha,\boldsymbol{\sigma})^2)\right) = \exp\left(\frac{1}{2\gamma^2}((b(\alpha,\boldsymbol{\sigma}) - \Gamma)^2 - b(\alpha,\boldsymbol{\sigma})^2)\right) \\ &= \exp\left(\frac{1}{2\gamma^2}(-2b(\alpha,\boldsymbol{\sigma})\Gamma) + \Gamma^2)\right) \\ &\leqslant \exp\left(\frac{1}{2\gamma^2}|2b(\alpha,\boldsymbol{\sigma})|\zeta + \zeta^2\right) \qquad (18)\end{aligned}$$

Note that for any random variable $Z \sim \mathcal{N}(0,1)$

$$\mathbb{P}(|Z| > t) \leqslant 2e^{-t^2/2}, \text{ for } t > 1.$$

Hence for, $b(\alpha, \boldsymbol{\sigma}) \sim \mathcal{N}(0, \gamma^2)$,

$$\mathbb{P}(|b(\alpha,\boldsymbol{\sigma})| \geqslant \gamma t) \leqslant 2e^{-t^2/2}, \text{ for any } t > 1.$$

Let $S_t := \{a \in \mathbb{R} : |a| \geqslant \gamma t\}$. Then it is easy to observe that for $\delta < \frac{2}{\sqrt{e}}$, and choosing $t_0 = \sqrt{2\log(2/\delta)}$ we get,

$$\mathbb{P}\left(b(\alpha,\boldsymbol{\sigma}) \in S_{t_0}\right) \leqslant \delta$$

Thus on the set $S := S_{t_0}^c$, using (18) and recalling the definition of γ from Algorithm 1 we find,

$$\frac{\nu(b(\alpha,\boldsymbol{\sigma});\varepsilon,\delta)}{\nu(b(\alpha,\boldsymbol{\sigma'});\varepsilon,\delta)} \leqslant \exp\left(\frac{1}{2\gamma^2}\left\{\gamma\zeta\sqrt{8\log\frac{2}{\delta}} + \zeta^2\right\}\right) \leqslant e^{\varepsilon/2}$$

which completes the proof of Lemma A.3 for (ε, δ) privacy. Now suppose we are using Algorithm 1 for $(\varepsilon, 0)$ privacy. Then by definition,

$$\frac{\nu(b(\alpha,\boldsymbol{\sigma});\varepsilon,0)}{\nu(b(\alpha,\boldsymbol{\sigma'});\varepsilon,0)} = \exp\left(\frac{\varepsilon}{2\zeta}\left(|b(\alpha,\boldsymbol{\sigma'})| - |b(\alpha,\boldsymbol{\sigma})|\right)\right) \leqslant \exp\left(\frac{\varepsilon}{2\zeta}\left|b(\alpha,\boldsymbol{\sigma'}) - b(\alpha,\boldsymbol{\sigma})\right|\right) \leqslant \exp\left(\frac{\varepsilon}{2}\right)$$

where the upper bound once again follows from Lemma A.1.

B Error Bound of `PrIsing` Algorithm

In this section we embark on a careful analysis of the `PrIsing` Algorithm and provide a detailed error bound on the performance of the same. The performance of the non-private MPLE was analysed Theorem 2.1 from Bhattacharya and Mukherjee (2018), where the authors concluded that the estimator is $\sqrt{a_n}$ consistent, where, under certain regularity conditions on the log-partition function, a_n is the Frobenius norm of the matrix $\boldsymbol{J}_n$. In the following result we recall the sufficient conditions for consistency of MPLE from Bhattacharya and Mukherjee (2018), and analyze the performance of `PrIsing` under the same.

Theorem B.1. *Let $\sup_{n \geqslant 1} \|\boldsymbol{J}_n\| < \infty$, and let $\beta_0 > 0$ be fixed. Suppose $\{a_n : n \geqslant 1\}$ is a sequence such that,*

$$a_n \overset{n\to\infty}{\longrightarrow} \infty$$

and for some $\vartheta > 0$,

$$0 < \liminf_{n\to\infty} \frac{1}{a_n} F_n(\beta_0 - \vartheta) \leqslant \limsup_{n\to\infty} \frac{1}{a_n} F_n(\beta_0 + \vartheta) < \infty.$$

Further assume that,

(i) $u_{n,K} := \mathbb{E}_{\beta_0}\left[\sum_{i=1}^n |m_i(\boldsymbol{\sigma})| \mathbb{1}\{|m_i(\boldsymbol{\sigma})| > K\}\right]$ is such that $\limsup_{K\to\infty} \limsup_{n\to\infty} \frac{1}{a_n} u_{n,K} = 0$, and

(ii) $\limsup_{n\to\infty} \frac{1}{a_n}\sum_{i,j=1}^{n} \boldsymbol{J}_n(i,j)^2 < \infty$

Then the private MPLE estimator $\hat{\beta}^{priv}$ from Algorithm 1 satisfies,

$$|\widehat{\beta}^{priv} - \beta_0| = O_p\left(\frac{1}{\sqrt{a_n}} + \frac{\sqrt{8\zeta^2\rho_{\varepsilon,\delta} + \varepsilon^2\Delta^2\beta_0^2}}{a_n\varepsilon}\right),$$

where

$$\rho_{\varepsilon,\delta} = \begin{cases} \log(2/\delta) + \varepsilon/2 & \text{if } \delta > 0 \\ 1 & \text{if } \delta = 0 \end{cases}$$

Proof. We prove Theorem B.1 by following the techniques developed in Bhattacharya and Mukherjee (2018). For notational convinience define,

$$k_{n,\delta} := \frac{n^2\varepsilon^2}{\zeta^2\left(8\log(2/\delta) + 4\varepsilon\right) + \varepsilon^2\Delta^2\beta_0^2} \text{ for all } \delta > 0,\ k_{n,0} := \frac{n^2\varepsilon^2}{8\zeta^2 + \varepsilon^2\Delta^2\beta_0^2}. \tag{19}$$

and consider,

$$s_{n,\delta}^2 := \frac{n^2 k_{n,\delta}}{\left(\sqrt{a_n k_{n,\delta}} + n\right)^2} = \begin{cases} \dfrac{1}{\frac{1}{\sqrt{a_n}} + \frac{\sqrt{\zeta^2(8\log(2/\delta)+4\varepsilon)+\varepsilon^2\Delta^2\beta_0^2}}{a_n\varepsilon}} & \text{if } \delta > 0 \\ \dfrac{1}{\frac{1}{\sqrt{a_n}} + \frac{\sqrt{8\zeta^2+\varepsilon^2\Delta^2\beta_0^2}}{a_n\varepsilon}} & \text{if } \delta = 0 \end{cases}.$$

Further we will use $(\varepsilon, 0)$ privacy in place of ε-privacy for consistency in notation and $\lesssim_\theta$ to denote less than equality upto a constant depending on a parameter θ.
By definition it is easy to observe that,

$$s_{n,\delta}^2 \leqslant \min\left\{\frac{n^2}{a_n}, k_{n,\delta}.\right\} \tag{20}$$

Recall that by Algorithm 1, $\widehat{\beta}^{\mathtt{priv}}$ is the solution to the equation,

$$\mathcal{L}_{\boldsymbol{\sigma}}(\beta, b) := L_{\boldsymbol{\sigma}}(\beta) + \frac{\Delta}{n}\beta + \frac{b}{n} = 0.$$

Then for $\delta > 0$ with (ε, δ) privacy,

$$\begin{aligned} \limsup_{n\to\infty} s_{n,\delta}^2 \mathbb{E}_{\beta_0, b\sim N(0,\gamma^2)}\left[\mathcal{L}_{\boldsymbol{\sigma}}(\beta_0, b)^2\right] &\lesssim \limsup_{n\to\infty} s_{n,\delta}^2 \mathbb{E}_{\beta_0} L_{\boldsymbol{\sigma}}(\beta_0)^2 + \frac{s_{n,\delta}^2}{n^2}\Delta^2\beta_0^2 + \frac{s_{n,\delta}^2}{n^2}\mathbb{E}_{\mathrm{N}(0,\gamma^2)} b^2 \\ &\leqslant \limsup_{n\to\infty} \frac{n^2}{a_n}\mathbb{E}_{\beta_0} L_{\boldsymbol{\sigma}}(\beta_0)^2 + 1 < \infty. \end{aligned} \tag{21}$$

where the finiteness follows by (Bhattacharya and Mukherjee, 2018, Lemma 5.2), the definition of $k_{n,\delta}$ from (19), γ, Δ from Algorithm 1 and observing that,

$$\frac{s_{n,\delta}^2}{n^2}\Delta^2\beta_0^2 + \frac{s_{n,\delta}^2}{n^2}\gamma^2 \leqslant \frac{k_{n,\delta}}{n^2}\Delta^2\beta_0^2 + \frac{k_{n,\delta}}{n^2}\gamma^2 \leqslant \frac{\varepsilon^2\Delta^2\beta_0^2 + \zeta^2(8\log(2/\delta) + 4\varepsilon)}{\zeta^2\left(8\log(2/\delta) + 4\varepsilon\right) + \varepsilon^2\Delta^2\beta_0^2} \leqslant 1. \tag{22}$$

where the first inequality follows from (20). Note that for $b \sim \mathrm{Lap}(0, 2\zeta/\varepsilon)$, $\mathbb{E}\left[b^2\right] = 8(\zeta/\varepsilon)^2$. Then for $\delta = 0$ by a similar computation as in (21) and (22) we get,

$$\limsup_{n\to\infty} s_{n,\delta}^2 \mathbb{E}_{\beta_0, b\sim \mathrm{Lap}(0,2\zeta/\varepsilon)}\left[\mathcal{L}_{\boldsymbol{\sigma}}(\beta_0, b)^2\right] < \limsup_{n\to\infty} \frac{n^2}{a_n}\mathbb{E}_{\beta_0} L_{\boldsymbol{\sigma}}(\beta_0)^2 + 1 < \infty. \tag{23}$$

Fix $\delta \geqslant 0$ and fix $\xi > 0$, then by Chebyshev inequality, (21) and (23) we can choose $K_1 = K_1(\xi) > 0$ such that,

$$\mathbb{P}\left(|\mathcal{L}_{\boldsymbol{\sigma}}(\beta_0, b)| > \frac{K_1}{s_{n,\delta}}\right) \leqslant \frac{s_{n,\delta}^2}{K_1^2}\mathbb{E}\mathcal{L}_{\boldsymbol{\sigma}}(\beta_0, b)^2 \lesssim_{\beta_0} \frac{1}{K_1^2} < \xi. \tag{24}$$

By (Bhattacharya and Mukherjee, 2018, Lemma 5.3) it is easy to observe that there exists $\nu := \nu(\xi)$ and $K_2 = K_2(\nu, \xi)$ such that,

$$\mathbb{P}_{\beta_0}\left(\sum_{i=1}^n m_i(\boldsymbol{\sigma})^2 \mathbf{1}\left\{|m_i(\boldsymbol{\sigma})| \leqslant K_2\right\} \geqslant \nu a_n\right) \geqslant 1 - \xi \tag{25}$$

for large enough n. Define,

$$T_n := \left\{(\boldsymbol{\sigma}, b) \in \{+1, -1\}^n \times \mathbb{R} : |\mathcal{L}_{\boldsymbol{\sigma}}(\beta_0, b)| \leqslant \frac{K_1}{s_{n,\delta}}, \sum_{i=1}^n m_i(\boldsymbol{\sigma})^2 \mathbf{1}\left\{|m_i(\boldsymbol{\sigma})| \leqslant K_2\right\} \geqslant \nu a_n\right\}.$$

Combining (24) and (25) and taking n large enough we conclude that,

$$\mathbb{P}(T_n) \geqslant 1 - 2\xi.$$

Now choosing $(\boldsymbol{\sigma}, b) \in T_n$ and recalling that the parameter $\beta \geqslant 0$ shows,

$$\begin{aligned}
\mathcal{L}'_{\boldsymbol{\sigma}}(\beta, b) := \frac{\partial}{\partial \beta}\mathcal{L}_{\boldsymbol{\sigma}}(\beta, b) &= \frac{1}{n}\sum_{i=1}^n m_i(\boldsymbol{\sigma})^2 \operatorname{sech}^2(\beta m_i(\boldsymbol{\sigma})) + \frac{\Delta}{n} \\
&\geqslant \frac{1}{n}\operatorname{sech}^2(\beta K_2)\sum_{i}^n m_i(\boldsymbol{\sigma})^2 \mathbf{1}\left\{|m_i(\boldsymbol{\sigma})| \leqslant K_2\right\} + \frac{\Delta}{n} \\
&\geqslant \nu\frac{a_n}{n}\operatorname{sech}^2(\beta K_2) + \frac{\Delta}{n}.
\end{aligned} \tag{26}$$

Thus,

$$\begin{aligned}
\frac{K_1}{s_{n,\delta}} \geqslant |\mathcal{L}_{\boldsymbol{\sigma}}(\beta_0, b)| &= |\mathcal{L}_{\boldsymbol{\sigma}}(\beta_0, b) - \mathcal{L}_{\boldsymbol{\sigma}}(\widehat{\beta}^{\texttt{priv}}, b)| \\
&\geqslant \int_{\widehat{\beta}^{\texttt{priv}} \wedge \beta_0}^{\widehat{\beta}^{\texttt{priv}} \vee \beta_0} \mathcal{L}'_{\sigma}(\beta, b)d\beta \\
&\geqslant \left|\nu\frac{a_n}{K_2 n}\tanh(K_2\widehat{\beta}^{\texttt{priv}}) + \frac{\Delta}{n}\widehat{\beta}^{\texttt{priv}} - \nu\frac{a_n}{K_2 n}\tanh(K_2\beta_0) - \frac{\Delta}{n}\beta_0\right| \\
&= \nu\frac{a_n}{K_2 n}\left|\tanh(K_2\widehat{\beta}^{\texttt{priv}}) + \frac{K_2\Delta}{\nu a_n}\widehat{\beta}^{\texttt{priv}} - \tanh(K_2\beta_0) - \frac{K_2\Delta}{\nu a_n}\beta_0\right|.
\end{aligned} \tag{27}$$

where the inequality in (27) follows from (26). Now recalling our choice of $(\boldsymbol{\sigma}, b) \in T_n$ we conclude,

$$\mathbb{P}\left(\frac{a_n s_{n,\delta}}{n}\left|\tanh(K_2\widehat{\beta}^{\texttt{priv}}) + \frac{K_2\Delta}{\nu a_n}\widehat{\beta}^{\texttt{priv}} - \tanh(K_2\beta_0) - \frac{K_2\Delta}{\nu a_n}\beta_0\right| \geqslant \frac{K_2}{\nu K_1}\right) \leqslant 2\xi$$

for large enough n. The proof is now complete by invoking Lemma D.1.

∎

B.1 Proof of Theorem 3.2

Note that all the assumptions of Theorem B.1 are satisfied. By Algorithm 1 the smallest permitted value of Δ is given by,

$$\Delta_0 = \max_j\left\{\frac{24}{\varepsilon n}\sum_{i=1}^n d_i \boldsymbol{J}_n(i, j)\right\}. \tag{28}$$

Fix $1 \leqslant j \leqslant n$. By the definition of $d_i, 1 \leqslant i \leqslant n$ from Algorithm 1 note that,

$$\sum_{i=1}^{n} d_i \boldsymbol{J}_n(i,j) = n \sum_{i=1}^{n}\sum_{k=1}^{n} \boldsymbol{J}_n(i,k)\boldsymbol{J}_n(i,j) = n \sum_{k=1}^{n}\sum_{i=1}^{n} \boldsymbol{J}_n(k,i)\boldsymbol{J}_n(i,j) = n \sum_{k=1}^{n} \boldsymbol{J}_n^2(k,j).$$

By (28) note that,

$$\varepsilon\Delta_0 = 24 \max_j \left\{ \sum_{k=1}^{n} \boldsymbol{J}_n^2(k,j) \right\} = 24 \left\| \boldsymbol{J}_n^2 \right\|_{1\to\infty}. \tag{29}$$

Now for ζ from Algorithm 1 we have,

$$\zeta = 8 \max_i \left\{ \frac{d_i}{n} \right\} = 8 \max_i \left\{ \sum_{j=1}^{n} \boldsymbol{J}_n(i,j) \right\} = 8 \max_j \left\{ \sum_{i=1}^{n} \boldsymbol{J}_n(i,j) \right\} = 8 \left\| \boldsymbol{J}_n \right\|_{1\to\infty} \tag{30}$$

where the penultimate equality follows since $\boldsymbol{J}_n$ is symmetric. Now recalling $\rho_{\varepsilon,\delta}$ from Theorem B.1, (29) and (30) shows,

$$8\zeta^2 \rho_{\varepsilon,\delta} + \varepsilon^2 \Delta_0^2 \beta_0^2 \lesssim \left\| \boldsymbol{J}_n \right\|_{1\to\infty}^2 \rho_{\varepsilon,\delta} + \left\| \boldsymbol{J}_n^2 \right\|_{1\to\infty}^2 \beta_0^2 \leqslant \max\left\{1, \left\| \boldsymbol{J}_n \right\|_{1\to\infty}^4 \right\} \left(\rho_{\varepsilon,\delta} + \beta_0^2 \right).$$

The result now follows from Theorem B.1.

C Additional Experiments

We report additional simulations to evaluate the cost of privacy. In Figure 12 we generate Ising model synthetic outcomes on Erdős-Rényi, HIV network and the Political Blogs network, and plot the MSE of the private-MPLE estimates across a wide range of ε, and in the regimes of $\beta > 1$ and $\beta < 1$. The cost, quantified by the MSE shows a decreasing trend with ε, with the difference in regimes being stark in the Erdős-Rényi network.

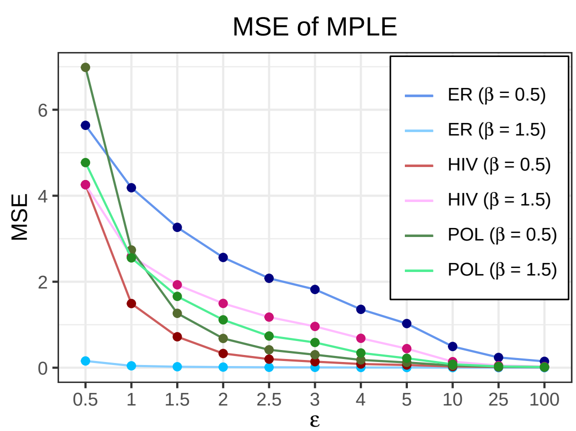

Figure 12: MSE of `PrIsing` estimates across ε for all networks in the paper

Next, we compare the privacy costs in a neighborhood of the estimated $\hat{\beta}$s. As noted in Sections 4.2.1 and Section 4.2.2 corresponding beta-hat turns out to be 1.8 and 2.85 respectively. As before, we generate Ising model synthetic outcomes with beta in a range around $\hat{\beta}$, and estimate $\hat{\beta}^{\texttt{priv}}$ 500 times to produce MSE values. We plot the results varying across epsilon, and plot the results in Figure 13(a) for HIV network and Figure 13(b) for political blogs network.

The results show a decreasing trend in MSE with increasing epsilon, re-ensuring that the MSE decreases as the privacy guarantee becomes weaker.

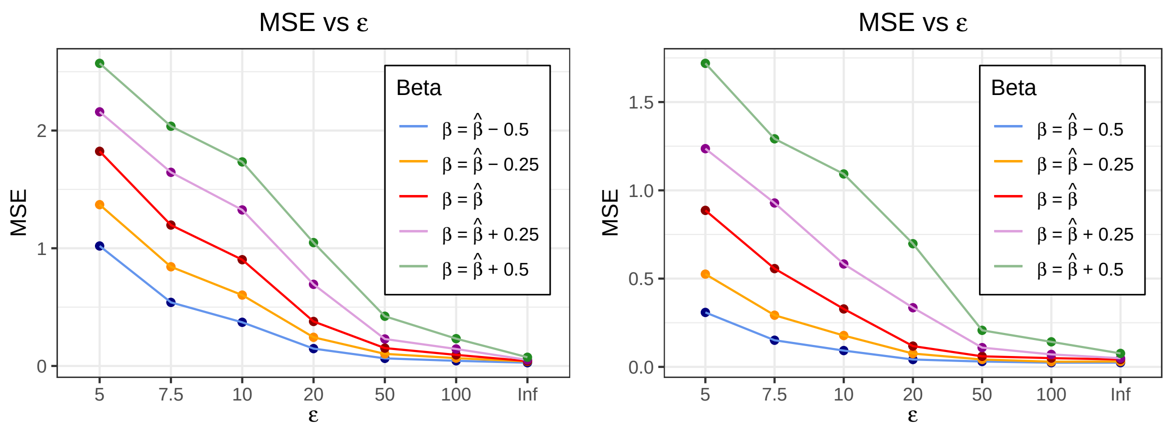

(a) (b)

Figure 13: MSE of `PrIsing` estimates across a range of β values around $\hat{\beta}$ using Ising Models on (a) HIV network and (b) Political Blogs network.

D Technical Details

In this section we provide technical lemmas needed for completing the proof of our results.

Lemma D.1. *Consider a sequence of random variables $\{X_n : n \geqslant 1\}$ and suppose that for every $\xi > 0$ there exists $K_1(\xi), K_2(\xi), K_3(\xi) > 0$ such that,*

$$\mathbb{P}\left(M_n \left|\tanh(K_1(\xi)X_n) + K_2(\xi)t_nX_n - \tanh(K_1(\xi)c) - K_2(\xi)t_nc\right| > K_3(\xi)\right) \leqslant \xi \tag{31}$$

for all $n \geqslant n_0(\xi)$, where $c > 0$ is a constant, $t_n > 0$ $\forall n \geqslant 1$ and $M_n \to \infty$ as $n \to \infty$. Then,

$$M_n\left|X_n - c\right| = O_p(1)$$

Proof. Observe that,

$$\begin{aligned}
&\Big|\tanh(K_1(\xi)X_n) + K_2(\xi)t_nX_n - \tanh(K_1(\xi)c) - K_2(\xi)t_nc\Big| \\
&\qquad\qquad = \Big|\tanh(K_1(\xi)X_n) - \tanh(K_1(\xi)c)\Big| + \Big|K_2(\xi)t_nX_n - K_2(\xi)t_nc\Big|
\end{aligned}$$

Then by (31) we get,

$$\begin{aligned}
&\mathbb{P}\left(M_n\Big|\tanh(K_1(\xi)X_n) - \tanh(K_1(\xi)c)\Big| > K_3(\xi)\right) \\
&\qquad \leqslant \mathbb{P}\left(M_n\left|\tanh(K_1(\xi)X_n) + K_2(\xi)t_nX_n - \tanh(K_1(\xi)c) - K_2(\xi)t_nc\right| > K_3(\xi)\right) \leqslant \xi
\end{aligned}$$

for all $n \geqslant n_0(\xi)$. Now for fixed $\xi > 0$ and using the mean value theorem,

$$\begin{aligned}
M_n|X_n - c| &= \frac{M_n}{K_1(\xi)}\left|\tanh^{-1}(\tanh(K_1(\xi)X_n)) - \tanh^{-1}(\tanh(K_1(\xi)c))\right| \\
&\leqslant \frac{M_n}{K_1(\xi)}\left|\frac{\tanh(K_1(\xi)X_n) - \tanh(K_1(\xi)c)}{1 - \zeta_\xi^2}\right|
\end{aligned} \tag{32}$$

where $\min\{\tanh(K_1(\xi)X_n), \tanh(K_1(\xi)c)\} \leqslant \zeta_\xi \leqslant \max\{\tanh(K_1(\xi)X_n), \tanh(K_1(\xi)c)\}$. By definition,

$$|1 - \zeta_\xi^2| = 1 - \zeta_\xi^2 \geqslant 1 - |\zeta_\xi| \geqslant 1 - |\tanh(K_1(\xi)c)| - |\tanh(K_1(\xi)X_n) - \tanh(K_1(\xi)c)|$$

Since $M_n \to \infty$, then there exists $n_1(\xi) > n_0(\xi)$ such that for all $n \geqslant n_1(\xi)$,

$$\frac{K_3(\xi)}{M_n} \leqslant K_4(\xi) := \frac{1}{2}(1 - |\tanh(K_1(\xi)c)|) \tag{33}$$

Note that on the event $|\tanh(K_1(\xi)X_n) - \tanh(K_1(\xi)c)| \leqslant K_4(\xi)$ with (33), we have,

$$|1 - \zeta_\xi^2| \geqslant K_4(\xi).$$

Hence recalling (32), on the event $|\tanh(K_1(\xi)X_n) - \tanh(K_1(\xi)c)| \leqslant K_4(\xi)$ we get,

$$M_n|X_n - c| \leqslant \frac{M_n}{K_1(\xi)K_4(\xi)} |\tanh(K_1(\xi)X_n) - \tanh(K_1(\xi)c)|$$

Now choosing $P(\xi) = \frac{K_3(\xi)}{K_1(\xi)K_4(\xi)}$ shows,

$$\begin{aligned}
\mathbb{P}\left(M_n\left|X_n - c\right| > P(\xi)\right) &\leqslant \mathbb{P}\left(M_n\left|X_n - c\right| > P(\xi), |\tanh(K_1(\xi)X_n) - \tanh(K_1(\xi)c)| \leqslant K_4(\xi)\right) \\
&\quad + \mathbb{P}\left(|\tanh(K_1(\xi)X_n) - \tanh(K_1(\xi)c)| > K_4(\xi)\right) \\
&\leqslant 2\mathbb{P}\left(|\tanh(K_1(\xi)X_n) - \tanh(K_1(\xi)c)| > \frac{K_3(\xi)}{M_n}\right) \leqslant 2\xi
\end{aligned}$$

for all $n \geqslant n_1(\xi)$, which completes the proof. ■