

UNIVERSITY COLLEGE LONDON

FACULTY OF SOCIAL AND HISTORICAL SCIENCES

DEPARTMENT OF ECONOMICS

**Essays on Individualized Treatment
Allocation and Network Spillover**

Guanyi Wang

*Submitted to University College London (UCL)
in fulfillment of the requirements for the degree
of Doctor of Philosophy in Economics*

April 2025

Declaration

I, Guanyi Wang, confirm that the work presented in my thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Guanyi Wang

April 14, 2025

Abstract

This thesis explores optimal individualized treatment allocation in social network settings.

The first chapter studies the individualized vaccine allocation under limited supply within a heterogeneous SIR network framework, leveraging social network data containing individual demographic characteristics and health status. By exploiting submodularity of the allocation problem, it devises a novel greedy algorithm to assign the treatment, with theoretical performance guarantees. Simulation results underline the importance of accounting for spillover effects when targeting vaccinations.

The second chapter focuses on treatment allocation in sequential decision games of interacting agents, where stationary distributions of outcomes follow Gibbs distributions. To overcome the analytical and computational challenges of direct optimization, it employs a variational approximation to characterize and estimate optimal treatment policies. I characterize the performance of the variational approximation, deriving a performance guarantee for the greedy optimization algorithm via a welfare regret bound. I implement our proposed method in simulation exercises and an empirical application using the Indian microfinance data (Banerjee et al., 2013), and show it delivers significant welfare gains.

The third chapter examines treatment allocation in large-scale simultaneous decision games with strategic complementarities. I introduce a maximin optimal treatment allocation rule that remains robust to the presence of multiple Nash equilibria. Remaining agnostic about the specific selection rule, I derive a closed-form expression for the boundary of the identified set of equilibrium outcomes. To address the incompleteness that emerges from unspecified selection, I propose a policy maximizing worst-case welfare. A greedy algorithm is used for implementation, with theoretical performance guarantees established through a welfare regret bound that accounts for both sampling uncertainty and the use of a greedy algorithm. Finally, I demonstrate this method with an application to the microfinance dataset of Banerjee et al. (2013).

Impact Statement

This thesis explores how *individualised treatment allocation on social networks* reshapes welfare, inequality and policy design. Because treatments received by one agent propagate through network spillovers, traditional “one-person-at-a-time” rules waste scarce resources and can even widen disparities. I study this problem across three complementary settings and show how modern optimisation and econometric tools translate into actionable guidance for health authorities, NGOs, and regulators.

Chapter 1. I develop a heterogeneous-SIR model that embeds detailed network data and prove that the planner’s welfare objective is *submodular*. This property lets us replace infeasible brute-force searches with a simple greedy algorithm that enjoys a 63% worst-case guarantee yet, in simulations, achieves virtually the global optimum. When applied to vaccine targeting, the rule lifts expected health welfare by 4–12% over random or demographically-based strategies. The result provides health agencies with a ready-to-use tool for distributing limited doses during pandemics or seasonal outbreaks.

Chapter 2. Turning to *sequential network games*—relevant for micro-finance adoption, agricultural innovations and other behaviours where decisions unfold over time—I introduce a variational approximation to the equilibrium Gibbs distribution, derive regret bounds, and embed both in a greedy assignment routine. Using Indian village data, the algorithm raises predicted loan uptake by an average 41% (and up to 137%) relative to the NGO’s historical practice. Practitioners can therefore amplify development interventions without additional subsidies, simply by re-ordering whom they target first.

Chapter 3. I tackle the pervasive problem of *multiple equilibria*. A maximin allocation rule is proposed that maximises *worst-case* welfare, together with closed-form expressions for the identified set of outcomes. This offers regulators and competition authorities a defen-

sible benchmark when uncertainty about strategic behaviour—e.g. in technology adoption or vaccination take-up—might otherwise stall action. Again using the Indian village data, our method achieves notably higher welfare levels, with average improvements of 116% even when multiple equilibria are present.

Societal benefits. Health ministries can embed the SIR-based routine in dashboards for real-time vaccine allocation. Development agencies and social enterprises can deploy the sequential-game tool to allocate interventions while accounting for stationary spillovers. Policymakers can remain agnostic about behavioural-equilibrium uncertainty yet still obtain robust welfare gains through the maximin rule.

Overall, the thesis provides scalable, open-source algorithms and rigorous performance guarantees that help translate complex network interactions into effective, equitable and resilient policy actions.

Research Paper Declaration Form

(i) **1. For a research manuscript that has already been published** (if not yet published, please skip to section 2):

(a) **What is the title of the manuscript?** Who Should Get Vaccinated? Individualized Allocation of Vaccines Over SIR Network

(b) **Please include a link to or doi for the work:** 10.1016/j.jeconom.2021.09.009

(c) **Where was the work published?** Journal of Econometrics

(d) **Who published the work?** Elsevier

(e) **When was the work published?** 15th September 2021

(f) **List the manuscript's authors in the order they appear on the publication:**
Toru Kitagawa, Guanyi Wang

(g) **Was the work peer reviewed?** Yes

(h) **Have you retained the copyright?** Yes

(i) **Was an earlier form of the manuscript uploaded to a preprint server (e.g. medRxiv)?** If 'Yes', please give a link or doi <https://arxiv.org/pdf/2012.04055>
If 'No', please seek permission from the relevant publisher and check the box next to the below statement:

☐ *I acknowledge permission of the publisher named under 1d to include in this thesis portions of the publication named as included in 1c.*

(ii) **For a research manuscript prepared for publication but that has not yet been published** (if already published, please skip to section 3):

(a) **What is the current title of the manuscript?**

(b) **Has the manuscript been uploaded to a preprint server 'e.g. medRxiv'?**
If 'Yes', please please give a link or doi:

(c) **Where is the work intended to be published?**

(d) **List the manuscript's authors in the intended authorship order:**

(e) **Stage of publication:**

(iii) **For multi-authored work, please give a statement of contribution covering all authors** (if single-author, please skip to section 4): All authors discussed the results and contributed to the final manuscript

(iv) **In which chapter(s) of your thesis can this material be found?** Chapter 1

e-Signatures confirming that the information above is accurate (this form should be co-signed by the supervisor/ senior author unless this is not appropriate, e.g. if the paper was a single-author work):

Candidate: Guanyi Wang

Date: 15 April 2025

Supervisor/Senior Author signature (where appropriate):

Date:

04/15/2025



Acknowledgments

I began my Ph.D. studies in 2018. The journey has not always been smooth sailing, but I have been fortunate to receive tremendous support from many people along the way.

First, I would like to thank my family for their unwavering financial and emotional support. My mother, Guiying, works tirelessly yet always reassures me that I need not worry about money. Although my grandmother, Chun, has difficulty speaking, she still does her best to talk with me whenever I call and never mentions the pain she endures from her illness. My grandfather, Zongxian, continually shows his love for me and even tries to save his own money so that he can help me even more. My aunts—Fang, Guirong, and Jingqiang—and my uncles—Hongchang, Mingqi, and Shunlin—always support me whenever I face important decisions or feel stressed, and every time I return home they welcome me with genuine kindness. My cousins, Danyang and Yuesong, provide invaluable support whenever I feel uncertain about the future, and I truly cherish our relationship.

Second, I would like to thank all of my supervisors. I first met Toru during my MRes econometrics course and appreciated his kindness in answering my many questions. Soon after he became my supervisor, the COVID-19 pandemic began, and I lost my motivation to study. I quarantined myself in a small room for about two months, during which my mental health deteriorated. Then I received an email from Toru asking how I was and whether I wanted to meet. I still remember how relieved I felt when I read it—it was as if an angel had stepped in. From that point on, we met weekly and produced our first co-authored paper, which became my first publication. We have collaborated ever since, and I have learned a great deal from the way he thinks and tackles problems. He also encouraged me to apply to and present at international conferences, where I gained valuable experience, met excellent scholars, and further advanced my academic development.

My first secondary supervisor, Martin, later moved to Oxford, yet he has remained a wonderful teacher of econometrics since my master's studies and continually encouraged me

while I prepared for the MPhil upgrade seminar.

Soon after that seminar, Aureo kindly agreed to become my second supervisor. At the time, both Toru and Martin had left UCL, and I felt I had no one to talk to; Aureo immediately took me under his wing. His guidance on my second paper helped me identify the research direction that ultimately became my job-market paper. I also value the regular emails he sends with interesting new work he thinks might be useful. Every conversation with him teaches me something new, and we recently became co-authors; I look forward to continuing to learn from him.

Although Andrei and Tim are not officially my supervisors, they have always treated me as if they were. Since Andrei joined UCL five years ago, he has attended my seminars and offered invaluable feedback; in fact, my job-market paper grew out of one of his questions at a UCL brown-bag seminar. He also encouraged me when I decided to spend an extra year in the Ph.D. programme. I met Tim while organising that seminar series, and I often reserve his student-meeting slots to discuss my research. Tim never complains, and I would like to express both my gratitude and my apologies for taking up so much of his time. I will also never forget the tacos we shared in the summer of 2024.

I would also like to thank all the faculty members in our department—especially Andrew, Ben, Daniel, Dennis, Duarte, Franck, Imran, Joao, Lars, Liyang, Martin, Pedro, Raffaella, Stephen, and Wei, among many others—for their consistent support. They have never hesitated to open their doors whenever I needed guidance. I am equally grateful to Daniella and Snjezana for their unfailing help with administrative matters, and to Nathan for his excellent IT support.

Throughout this journey, I have been fortunate to form lasting friendships with many fellow students—Chen-Wei, Dongwoo, Ertian, Hao, Hugo, Jeff, Jiacheng, Linda, Morgane, Nick, Nathan, Riccardo, Ruimin, Silvia, Tian, Thomas, Weisheng, Wenhao, Yanziyi, Yikai, Yiming, Yilun, Yuanqi, Ziyuan, and Ziyu, as well as everyone else I have met at UCL. I will always remember the comfort they offered whenever I felt discouraged; the hours they spent helping

me rehearse my job-market talks; the moments we untangled new proofs together; the drops of sweat we shed at the gym; the mountains we hiked; the papers we read and debated; and every other moment we shared.

To all of you, thank you.

Guanyi Wang

April 17, 2025

Contents

1	Who Should Get Vaccinated? Individualized Allocation of Vaccines Over SIR Network	3
1.1	Introduction	3
1.2	Setup and Identification	9
1.3	Estimation	15
1.4	Optimization	18
1.5	Regret Bounds	25
1.6	Simulation Exercises	30
1.7	Conclusion	38
A.1	The Transmission Term	40
A.2	Lemmas	41
A.3	Proofs for Theorems	46
2	Individualized Treatment Allocation in Sequential Network Games	51
2.1	Introduction	51
2.2	Model	58
2.3	Treatment Allocation	69
2.4	Simulation Exercises	85
2.5	Empirical Application	90
2.6	Conclusion	94

A.1	Lemma and Proposition	98
A.2	Theorem	117
A.3	Results from Previous Literature	122
3	Robust Network Targeting with Multiple Nash Equilibria	126
3.1	Introduction	126
3.2	Model	135
3.3	Counterfactual Analysis for the Target Population	144
3.4	Treatment Allocation	149
3.5	Theoretical Analysis	157
3.6	Empirical Application	168
3.7	Extension	172
3.8	Conclusion	175
A.1	Chi-Square Goodness of Fit Test	177
A.2	Lemmas	177
A.3	Theorems and Propositions	185
A.4	Preliminary Lemmas	192
A.5	Results from Previous Literature	201

List of Figures

- 1.1 Comparison between Greedy Algorithm and Random allocation 35
- 1.2 Comparison between Greedy Algorithm and Targeting Without Network Information 37
- 2.1 Comparison between the greedy algorithm and random allocation for the parameter set 1 88
- 2.2 Comparison between the greedy algorithm and random allocation for the parameter set 2 88
- 3.1 Comparisons between three approaches 172

List of Tables

1.1	Summary of the SIR parameter values	31
1.2	The value of welfare averaged over 100 random networks: Greedy Algorithm versus Brute Force	32
1.3	The value of welfare averaged over 100 random networks: Greedy Algorithm versus Random allocation	34
1.4	The value of welfare averaged over 100 random networks: Greedy Algorithm versus TWNI allocation	36
1.5	The percentage of vaccinated younger units in the second period under the vaccine allocation policies obtained by Greedy Algorithm	38
2.1	Summary of the parameter values	86
2.3	Comparison between the average welfare computed using Gibbs sampling and variational approximation for parameter set 1	90
2.4	Comparison between the average welfare computed using Gibbs sampling and variational approximation for parameter set 2	90
2.2	Comparison for parameter set 1	96
2.5	Welfare Gain compares the equilibrium welfare simulated under Greedy Allo- cation and the equilibrium welfare simulated under Original Allocation imple- mented by BSS	97

3.1	Minimal Welfare Gain compares the minimal simulated equilibrium welfare under the Robust allocation method and the simulated equilibrium welfare under the Original allocation implemented by BSS	171
A.2	Chi-square values based on number of rooms for 43 Indian villages microfinance data from Banerjee et al. (2013)	177

Chapter 1

Who Should Get Vaccinated?

Individualized Allocation of Vaccines Over SIR Network

1.1 Introduction

Allocation of a resource over individuals who interact within a social network is an important task in many fields, such as economics, medicine, education, and engineering ([Lee et al. \(2020\)](#), [Banerjee et al. \(2013\)](#), among others). One of the important policy decisions of this sort in pandemic times is how to allocate vaccines over heterogeneous individuals to control the spread of disease and protect the lives of vulnerable. It is crucial for the vaccine allocation rule to take into account the spillover effect of cutting transmission of the disease.

Since the start of COVID-19 pandemic, governments around the world have gone to great lengths to collect network data in which one can trace who is contacting whom. Motivated by these observations, we study how to estimate optimal individualized allocations of vaccines under capacity constraint, using micro-level social network data. Data is informative about the covariates of N units, their health status, and their associated neighbors. Using in-

sample information, we evaluate the risk to each unit, calculated from its own covariates and spillovers from its heterogeneous neighbors, using an *individualized Susceptible-Infectious-Recovered* model. The purpose of vaccine allocation is to maximize public health, by selecting units to be vaccinated. Obtaining an optimal assignment is, however, challenging since whether a treatment is optimal for an individual depends on which treatments are given to her neighbors. This implies that the search for an optimal allocation has to be performed over the entire network jointly, not individually.

This paper makes two main contributions. The first contribution is to develop methods to estimate vaccine assignment policies that exploit network information at the micro-level. The second contribution is to show that the empirical welfare criterion built upon the SIR spillover structure delivers a submodular objective function, which we exploit to obtain computationally attractive algorithms to solve the welfare optimization problem. Distinct from the existing approach of estimating individualized allocation policies under network interference (Viviano, 2019; Ananth, 2020a), our setting does not assume the availability of Randomized Control Trial (RCT) data. Instead, we assume the availability of estimated values of these spillover parameters from other sources, which we plug into our SIR model. Exploiting already estimated SIR parameter values for immediate targeting and allocation is useful when time is of the essence and the need for policy action is pressing, and avoids the cost of running an RCT.

To optimize the empirical welfare of allocation policies, one naive approach is to evaluate the value of empirical welfare exhaustively for all possible combinations of vaccine allocations over individuals. We refer to this as the brute-force approach. Although the brute-force approach is guaranteed to optimize the empirical welfare, it is not practicable since the number of possible combinations grows exponentially as the number of individuals in the network increases. On the other hand, giving up on optimization entirely and implementing random allocation is indeed practicable, but leads to a significant waste of the vaccine supply, which we show in our simulation exercises.

Given the challenge in optimizing the empirical welfare, what we recommend in this paper is an allocation policy obtained by greedy optimization. A greedy optimization algorithm in the current setting is to sequentially allocate a vaccine to an individual in the network who is most influential for improving the social welfare. In general, greedy algorithms are not guaranteed to yield an optimum. With the current welfare criterion built upon the SIR spillover structure, however, we can obtain a non-decreasing submodular objective function. Relying on the seminal result in discrete convex analysis shown by [Nemhauser et al. \(1978\)](#), we show that the greedy algorithm delivers an allocation policy at which the value of the objective function is worse than the optimum only up to a universal constant factor, independent of the spillovers, size, and density of the SIR networks. Our derivation of the population welfare regret of the greedily estimated allocation policy reflects the potential loss of welfare due to non-feasibility of obtaining the brute-force allocation policy.

We further illustrate the advantages of our method in our simulation exercises. In a small network setting (up to 35 individuals in the network), comparisons with the brute-force allocation rules reveal that our proposed greedy allocation rules leads to an optimal solution. In a large network setting, we evaluate the performance of our method versus two different assignment rules: random assignment, and targeting without considering network information. The welfare improvement relative to these two baselines ranges over 4% - 12%, and this result is insensitive to the values of SIR parameters and the size and density of network.

To assess how uncertainty in the SIR parameter estimates affect the welfare performance of the estimated policy, we derive a uniform upper bound of the welfare regret of our vaccine allocation rule and its convergence rate with respect to the size of the sample used for obtaining the SIR parameter estimates. The uniform upper bound of regret depends upon two things. Firstly, n , which is the sample size of the separate dataset used to estimate the SIR parameters. Secondly, the ratio of the network data sample size N to the maximum number of neighbors N_M plus the minimum between the number of infected units N_I and the number

of available vaccine doses d (i.e., $(d \min\{N_M, d\} + 2dN_M + \min\{N_I, d\})/N$). As N_M and N_I grow, the complexity and risk of the social network increase, which can reduce the welfare regret performance of the estimated vaccine allocation rule.

The remainder of this paper is organized as follows. We first discuss the relevant literature in the rest of this section. Section 2 details various models, and the HI-SIR model in particular, and the wider setting. Section 3 is concerned with estimation, including the estimation of SIR parameters and the construction of the QIP problem. The optimization procedure is contained in section 4. Section 5 contains the theoretical results. Simulation details are shown in Section 6, and Section 7 concludes. All proofs and derivations are shown in the appendix.

1.1.1 Related Literature

Our work contributes to the literature on statistical treatment rules, which was first introduced into econometrics by [Manski \(2004\)](#). The optimal treatment allocation regime has been studied in many fields, such as medical statistics ([Zhao et al., 2012, 2015](#)), operational research ([Loiola et al., 2007](#)) and economics. Following the pioneering works of [Hannan \(1957\)](#) and [Savage \(1951\)](#),¹ researchers in econometrics and machine learning often use regret to evaluate the performance of decision rules. The recent literature of statistical treatment rules includes [Dehejia \(2005\)](#), [Hirano and Porter \(2009\)](#), [Stoye \(2009, 2012\)](#), [Tetenov \(2012\)](#), [Bhattacharya and Dupas \(2012\)](#), [Kitagawa and Tetenov \(2018\)](#), [Zhou et al. \(2018\)](#), [Manski \(2019\)](#), [Kasy and Sautmann \(2019\)](#), [Athey and Wager \(2020\)](#), [Kock et al. \(2020\)](#), [Mbakop and Tabord-Meehan \(2021\)](#), [Manski and Tetenov \(2021\)](#), [Sakaguchi \(2021\)](#), and [Kitagawa et al. \(2021\)](#) among others. The planner’s objective function in the majority of these works is a sum of individual outcomes under the no-interference assumption (i.e., Stable Unit Treatment Value Assumption of [Rubin \(1974\)](#)). This assumption does not hold in this

¹[Hannan \(1957\)](#) considers regret-minimizing strategies in the context of zero-sum and sequential games. [Savage \(1951\)](#) introduces minimax-regret rules to the statistical decision theory.

study because of the network spillover effects that are present. To our knowledge, there are only two other papers that also consider the network setting in statistical treatment choice, which are [Viviano \(2019\)](#) and [Ananth \(2020a\)](#). These two papers assume the availability of pilot data from RCT studies performed over networks in order to form empirical welfare criteria. Their frameworks are not restricted to the SIR setting of the current paper and cover spillover structures commonly assumed in social science applications. In contrast, our approach forms welfare estimates by imposing the HI-SIR model structure and plugging in values of the primitive spillover parameters that are estimated or calibrated in some external study (e.g., [Baqae et al. \(2020\)](#)). Another notable difference is that we consider allocation policies that are not constrained other than via the capacity constraint, while [Viviano \(2019\)](#) and [Ananth \(2020a\)](#) assume the class of implementable allocation policies has a finite VC-dimension to control overfitting to the training RCT sample.

The SIR model was originally proposed by [Kermack and McKendrick \(1927\)](#), and is now the workhorse model in the epidemiological literature. Many extended versions have been studied in epidemiological analyses, such as the Susceptible-Infected-Susceptible model ([Nåsell, 1996](#)) and the Susceptible-Exposed-Infected-Recovered model ([Li and Muldowney, 1995](#)). During the global pandemic, an epidemiological literature has sprung up within economics. [Atkeson \(2020\)](#) and [Stock \(2020\)](#) introduced the SIR model into economics to study the implications of the current pandemic on the US economy. We introduce heterogeneity into the SIR model, which is similar to what [Acemoglu et al. \(2020\)](#) does in studying the Multi-Risk SIR model. That paper assumes, however, that the infection rate after the release of a vaccine equals zero, which means it does not consider the vaccine allocation problem. Our work contributes to the current literature by studying micro-level vaccine assignment rules in a heterogeneous SIR model with network information. In contrast, the existing works analyzing vaccine allocation rules focus on solving for the optimal proportion of vaccinated units in the population ([Pastor-Satorras and Vespignani \(2002\)](#), [Manski \(2010, 2017\)](#)). [Chen et al. \(2020\)](#) analyzes vaccine allocation using a heterogeneous SIR model, while they consider

vaccine allocation policies at the group-level rather than the individual-level.

We build a connection to the literature on using a submodular function to solve an optimization problem. The performance guarantee of a general greedy algorithm for solving submodular maximization problems with a cardinality constraint was first established by [Nemhauser et al. \(1978\)](#). The later literature links the cardinality constraint to a more general constraint: Matroid constraint ([Fisher et al., 1978](#); [Cunningham, 1985](#)). See [Bach \(2011\)](#) and [Krause and Golovin \(2014\)](#) for overviews of papers studying optimization of submodular functions. In this work, we discuss a submodular function with a uniform matroid constraint (i.e., capacity constraint) and a more general partition matroid constraint.


We notice that our approach to vaccine allocation problem is related to the influence maximization problem first formulated by [Kempe et al. \(2003\)](#). [Chen et al. \(2010\)](#) investigates submodularity of objective functions and greedy optimization algorithms in this problem. Applications of the influence maximization problem include targeting for viral marketing ([Domingos and Richardson, 2001](#)) and optimal information spread in social network ([Bakshy et al., 2011](#)). There are two widely studied information diffusion models in this literature: *Independent Cascade Model* ([Goldenberg et al., 2001](#)) and *Linear Threshold Model* ([Granovetter, 1978](#)). Despite some similarity between the diffusion models and our SIR model, this literature has not considered individualized vaccine allocation problem.

We also note that there is a growing literature on estimation of treatment effects under network interference. [Manski \(2013\)](#) discusses identification of treatment effects and spillover effects under a deterministic interference graph and a set of relevant potential outcomes. The increased number of network datasets that have recently become available has motivated further work on this topic, including [Sävje et al. \(2017\)](#), [Aronow et al. \(2017\)](#), [Athey et al. \(2018\)](#), [Basse et al. \(2019\)](#), and [Leung and Moon \(2019\)](#). [Li and Wager \(2020\)](#) non-parametrically estimates direct and indirect effects of treatment in a random network setting. [Vazquez-Bare \(2020\)](#) analyzes estimation of spillover effects using an instrument variable. See [Kline and Tamer \(2020\)](#) and [Graham and De Paula \(2020\)](#) for recent reviews

on econometric analysis in the presence of social interactions.

1.2 Setup and Identification

We consider a basic model to study the vaccination allocation problem. Let us first introduce the timeline and data setting that we consider in this work.



As shown in the illustration, we suppose there are two periods. At $t = 0$, policymakers initially observe the network structure A (i.e., adjacency matrix) linking N individuals, for which we provide further details below. Policymakers then observe covariates $X_i \in \mathcal{X} \subset \mathbb{R}^{d_x}$ and current period health state $H_0 \in \{S, I, R\}$ for each of the N individuals. The health states $\{S, I, R\}$ stand for *Susceptible*, *Infected*, and *Recovered*. We assume the network structure A is observed before personal health status to avoid the impact of self-isolation on the network structure. At $t = 1$, policymakers start to assign the vaccine. After a short vaccination period, people begin to meet their neighbors, which we call the interaction period. The health state during that period is defined as $H_1 \in \{S, I, R, D\}$, where D stands for death. Since at the time of assigning vaccination, H_1 is not yet observed by researchers, a stochastic health state will be used to evaluate personal risk. The ultimate goal of policymakers is to maximize the expected social health situation via the allocation of vaccines.

In our setting, units are connected through a social network. We assume the following property on network structure holds:

Assumption 1. (*Undirected Relationships*) The interference graph is undirected. i.e., $A_{ij} = A_{ji}$.

The symmetric $N \times N$ adjacency matrix A specifies who contacts with whom, with the (i, j) th element of A , denoted by A_{ij} , equal to one if unit i and unit j has positive contact time, and

zero otherwise. By convention, all the diagonal elements A_{ii} are equal to zero. If $A_{ij} = 1$, then we say that i and j are neighbors. Let N_i indicate the neighbors of unit i , then we write $A_{ij} = 1$ if $j \in N_i$ and $i \in N_j$. The size of spillover (i.e., the probability of disease transmission) between the units i and j depends not only on A_{ij} but also on the amount of their contact time and the transmission rates which are allowed to be asymmetric between them. We accordingly have a directed weighted network structure for the spillovers, as shown in later sections.

Now, let us introduce the notation that we use in the following sections. First, v_i is the individual vaccine assignment rule (i.e., $v_i = 1$ if unit i gets the vaccine). Let \mathbf{v} denote $(v_1, \dots, v_N) \in \{0, 1\}^N$, and \mathbf{X} denote $(X_1, \dots, X_N) \in \mathbb{R}^{N \times d_x}$. Let S_i be the susceptible state indicator in the first period (i.e., $S_i = \mathbb{1}_{\{H_{0i}=S\}}$), let I_i be the infected state indicator in the first period (i.e., $I_i = \mathbb{1}_{\{H_{0i}=I\}}$), and let R_i be the recovered state indicator in the first period (i.e., $R_i = \mathbb{1}_{\{H_{0i}=R\}}$). Moreover, let $|N_i|$ denote the number of neighbors of unit i (i.e., $|N_i| = \sum_j A_{ij}$).

1.2.1 Heterogeneous-Interacted-SIR model

To measure the personalized transition probability, we use a **HI-SIR** model. Our model is defined in discrete time within a simplified setting of two time periods. In the first period, we observe the health state of each unit H_0 , which belongs to S (Susceptible), I (Infected), or R (Recovered),

$$S_i + I_i + R_i = 1.$$

In the second period, the state variable is H_1 . Compared with H_0 , H_1 includes one more state D (Death).

$$\mathbb{1}_{\{H_{1i}=S\}} + \mathbb{1}_{\{H_{1i}=I\}} + \mathbb{1}_{\{H_{1i}=R\}} + \mathbb{1}_{\{H_{1i}=D\}} = 1.$$

Without the vaccine, the state can move from susceptible to infected, then to either recovery or death. Now, we consider the setting after introducing the vaccine. Generally, vaccination

has two purposes: the first is limiting the spread of disease, and the second is treatment. Vaccination builds up the immune system, which leads to recovery. However, the effectiveness of vaccination (i.e., the percentage of vaccinated units that recover) is not clear. For simplicity, we assume that assumption 2 holds.

Assumption 2. (PT) *Perfect Treatment*: A vaccinated unit enters the Recovered state, regardless of its previous state (i.e., $\Pr(H_{1i} = R | v_i = 1) = 1$).

To further simplify the setting, we split all units into a finite number of disjoint groups based on their characteristics. The infection rate between each group varies. This setting could be extended to the individual level, but the micro level infection rate would need to be known in this case. Here, we consider two groups and use **age** as a binary indicator: G_1 (Young) and G_2 (Old). We now define a_i and b_i as the group indicators (i.e., $a_i = \mathbb{1}_{\{X_i \in G_1\}}$ and $b_i = \mathbb{1}_{\{X_i \in G_2\}}$).

We specify one of the key components in SIR models, the infection rate of unit i , as:

$$q_i = \left[\beta_{11} \frac{\sum_{j \in N_i} I_j(1 - v_j)a_j}{|N_i|} + \beta_{12} \frac{\sum_{j \in N_i} I_j(1 - v_j)b_j}{|N_i|} \right] \cdot a_i + \left[\beta_{21} \frac{\sum_{j \in N_i} I_j(1 - v_j)a_j}{|N_i|} + \beta_{22} \frac{\sum_{j \in N_i} I_j(1 - v_j)b_j}{|N_i|} \right] \cdot b_i, \quad (1.1)$$

where $\beta_{sk} = -\kappa_s \ln(1 - c_{sk})$, c_{sk} is the probability of successful disease transmission following a contact between group s and group k (i.e., c_{11} measures the transmission probability from one unit to another within the young group, c_{12} is the corresponding probability of transmission from a unit in the old group to a unit in the young group, with similar definitions for c_{21} and c_{22}), and κ_s is the average number of contacts in group s at each time period. β_{sk} describes the effective contact rate of the disease between group s and group k . The derivation of equation 1.1 can be found in the appendix.

In the above expression, $I_j(1 - v_j)$ means a susceptible individual can only be infected by

neighbors who were infected and not vaccinated.² Those neighbors may come from various groups. We calculate the fraction of neighbors in each group and multiply them by the associated risk parameters. The risk parameter β_{sk} measures the probability that a susceptible individual in group k is infected by an infected individual from group s in one time period.

We now define $\{\gamma_1, \gamma_2\}$ as the recovery rate and $\{\delta_1, \delta_2\}$ as the mortality from infection in group 1 and group 2 respectively. Given this, we can formulate the probability of staying in the infection state for the infected unit i as:

$$p_i = 1 - a_i(\gamma_1 + \delta_1) - b_i(\gamma_2 + \delta_2). \quad (1.2)$$

Since the probability of recovery and death purely depend on personal physical fitness,³ there is no interactive part in equation 1.2. The transition probability to the infected state is then:

$$P_{Ii}(\mathbf{v}) \equiv \Pr(H_{1i} = I | X, \mathbf{v}, A, H_0) = S_i q_i \cdot (1 - v_i) + I_i p_i \cdot (1 - v_i).$$

In the above expression, the probability of an unvaccinated unit being infected has two components. The first is the probability of a healthy unit being infected. The second is the probability of staying in the infected state for those infected in the first period. Under Assumption 2, a vaccinated unit has zero probability of being infected. Similarly, the transition probability to the susceptible state is:

$$P_{Si}(\mathbf{v}) \equiv \Pr(H_{1i} = S | X, \mathbf{v}, A, H_0) = [1 - v_i - q_i(1 - v_i)] \cdot S_i.$$

An unvaccinated unit can only exit the susceptible state by infection. Therefore, the probability of staying in the susceptible state decreases with the risk parameter β_{sk} , which depends on

²This can also be thought of as an underlying assumption, which is commonly used in the epidemiological literature (e.g., [Pesaran and Yang \(2020\)](#)).

³This can be thought as a simplified assumption, which indicates the death rate does not depend on the availability of hospital spare capacity.

the number of infected neighbors and the number of contacts with them. The remaining two states do not rely on the network structure. First, the transition probability to the recovered state :

$$P_{Ri}(\mathbf{v}) \equiv \Pr(H_{1i} = R | X, \mathbf{v}, A, H_0) = v_i + [R_i + I_i(a_i\gamma_1 + b_i\gamma_2)] \cdot (1 - v_i).$$

In the above expression,⁴ recovery has two different sources. One is the vaccine, and the other is self-immunity. The effect of self-immunity is heterogeneous and varies with personal characteristics. The probability of building immunity in each group is γ_1, γ_2 . The last state is death, which occurs with probability

$$P_{Di}(\mathbf{v}) \equiv \Pr(H_{1i} = D | X, \mathbf{v}, A, H_0) = I_i(a_i\delta_1 + b_i\delta_2) \cdot (1 - v_i).$$

1.2.2 Optimal Vaccine Allocation Problem

In [Emanuel et al. \(2020\)](#), a group of medical ethics experts suggest a successful vaccine is needed to reduce death and morbidity from infection, and is also needed for the restoration of economic and social activity. Following that suggestion, we choose our baseline outcome variable as the weighted average of the probability of being healthy in the second period. The idea of using weighted probability is to allow a flexible policy target of the planner. For example, if the planner wants to incorporate the importance of economic recovery into the policy objective, she may want to weight more the probabilities of being healthy of those who can contribute more to the economic output. For instance, the planner could specify the weights on the individuals to depend on their individuals characteristics including working hours and other socioeconomic characteristics (i.e., $g_i = g(X_i)$).⁵ We assume the weight is

⁴A maintained assumption in this equation is that the probability of being reinfected for the recovered units is zero. We relax this assumption in Section 7.

⁵A maintained assumption in this expression is that for every unit, the weight is same for both susceptible and recovered states. This is a simplifying assumption that can be relaxed if we want to weight differently the susceptible and recovered states.

non-negative for every unit. Taking these into consideration, equation (1.3) specifies the goal of the vaccine allocation policy as a constrained optimization problem:

$$\max_{\mathbf{v} \in \{0,1\}^N} \frac{1}{N} \sum_{i=1}^N g_i \sum_{h \in \{S,R\}} P_{hi}(\mathbf{v}), \quad (1.3)$$

s.t.

$$\sum_{i=1}^N v_i \leq d,$$

where

$$P_{hi}(\mathbf{v}) = \Pr(H_{1i} = h | X, \mathbf{v}, A, H_0),$$

and $d \geq 1$ is a positive integer for the exogenous cardinality constraint. The main idea of the above objective function is to maximize the weighted probability of being in the susceptible or recovered state in the second period by appropriately assigning the d doses of vaccine at the end of the first period.

In equation (1.3), P_{hi} is the heterogeneous state transition function, which describes the probability of $h \in \{S, R\}$ in the second period. This transition probability depends on the individuals' covariates and previous state including whether being vaccinated or not, and the associated network structure. We adopt the HI-SIR model to formulate the above transition function, which has been provided in the previous subsection.

One relevant question is: *Will vaccine allocation change the network structure?* Yes, it would change the behaviour of vaccinated units. For example, vaccinated units prefer to go out as compared to unvaccinated units. Given this, the number of contacts at each time period κ_s and the network structure A would change after the vaccine allocation. Our framework allows the network structure to vary without affecting the optimal allocation of vaccines in a special case where *only the vaccinated units* change their behaviours. This is because under our perfect treatment assumption, the vaccinated units no longer spread the disease or be infected, and their behavioral changes do not affect the health statuses of the neighbors

and themselves. On the other hand, our framework cannot accommodate a general case where the unvaccinated units also change their behaviours, since if so the heterogeneous SIR parameters in the objective function change in response to the vaccine allocation. To allow this scenario, we could incorporate uncertainty as to the values of κ_s and A in the second period, for instance, by optimizing an objective function that takes the expectation of the SIR parameters the adjacency matrices conditional on \mathbf{v} . We do not, however, consider such an extension in this paper and leave this topic for future research.

1.3 Estimation

In order to measure the individual risk level using the HI-SIR model, we need to know the associated SIR parameters: transmission rate (i.e., $\beta_{11}, \beta_{12}, \beta_{21}, \beta_{22}$), and recovery rate (i.e., γ_1, γ_2). Given that we cannot observe the true value of those parameters, it is infeasible to evaluate the objective function (1.3) based on the in-sample information of (H_0, X) and A of the target network. We therefore assume access to a separate dataset with sample size n or an external study analyzing it, from which we can form estimates for these exogenous parameters. We construct an empirical version of the population welfare (1.3) and maximize over the feasible allocation policies. To reflect the precision of the SIR parameter estimates in the welfare performance of an estimated allocation rule, we explicitly take into account the sampling uncertainty of the parameter estimates in our derivation of the welfare regret upper bound.

1.3.1 Estimation of SIR Parameters

The estimation of infection rate and death rate always faces severe missing data problems as discussed in [Manski and Molinari \(2020\)](#). [Keeling and Rohani \(2011\)](#) points out that, usually, researchers first estimate the reproductive ratio \mathcal{R}_0 , which is the average number of

individuals that one sick person infects.

$$\mathcal{R}_0 = \beta \times \frac{1}{\gamma}.$$

Then, the infection rate can be derived from the estimated recovery rate $\hat{\gamma}$ and $\hat{\mathcal{R}}_0$. In our case, the reproductive ratio is heterogeneous at group level.

$$\mathcal{R}_{0sk} = \beta_{sk} \times \frac{1}{\gamma_s} \quad \forall s, k = 1, 2, \quad (1.4)$$

where \mathcal{R}_{0sk} is the number of infectious individuals in group s resulting from one sick person in group k . We need to estimate the average number of younger infectious and older infectious from one sick person in group 1 and group 2, and also the recovery rate in each group. Given these values, we can estimate $\beta_{11}, \beta_{12}, \beta_{21}, \beta_{22}$ from equation (1.4).

Remark 1.3.1. We do not discuss what is a desirable procedure for estimating the model parameters in this work, since the choice of estimator depends on the type of data (e.g., Seroprevalence data, Reported cases data, etc.). See [Keeling and Rohani \(2011\)](#) for further details. For the COVID-19 transmissions, estimation of homogeneous \mathcal{R}_0 and other SIR parameters has been performed in several papers including [Fernández-Villaverde and Jones \(2020\)](#), [Ferguson et al. \(2020\)](#), and [Korolev \(2021\)](#). They note the difficulty in calibrating critical parameters at an early stage of the pandemic due to the lack of credible data, which motivates partial identification analysis of [Manski and Molinari \(2020\)](#) and [Stoye \(2021\)](#). Our approach, however, assumes availability of credible point estimates and does not allow identified-set estimates for the SIR parameters. See [Ellison \(2020\)](#) and [Akbarpour et al. \(2020a\)](#) for recent estimates of heterogeneous SIR parameters.

1.3.2 Quadratic Integer Programming

Plugging the parameter estimates into our **HI-SIR** model, we now have the sample analog of the population maximization problem (1.3), which is

$$\begin{aligned} \max_{\mathbf{v} \in \{0,1\}^N} \mathcal{W}_n(\mathbf{v}), \quad \text{s.t.} \quad \sum_{i=1}^N v_i \leq d, \quad \text{where} \\ \mathcal{W}_n(\mathbf{v}) = \frac{1}{N} \sum_{i=1}^N g_i \sum_{h \in \{S,R\}} \hat{P}_{hi}(\mathbf{v}). \end{aligned}$$

We can formulate this optimization as a quadratic integer programming (**QIP**) problem, which in the context of an assignment problem over a network is synonymous with the Quadratic Assignment Problem (**QAP**) of [Koopmans and Beckmann \(1957\)](#). We can express $\mathcal{W}_n(\mathbf{v})$ as

$$\mathcal{W}_n(\mathbf{v}) = \frac{1}{N} \sum_{i=1}^N g_i \underbrace{\left[v_i + \left[R_i + (a_i \hat{\gamma}_1 + b_i \hat{\gamma}_2) I_i \right] (1 - v_i) + M_i S_i (1 - v_i) \right]}_{\text{Probability of being healthy}} \quad (1.5)$$

where

$$M_i = 1 - \frac{\sum_{j=1}^N (\hat{\beta}_{11} a_i a_j + \hat{\beta}_{12} a_i b_j + \hat{\beta}_{21} b_i a_j + \hat{\beta}_{22} b_i b_j) A_{ij} I_j (1 - v_j)}{|N_i|}.$$

For the probability of being healthy in equation (1.5), there are two linear terms and one quadratic term in \mathbf{v} . The first term measures the direct effect of vaccination. A vaccinated unit is safe from infection with 100% probability. The last two terms describe the probability of being free of infection for unvaccinated units. Infected units naturally recover with probability $\{\gamma_1, \gamma_2\}$, which depends on their own characteristics. For those units who are already recovered in the first period, they are free from infection in the second period. The last component takes into account the indirect effect of vaccination. For susceptible units, the probability of being infected by their infected neighbors is summarized by the interaction term.

After removing all the constant parts in equation (1.5), we obtain a simplified objective

function (i.e., $\mathcal{W}_n(\mathbf{v}) = F_n(\mathbf{v}) + \text{constant}$):

$$F_n(\mathbf{v}) = \sum_{i=1}^N \hat{c}_i v_i + \frac{1}{N} \sum_{i=1}^N \frac{1}{|N_i|} \sum_{j=1}^N T_{ij} (v_i + v_j - v_i v_j), \quad (1.6)$$

where

$$\hat{c}_i = g_i [1 - R_i - (a_i \hat{\gamma}_1 + b_i \hat{\gamma}_2) I_i - S_i] / N,$$

$$T_{ij} = g_i (\hat{\beta}_{11} a_i a_j + \hat{\beta}_{12} a_i b_j + \hat{\beta}_{21} b_i a_j + \hat{\beta}_{22} b_i b_j) A_{ij} I_j S_i.$$

Since F_n differs from \mathcal{W}_n only by an additive constant (conditional on the network structure and individual characteristics in the first period), maximizing F_n is equivalent to maximizing the original empirical welfare function \mathcal{W}_n . Therefore, from now on, we will focus on $F_n(\mathbf{v})$ as our new objective function. Within $F_n(\mathbf{v})$, there is a quadratic term plus linear components in \mathbf{v} . Current software is available to solve general QIP problems, such as CPLEX and Gurobi. However, both applications require a **symmetric weighting matrix**, which does not hold in our case. This asymmetric property comes from the infectious process, since disease can only be transmitted from infected units to susceptible units, but the reverse is not true. We discuss how to solve this QIP problem with showing and exploiting the submodular property of our objective function in the next section.

1.4 Optimization

1.4.1 Submodularity

We showed in the last section that we can formulate our objective function as QAP. This kind of problem is well known as an *NP-hard* and *NP-hard to approximate* problem (Cela, 2013). In general, we cannot solve QAP in polynomial time, which is an issue in practice. We shall, however, show that the quadratic integer programming in our vaccine allocation problem can

be linked to the submodular optimization problem. The benefit of submodularity is that there exist off-the-shelf algorithms that can solve a submodular minimization problem in exact polynomial time and *approximately* solve a submodular maximization problem with capacity constraint in polynomial time. The seminal result of [Nemhauser et al. \(1978\)](#) provides a universal bound for the quality of approximation as detailed below in Section 1.4.2.

Definition 1.4.1 (Submodular function). Let $\mathcal{N} = \{1, 2, \dots, N\}$. A real-valued set-function $F: 2^{\mathcal{N}} \rightarrow \mathbb{R}$ is submodular if and only if, for all subsets $A, B \subseteq \mathcal{N}$, we have: $F(A) + F(B) \geq F(A \cap B) + F(A \cup B)$.

In simple terms, submodularity describes the diminishing returns property. The marginal increase in the average probability of being healthy decreases in the number of vaccinated units. This property is crucial for the maximization algorithm. For ease of exposition, we express the simplified empirical welfare F_n as a set function with argument $V \in 2^{\mathcal{N}}$, where the binary vector of vaccine allocation $\mathbf{v} \in \{0, 1\}^N$ and V correspond by $V = \{i \in \mathcal{N} : v_i = 1\}$:

$$F_n(V) = \mathbf{v}^\top \hat{W} \mathbf{v} + \hat{C}^\top \mathbf{v} - \mathbf{1}_{N \times 1}^\top \hat{W} \mathbf{v} - \mathbf{v}^\top \hat{W} \mathbf{1}_{N \times 1}, \quad (1.7)$$

where

$$\hat{C} = \begin{bmatrix} \hat{c}_1 \\ \vdots \\ \hat{c}_N \end{bmatrix}, \quad \mathbf{1}_{N \times 1} = \begin{bmatrix} 1 \\ \vdots \\ 1 \end{bmatrix}, \quad \hat{W} = \begin{bmatrix} \hat{w}_{11} & \hat{w}_{12} & \cdots & \hat{w}_{1N} \\ \hat{w}_{21} & \hat{w}_{22} & \cdots & \hat{w}_{2N} \\ \vdots & \vdots & \ddots & \vdots \\ \hat{w}_{N1} & \hat{w}_{N2} & \cdots & \hat{w}_{NN} \end{bmatrix},$$

$$\hat{w}_{ij} = -\frac{A_{ij}g_i}{|N_i|N}(\hat{\beta}_{11}a_iS_ia_jI_j + \hat{\beta}_{12}a_iS_ib_jI_j + \hat{\beta}_{21}b_iS_ia_jI_j + \hat{\beta}_{22}b_iS_ib_jI_j).$$

We then denote the class of feasible allocation sets V subject to the cardinality constraint $|V| \leq d$ by $\mathcal{V}_d \equiv \{V \in 2^{\mathcal{N}} : |V| \leq d\}$. Since vaccinating additional individuals cannot reduce welfare, F_n is a *non-decreasing* set function, i.e., for any $V \subset V'$, $F_n(V) \leq F_n(V')$.⁶

⁶See the proof of Theorem 1.4.1 for a formal proof for the non-decreasing property of F_n .

The quadratic functional form of F_n shown in (1.7) can be linked to one classic submodular function called a *cut function*. Cut functions have been well studied in combinatorial optimization and graph theory. We apply some of the results from that literature (e.g., [Bach \(2011\)](#)).

Lemma 1.4.1. *Let $\hat{W} \in \mathbb{R}^{N \times N}$ and $\hat{C} \in \mathbb{R}^N$. Then the set function $F_n: V \mapsto \mathbf{v}^\top \hat{W} \mathbf{v} + \hat{C}^\top \mathbf{v} - \mathbf{1}_{N \times 1}^\top \hat{W} \mathbf{v} - \mathbf{v}^\top \hat{W} \mathbf{1}_{N \times 1}$ is submodular if and only if $\hat{w}_{ij} \leq 0 \ \forall i \neq j$.*

The proof is shown in the appendix. Note that the necessary and sufficient condition for submodularity shown in this lemma is distinct from negative semidefiniteness of the matrix \hat{W} . Since all the parameters in \hat{w}_{ij} are non-negative, we must have $\hat{w}_{ij} \leq 0, \forall i, j = 1, \dots, N$. This immediately leads to the following theorem:

Theorem 1.4.1. *The objective function $F_n(V)$ is a non-decreasing submodular function for any adjacency matrix, covariate values, and parameter estimates.*

Theorem 1.4.1 is the key result in our paper. It describes two important properties of our objective function; monotonicity and submodularity. We exploit these two properties to justify the uses of greedy maximization algorithms shown in the next subsection.

1.4.2 Greedy Maximization Algorithm

Greedy maximization algorithms for submodular functions have been studied and frequently used for well over forty years. The performance guarantee of the algorithm that we study was first introduced by [Nemhauser et al. \(1978\)](#). This algorithm essentially uses the diminishing returns property of the submodular function. The idea is to iteratively select the most valuable element until the capacity constraint is reached. At each round, the algorithm evaluates $\mathcal{O}(N)$ functions to identify the marginal gain of each element. The number of rounds depends on the capacity constraint d . As a result, the computational complexity of the greedy algorithm is of order $\mathcal{O}(N \cdot d)$, well below the computational complexity of the brute-force

search. Algorithm 1 presents the greedy maximization algorithm applied to maximization of the empirical welfare (1.7).

Algorithm 1: Capacity Constrained Greedy Algorithm

```

1: Input: Dataset  $\{S_i, I_i, R_i, a_i, b_i\}_{i=1}^N$ ,  $\{A_{ij}\}_{i,j=1}^N$ , estimated parameters
    $\{\hat{\beta}_{11}, \hat{\beta}_{12}, \hat{\beta}_{21}, \hat{\beta}_{22}, \hat{\gamma}_1, \hat{\gamma}_2\}$ , weight  $\{g_i\}_{i=1}^N$  and capacity constraint  $d$ ;
2: Initialization: Starting from the empty set  $V = \emptyset$ ;
if  $|V| < d$  then
    3: for each  $i \in \mathcal{N} \setminus V$  do
    4: Compute the marginal gain  $F_n(V + \{i\}) - F_n(V)$ ;
    5: Select  $i$  which maximizes the marginal gain and add it into the set  $V$ ;
else
    return the set  $V$ ;
end

```

In general, there is no performance guarantee of the greedy algorithm. However, as shown by Nemhauser et al. (1978) for a non-decreasing submodular function with cardinality constraint (i.e., capacity constraint in our case), the greedy maximization algorithm is guaranteed to yield an allocation rule $\hat{V} \in \mathcal{V}_d$ that satisfies $F_n(\hat{V}) \geq (1 - \alpha_d)F_n(\hat{V}^*)$, where $\hat{V}^* \in \mathcal{V}_d$ is a constrained optimum under the capacity constraint, and α_d is a positive constant that depends only on $d \geq 1$ and $\alpha_d \geq 1/e$ for all $d \geq 1$. This seminal result implies that the greedy maximization algorithm provides a universal optimization guarantee for non-decreasing submodular functions, $F_n(\hat{V}) \geq (1 - 1/e)F_n(\hat{V}^*) \approx 0.63F_n(\hat{V}^*)$. Since we show in Theorem 1.4.1 that our objective function is non-decreasing and submodular, we obtain the following theorem as an immediate corollary of our Theorem 1.4.1 and Nemhauser et al. (1978).

Theorem 1.4.2 (Nemhauser et al. 1978). *Let $F_n : 2^{\mathcal{N}} \rightarrow \mathbb{R}$ be the simplified empirical welfare function as defined in (1.7) and $\hat{V}^* \in \arg \max_{V \in \mathcal{V}_d} F_n(V)$, $d \geq 1$. The greedy algorithm shown in Algorithm 1 outputs $\hat{V} \in \mathcal{V}_d$ such that*

$$F_n(\hat{V}) \geq (1 - \alpha_d)F_n(\hat{V}^*) \geq (1 - 1/e)F_n(\hat{V}^*),$$

where $1 - \alpha_d \equiv 1 - \left(1 - \frac{1}{d}\right)^d$ is monotonically decreasing in d and converges to $1 - e^{-1}$ as $d \rightarrow \infty$.

1.4.3 Targeting Constraint

Up until now, we have only considered a simple capacity constraint in the vaccine assignment rule. In reality, Beyond the weight specification in the objective function, policymakers may want to prioritize some group over the others by limiting the number of vaccines that are administered in each group.⁷ For example, policymakers may limit access to vaccines for those people that can work at home. If we are able to divide individuals into two groups based on their job categories, into a group that can work at home and a group that cannot say, then policymakers can set an upper bound on the number of vaccines that are available for the work at home group.

We call this kind of constraint a *targeting constraint*, and impose it in our model in such way that each of the two age groups has a capacity constraint for the number of available vaccines:

$$\sum_{i: X_i \in G_1} v_i = \sum_{i=1}^N a_i v_i \leq d_1, \quad \sum_{i: X_i \in G_2} v_i = \sum_{i=1}^N b_i v_i \leq d_2.$$

This targeting constraint belongs to a general class of constraints: the so called *matroid* class. First, we use \mathcal{I} to describe the subset of $2^{\mathcal{N}}$ that is compatible with all of the constraints imposed. If we restrict the set of vaccinated agents V to belong to \mathcal{I} , which is part of a matroid $(\mathcal{Y}, \mathcal{I})$, this constraint is called a matroid constraint.

Definition 1.4.2 (Matroid). Let \mathcal{I} be a nonempty family of allowable subsets of \mathcal{N} . Then the tuple $(\mathcal{N}, \mathcal{I})$ is a matroid if it satisfies:

- (Hereditiy) For any $D \subset E \subset \mathcal{N}$, if $E \in \mathcal{I}$, then $D \in \mathcal{I}$.
- (Augmentation) For any $D, E \in \mathcal{I}$, if $|D| < |E|$, then there exists an $x \in E \setminus D$ such that $D \cup \{x\} \in \mathcal{I}$.

⁷A group in this section does not need to coincide with the group defining the heterogeneity of the SIR parameters. For example, we could divide units based on their job category, geographical location, or community.

Let \mathcal{N}_1 and \mathcal{N}_2 be the disjoint subsets partitioned by X_i ($\mathcal{N}_1 \cup \mathcal{N}_2 = \mathcal{N}$). We can represent the targeting constraint by

$$\mathcal{I} \equiv \{V : V \subseteq \mathcal{N}, |V \cap \mathcal{N}_1| \leq d_1, |V \cap \mathcal{N}_2| \leq d_2\}.$$

We can show that this $(\mathcal{N}, \mathcal{I})$ is a matroid referred to as a *partition matroid*. First, we show heredity. For any $D \subset E$, we must have $|D \cap \mathcal{N}_1| \leq |E \cap \mathcal{N}_1|$ and $|D \cap \mathcal{N}_2| \leq |E \cap \mathcal{N}_2|$. If $E \in \mathcal{I}$, then it means D must satisfy the targeting constraint in \mathcal{I} . Next, for any $D, E \in \mathcal{I}$, we must have $|D \cap \mathcal{N}_1|, |E \cap \mathcal{N}_1| \leq d_1$ and $|D \cap \mathcal{N}_2|, |E \cap \mathcal{N}_2| \leq d_2$. If $|D| < |E|$, then either $|D \cap \mathcal{N}_1| < |E \cap \mathcal{N}_1|$ or $|D \cap \mathcal{N}_2| < |E \cap \mathcal{N}_2|$ or both. As a result, there must exist an element x that belongs to $E \setminus D$ such that $|D \cup \{x\} \cap \mathcal{N}_1| \leq d_1$ and $|D \cup \{x\} \cap \mathcal{N}_2| \leq d_2$.

This problem of optimal treatment assignment subject to a partition matroid constraint is to maximize $F_n(V)$ over $V \in \mathcal{I}$. The following Algorithm 2 is guaranteed to produce a solution $\hat{V}' \in \mathcal{I}$. Greedy maximization algorithms subject to a partition matroid constraint performed for non-decreasing submodular functions attain at least 50% of the optimal welfare.

Algorithm 2: Targeting Constraint Greedy Algorithm

```
1: Input: Dataset  $\{S_i, I_i, R_i, a_i, b_i\}_{i=1}^N$ ,  $\{A_{ij}\}_{i,j=1}^N$ , estimated parameters  
    $\{\hat{\beta}_{11}, \hat{\beta}_{12}, \hat{\beta}_{21}, \hat{\beta}_{22}, \hat{\gamma}_1, \hat{\gamma}_2\}$ , weight  $\{g_i\}_{i=1}^N$ , capacity constraint  $d$ , and targeting  
   constraints  $d_1, d_2$ ;  
2: Initialization: Starting from the empty set  $V = \emptyset$  ;  
if  $|V| < d$  then  
    3: for each  $i \in \mathcal{N} \setminus V$  do  
    4: Compute the marginal gain  $F_n(V + \{i\}) - F_n(V)$ ;  
    5: Sort  $i$  in order of decreasing marginal gain  
    6. if  $\sum_{j \in V} a_j + a_{i(1)} \leq d_1 \cap \sum_{j \in V} b_j + b_{i(1)} \leq d_2$  then  
        7: Add the 1st element of  $i$  into  $V$ ;  
    else  
        8: Repeat step 6 with remaining  $i$ ;  
    end  
else  
    return the set  $V$ ;  
end
```

Proposition 1.4.1 (Fisher et al. 1978). Let $F_n : 2^{\mathcal{N}} \rightarrow \mathbb{R}$ be the simplified empirical welfare function as defined in (1.7) and $\hat{V}^{**} \in \arg \max_{V \in \mathcal{I}} F_n(V)$. The greedy maximization algorithm shown in Algorithm 2 outputs $\hat{V}' \in \mathcal{I}$ such that

$$F_n(\hat{V}') \geq \frac{1}{2} F_n(\hat{V}^{**}).$$

The performance guarantee of the greedy algorithm with targeting constraint is worse than the performance guarantee of Algorithm 1. This implies a trade-off between additional constraints and the accuracy of computation. In the next section, we discuss the welfare regret bounds of the allocation rules estimated by the above greedy algorithms.

1.4.4 Perfect Treatment Assumption and Submodularity

Recall Assumption 2 (Perfect Treatment): A vaccinated unit enters the Recovered state, regardless of its previous state (i.e., $\Pr(H_{1i} = R | v_i = 1) = 1$). There are three possible ways to relax this assumption:

- The recovered units can still spread disease.
- The recovered units will become susceptible after one period (few periods).
- Some percentage of vaccinated units remain susceptible or infected.

In the first case, if the person is recovered at H_0 , she will spread the disease during the first period. In that case, the recovered neighbors of unit i will be taken into account by the infection rate q_i . This will not, however, change the sign of our weighting matrix, hence submodularity (by Theorem 1.4.1) still holds. In the second case, if unit i is recovered in the first period (i.e., $H_{0i} = R$), she could become susceptible in the second period (i.e., $H_{1i} = S$). Then, she may be infected in the next period (i.e., $H_{2i} = I$). However, we only consider a one time period setting in this work, which rules out this risk. In the third case, varying this percentage only affects the coefficient of the linear term in the objective function (i.e., \hat{c}_i in equation 1.6), which is irrelevant to submodularity.

1.5 Regret Bounds

Following Manski (2004) and the subsequent literature on statistical treatment rules, we use regret to evaluate the performance of our algorithm for vaccine allocation. Let $F : 2^{\mathcal{N}} \rightarrow \mathbb{R}$ be the population analogue of $F_n(\cdot)$ in (1.7), where the estimated parameters are replaced by the truth. The expected regret measures the average difference in the welfare between using the constrained optimal assignment rule $V^* \in \arg \max_{V \in \mathcal{V}_d} F(V)$ and using the constrained estimated greedy algorithm \hat{V} obtained from Algorithm 1:

$$F(V^*) - \mathbb{E}_{P^n} [F(\hat{V})] = \mathbb{E}_{P^n} [F(V^*) - F(\hat{V})] \geq 0,$$

where \mathbb{E}_{P^n} is the expectation with respect to the sampling uncertainty of the parameter estimates in the external studies.

In this work, we assume that consistent estimators of effective contact rate and recovery rate are available from other studies. Generally, there is no requirement on the estimator except that Assumption 3 needs to hold.

Assumption 3. Let $\hat{\beta}_{sk}$ denote the estimate of effective contact rate between group s and group k , and $\hat{\gamma}_s$ denote the estimate of recovery rate in group s . The following properties need to hold:

$$\mathbb{P}\left\{\left|\hat{\beta}_{sk} - \beta_{sk}\right| \geq \epsilon\right\} \leq 2e^{-2n\epsilon^2} \quad \forall s, k = 1, 2.$$

$$\mathbb{P}\left\{\left|\hat{\gamma}_s - \gamma_s\right| \geq \epsilon\right\} \leq 2e^{-2n\epsilon^2} \quad \forall s = 1, 2,$$

where \mathbb{P} is the sampling distribution in another study that has sample size n .

The above assumption is an exponential tail bound obtained by applying Hoeffding's large deviation inequality (Hoeffding, 1963). Since β_{sk} is the effective contact rate of the disease between group s and k , and γ_s is the recovery rate in group s , both are naturally bounded in $[0, 1]$. Hence, common estimators (e.g., sample analog) meet the above condition. However, other tail bounds might apply for some other estimators, which do not necessarily have the same form as the above tail bound. Our approach can accommodate various tail bounds, such as the tail bound associated with the maximum likelihood estimator (Miao, 2010).

The estimators for the contact rates and recovery rates may come from different studies with different sample sizes. In this case, we can view n in Assumption 3 as the smallest sample size among the studies.

In order to derive the uniform convergence rate of the welfare regret, we decompose regret into three components as follows.

$$F(V^*) - F(\hat{V}) = \underbrace{F(V^*) - F_n(\hat{V}^*)}_{\textcircled{1}} + \underbrace{F_n(\hat{V}^*) - F_n(\hat{V})}_{\textcircled{2}} + \underbrace{F_n(\hat{V}) - F(\hat{V})}_{\textcircled{3}},$$

where V^* is an oracle optimum $V^* = \arg \max_{V \in \mathcal{V}_d} F(V)$, \hat{V}^* is a constrained optimal solution

to the estimated welfare $\hat{V}^* = \arg \max_{V \in \mathcal{V}_d} F_n(V)$, and \hat{V} is the output from the greedy maximization algorithm under the capacity constraint. Therefore, ① describes the regret we would attain if the constrained optimum could be computed exactly. ② measures the welfare loss introduced by the greedy algorithm. ③ indicates the loss from using the estimated objective function instead of the true objective function. We compute the upper bound of each component separately and then combine them.

First, we start from the derivation of the upper bound of ①. This part is similar to the approach in [Kitagawa and Tetenov \(2018\)](#). Before looking at V^* , consider the following inequality, which holds for any $\tilde{V} \in \mathcal{V}_d$:

$$\begin{aligned} F(\tilde{V}) - F_n(\hat{V}^*) &\leq F(\tilde{V}) - F_n(\tilde{V}) \\ &\quad (\because F_n(\hat{V}^*) \geq F_n(\tilde{V})) \\ &\leq \sup_{V \in \mathcal{V}_d} |F_n(V) - F(V)|. \end{aligned}$$

Since the above inequality applies to $F(\tilde{V})$ for all \tilde{V} , it also applies to V^* :

$$F(V^*) - F_n(\hat{V}^*) \leq \sup_{V \in \mathcal{V}_d} |F_n(V) - F(V)|.$$

For the second component, we can obtain an upper bound by applying Theorem [1.4.2](#):

$$\begin{aligned} F_n(\hat{V}^*) - F_n(\hat{V}) &\leq \frac{1}{e} F_n(\hat{V}^*) \\ &\leq \frac{1}{e} (F_n(\hat{V}^*) - F(\hat{V}^*)) + \frac{1}{e} F(V^*) \\ &\leq \frac{1}{e} |F_n(\hat{V}^*) - F(\hat{V}^*)| + \frac{1}{e} F(V^*) \\ &\leq \frac{1}{e} \sup_{V \in \mathcal{V}_d} |F_n(V) - F(V)| + \frac{1}{e} F(V^*). \end{aligned} \tag{1.8}$$

Similarly to the first component, the third component can be bounded as:

$$F_n(\hat{V}) - F(\hat{V}) \leq |F_n(\hat{V}) - F(\hat{V})| \leq \sup_{V \in \mathcal{V}_d} |F_n(V) - F(V)|.$$

Combining all the previous results, we obtain the upper bound of regret:

$$F(V^*) - F(\hat{V}) \leq \left(2 + \frac{1}{e}\right) \sup_{V \in \mathcal{V}_d} |F_n(V) - F(V)| + \frac{1}{e} F(V^*). \quad (1.9)$$

Compared with the regret upper bound when one could compute \hat{V}^* , the regret upper bound shown in (1.9) has one additional term $\frac{1}{e} \sup_{V \in \mathcal{V}_d} |F_n(V) - F(V)| + \frac{1}{e} F(V^*)$. This additional term comes from equation (1.8) and captures the welfare loss induced by the use of greedy algorithm. As we characterize below, the first term converges to zero as $n \rightarrow \infty$ under Assumption 3, while the second term remains independent of the accuracy of the parameter estimates. A simulation study in Section 6 assesses the magnitude of the optimization error of the greedy algorithm, and shows numerically that the greedy algorithm yields an exact optimum for small network cases ($N = 35$) at least. Based on this, we believe that the optimization error term of the greedy algorithm is much smaller than the universal theoretical bound $\frac{1}{e} F(V^*)$.

In the partition matroid (targeting constraint) case, by applying Proposition 1.4.1 and repeating the arguments to derive (1.9), we obtain

$$F(V^{**}) - F(\hat{V}') \leq \frac{5}{2} \sup_{V \in \mathcal{I}} |F_n(V) - F(V)| + \frac{1}{2} F(V^{**}),$$

where V^{**} is an oracle optimum under the targeting constraint, $V^{**} \in \arg \max_{V \in \mathcal{I}} F(V)$.

In order to bound $\sup_{V \in \mathcal{V}_d} |F_n(V) - F(V)|$, we use the triangle inequality to find the bound

of $|F_n(V) - F(V)|$:

$$\begin{aligned}
|F_n(V) - F(V)| &= \left| \mathbf{v}^\top (\hat{W} - W) \mathbf{v} + (\hat{C}^\top - C^\top) \mathbf{v} - \mathbf{1}_{N \times 1}^\top (\hat{W} - W) \mathbf{v} - \mathbf{v}^\top (\hat{W} - W) \mathbf{1}_{N \times 1} \right| \\
&\leq \left| \mathbf{v}^\top (\hat{W} - W) \mathbf{v} \right| + \left| (\hat{C}^\top - C^\top) \mathbf{v} \right| + \left| \mathbf{1}_{N \times 1}^\top (\hat{W} - W) \mathbf{v} \right| + \left| \mathbf{v}^\top (\hat{W} - W) \mathbf{1}_{N \times 1} \right| \\
&\leq \mathbf{v}^\top |\hat{W} - W| \mathbf{v} + |(\hat{C}^\top - C^\top)| \mathbf{v} + \mathbf{1}_{N \times 1}^\top |\hat{W} - W| \mathbf{v} + \mathbf{v}^\top |\hat{W} - W| \mathbf{1}_{N \times 1},
\end{aligned}$$

where the absolute value of a matrix or vector stands for the element-wise absolute values.

Therefore, we can decompose the maximal deviation $\sup_{V \in \mathcal{V}} |F_n(V) - F(V)|$ into four parts :

$$\begin{aligned}
\sup_{V \in \mathcal{V}_d} |F_n(V) - F(V)| &\leq \sup_{V \in \mathcal{V}_d} \mathbf{v}^\top |\hat{W} - W| \mathbf{v} + \sup_{V \in \mathcal{V}_d} |(\hat{C}^\top - C^\top)| \mathbf{v} \\
&\quad + \sup_{V \in \mathcal{V}_d} \mathbf{1}_{N \times 1}^\top |\hat{W} - W| \mathbf{v} + \sup_{V \in \mathcal{V}_d} \mathbf{v}^\top |\hat{W} - W| \mathbf{1}_{N \times 1}.
\end{aligned} \tag{1.10}$$

Under Assumption 3, we can obtain an upper bound for the mean of each element in $\hat{W} - W$ and $\hat{C} - C$, as shown in the next lemma.

Lemma 1.5.1. *Under Assumption 1, 2, and 3, we have*

$$\mathbb{E}_{P^n} |\hat{w}_{ij} - w_{ij}| \leq \sqrt{\frac{1 + \ln(2)}{2n}} \frac{A_{ij} g_i}{N}, \quad \mathbb{E}_{P^n} |\hat{c}_i - c_i| \leq \sqrt{\frac{1 + \ln(2)}{2n}} \frac{I_i g_i}{N}.$$

Combining this lemma with equations (1.9) and (1.10), we obtain the following theorem:

Theorem 1.5.1. *Let $N_M = \max_{i \in \mathcal{N}} |N_i|$, N_I be the total number of infected units, and $g = \max_{i \in \mathcal{N}} g_i$. Under Assumptions 1, 2, and 3, we have*

$$\mathbb{E}_{P^n} [F(V^*) - F(\hat{V})] \leq \bar{C} \cdot \frac{g[d \min\{N_M, d\} + 2dN_M + \min\{N_I, d\}]}{N} \sqrt{\frac{1}{n}} + \frac{1}{e} F(V^*),$$

where \bar{C} is a universal constant and d is the number of available vaccine doses.

Proof of the above theorem is shown in the appendix. In Theorem 1.5.1, we provided a distribution-free upper bound on the expected regret. We show that the convergence rate of the upper bound depends on the network data sample size N and also the sample size n for estimating the SIR parameters. At the same time, the regret upper bound is increasing

in the complexity and the riskiness of the network. The intuition is that our algorithm finds it harder to identify the most valuable units when the maximum number of edges and the number of infected individuals in the network increases. The maximum individual weight g also boosts the upper bound of regret. Moreover, our algorithm finds it harder to identify the best allocation rule when the number of possible combinations increase, which occurs when the capacity constraint is relaxed. This also implies the benefit of quarantine. Since quarantine controls the maximum number of connections in the network, the effectiveness of vaccine allocation is boosted by such government policy. Therefore, there is advantage to complementing a vaccine assignment policy with quarantine, which is evidenced by our simulation exercises.

1.6 Simulation Exercises

In this section, we use an Erdős-Renyi model to generate random social networks. In each of the following tables, we use 100 different networks and take the average of the outcome variable across all of the networks. We further show the standard deviation of in-sample welfare to understand the variation of network structure. We choose the probability of allocating a unit to group 1 to be 40% and the probability of allocating a unit to group 2 to be 60% (i.e., $\mathbb{P}(X_i = G_1) = 0.4$ and $\mathbb{P}(X_i = G_2) = 0.6$). In the epidemiological literature, researchers usually find the steady state of the SIR parameters. In order to identify the impact of varying the SIR parameters, we choose two different sets of parameter values to run the simulation. Throughout our simulation studies, we do not consider sampling errors in the parameter estimates and focus on optimizing welfare with the true parameter values plugged in. Table [1.1](#) summarizes all the values of the SIR parameters that we have used.

<i>Parameters</i>	<i>set 1</i>	<i>set 2</i>	<i>Parameters</i>	<i>set 1</i>	<i>set 2</i>
β_{11}	0.7	0.8	β_{12}	0.5	0.5
β_{21}	0.5	0.7	β_{22}	0.6	0.7
γ_1	0.1	0.1	γ_2	0.05	0.025

Table 1.1: Summary of the SIR parameter values

In addition, we choose three different densities, 0.1, 0.5 and 1, in order to identify the effect of network complexity. Here, *density* = 1 means that the network is fully connected (i.e., complete graph). We choose *full* to understand the behaviour of our heuristic algorithm not only in the sparse network case but also in the densest case. We also compare three capacity constraints, $d = 7\%N, 10\%N, 20\%N$, to evaluate the marginal performance gain of our greedy algorithm. We choose equal weight in the following comparisons. We, however, show the impact of changing weights on the number of vaccinated younger units in Table 1.5.

In the following sections, we compare our greedy algorithm with three familiar allocation rules. We first compare our algorithm with a brute force method in order to find the difference between the potentially sub-optimal greedy solution and the brute-force optimal solution. However, the number of possible combinations dramatically increases with the number of nodes and the capacity constraint. We cannot use a large number of agents to compute the brute force optimum in the simulation. Given this, in Section 1.6.2, we use a random assignment rule as a baseline to evaluate the performance of our algorithm in a large network setting. The third allocation rule that we compare our greedy algorithm with is an allocation rule which assigns the vaccine without considering network information. We compare the greedy algorithm with this third rule in Section 1.6.3.

1.6.1 Comparing with Brute Force

<i>Allocation Rule</i>	<i>Greedy Algorithm</i>			<i>Brute Force</i> ⁸		
<i>Capacity Constraint</i>	$d = 7\%N$	$d = 10\%N$	$d = 20\%N$	$d = 7\%N$	$d = 10\%N$	$d = 20\%N$
Parameter set 1						
$N = 500, \text{density} = 0.1$	0.60 (0.21)	0.65 (0.22)	0.77 (0.26)	0.60 (0.21)	0.65 (0.22)	0.77 (0.26)
$N = 500, \text{density} = 0.5$	0.47 (0.40)	0.51 (0.39)	0.63 (0.39)	0.47 (0.40)	0.51 (0.39)	0.63 (0.39)
$N = 500, \text{density} = 1^*$	0.33	0.37	0.49	0.33	0.37	0.49
$N = 800, \text{density} = 0.1$	0.58 (0.23)	0.66 (0.25)	0.76 (0.26)	0.58 (0.23)	0.66 (0.25)	0.76 (0.26)
$N = 800, \text{density} = 0.5$	0.44 (0.38)	0.52 (0.38)	0.62 (0.38)	0.44 (0.38)	0.52 (0.38)	0.62 (0.38)
$N = 800, \text{density} = 1^*$	0.30	0.36	0.46	0.30	0.36	0.46
Parameter set 2						
$N = 500, \text{density} = 0.1$	0.59 (0.27)	0.64 (0.29)	0.77 (0.32)	0.59 (0.27)	0.64 (0.29)	0.77 (0.32)
$N = 500, \text{density} = 0.5$	0.42 (0.49)	0.46 (0.49)	0.59 (0.49)	0.42 (0.49)	0.46 (0.49)	0.59 (0.49)
$N = 500, \text{density} = 1^*$	0.25	0.29	0.41	0.25	0.29	0.41
$N = 800, \text{density} = 0.1$	0.57 (0.28)	0.65 (0.30)	0.76 (0.31)	0.57 (0.28)	0.65 (0.30)	0.76 (0.31)
$N = 800, \text{density} = 0.5$	0.40 (0.47)	0.47 (0.47)	0.58 (0.47)	0.40 (0.47)	0.47 (0.47)	0.58 (0.47)
$N = 800, \text{density} = 1^*$	0.22	0.29	0.40	0.22	0.29	0.40

Table 1.2: The value of welfare (the sum of probabilities of being healthy in the second period) averaged over 100 random networks (standard errors in parentheses). We use the Greedy Algorithm or the Brute Force to determine who in each network should be vaccinated.

* When the density of Erdős-Renyi network equals 1, standard errors are 0 since there is no randomness in the network structure.

Since Theorem 1.4.2 shows the gap between the optimal solution and the heuristic result is at most 37%, we want to explore this theoretical difference using numerical study. We list all the possible combinations and use brute force to search for the optimal solution given a manageable number of units. We specify the maximum number of units to be $N = 35$, which is limited by computer performance. As the number of nodes increases, the possible number of combinations grows exponentially. The memory requirement and running time become impractical in a more realistic case. We recognize that the results from a small network may

⁸We compare all possible combinations given the capacity constraint and select the set V that maximizes \mathcal{W}_n .

not be accurate in a large network setting, but help us to understand the regret of our greedy algorithm to some degree. We summarize the in-sample welfare \mathcal{W}_n of these two approaches in Table 1.2.

In the small network case, we find that our greedy algorithm finds optimal allocation rules in all cases that we consider, which indicates a good performance of our method. We also notice that the welfare that is associated with the optimum decreases with the number of edges. As we relax the capacity constraint, welfare increases rapidly. The main purpose of this comparison is to get an idea of how much worse the empirical welfare at the greedy solution can be relative to the brute force optimum. More results are illustrated in the following two sections.

1.6.2 Comparing With Random Assignment

In this section, we use a random assignment rule to define the baseline of vaccine allocation. We randomly draw an allocation 10,000 times and calculate the average value of the outcome variable. Random allocation is one common assignment rule for policymakers. The purpose of this simulation is to learn about the improvement of our greedy allocation rule. In order to evaluate its performance in a relatively large network setting, we choose $N = 500$ and 800. Table 1.3 records the main differences in terms of in-sample welfare between these two methods.

From Table 1.3, we find that the performance of both methods decreases with the number of edges, which is also true for the first comparison. As the number of edges increase, the greedy algorithm finds it harder to identify who is relatively crucial in the network, which supports our interpretation of Theorem 1.5.1 in the previous section. This effect becomes more pronounced as the capacity constraint is relaxed. In the most extreme case, when everyone is connected with each other, the performance of our method is still better than the random assignment rule. This performance gap widens with the capacity constraint. We also find that the average welfare increases by 12% when the capacity constraint increases

by $0.1N$. Moreover, this improvement is robust with respect to the variation of number of nodes and the changes of density levels of network. The number of nodes decreases the performance of our method in a sparse network setting. For $N = 800$, welfare in the densest network is 14% lower than the welfare with $density = 0.5$, no matter which capacity constraint and parameter set we use.

Allocation Rule Capacity Constraint	Greedy Algorithm			Random Assignment⁹		
	$d = 7\%N$	$d = 10\%N$	$d = 20\%N$	$d = 7\%N$	$d = 10\%N$	$d = 20\%N$
Parameter set 1						
$N = 500, density = 0.1$	0.61 (<0.01)	0.65 (0.01)	0.77 (0.01)	0.57 (<0.01)	0.59 (<0.01)	0.66 (<0.01)
$N = 500, density = 0.5$	0.61 (<0.01)	0.64 (<0.01)	0.76 (<0.01)	0.57 (<0.01)	0.59 (<0.01)	0.66 (<0.01)
$N = 500, density = 1^*$	0.61	0.64	0.76	0.57	0.59	0.66
$N = 800, density = 0.1$	0.59 (0.04)	0.63 (0.04)	0.75 (0.04)	0.55 (0.04)	0.57 (0.04)	0.64 (0.04)
$N = 800, density = 0.5$	0.48 (0.07)	0.51 (0.07)	0.63 (0.07)	0.44 (0.07)	0.46 (0.07)	0.53 (0.07)
$N = 800, density = 1^*$	0.34	0.37	0.49	0.30	0.32	0.39
Parameter set 2						
$N = 500, density = 0.1$	0.60 (0.01)	0.64 (0.01)	0.77 (0.01)	0.56 (<0.01)	0.59 (<0.01)	0.66 (0.01)
$N = 500, density = 0.5$	0.60 (<0.01)	0.64 (<0.01)	0.76 (<0.01)	0.56 (<0.01)	0.59 (<0.01)	0.66 (<0.01)
$N = 500, density = 1^*$	0.60	0.64	0.76	0.56	0.59	0.66
$N = 800, density = 0.1$	0.58 (0.05)	0.62 (0.05)	0.74 (0.05)	0.54 (0.05)	0.56 (0.05)	0.63 (0.05)
$N = 800, density = 0.5$	0.44 (0.08)	0.48 (0.08)	0.60 (0.08)	0.40 (0.08)	0.43 (0.08)	0.50 (0.08)
$N = 800, density = 1^*$	0.26	0.30	0.42	0.23	0.25	0.32

Table 1.3: The value of welfare (the sum of probabilities of being healthy in the second period) averaged over 100 random networks (standard errors in parentheses). We use the Greedy Algorithm or the Random allocation to determine who in each network should be vaccinated.

* When the density of Erdős-Renyi network equals 1, standard errors are 0 since there is no randomness in the network structure.

If we look at the random assignment rule in Table 1.3, its performance is much worse than the performance of the greedy algorithm. This difference increases when the complexity of and the number of nodes in the network increase. The performance of the random

⁹In Random allocation, we randomly select an assignment 10,000 times for each network and take the average value of the outcome variable.

assignment rule improves as we relax the capacity constraint. However, this improvement is only about 7% when the capacity constraint increases by $0.1N$. Compared with the greedy algorithm, random assignment is less effective. Given its scarcity, we waste considerable resources by randomly assigning the vaccine. Looking at the situation of full edges, the performance of random allocation is inferior. The ratio of the welfare attained by random allocation to the welfare attained by the greedy algorithm is illustrated in Figure 1.1. This ratio increases slowly with the number of edges and decreases with the number of nodes in the network. In addition, the ratio decreases in an obvious way with the number of vaccines that are available.




Figure 1.1: Comparison between Greedy Algorithm and Random allocation

1.6.3 Comparing With Targeting Without Network Information

Allocation Rule Capacity Constraint	Greedy Algorithm			TWN ¹⁰		
	$d = 7\%N$	$d = 10\%N$	$d = 20\%N$	$d = 7\%N$	$d = 10\%N$	$d = 20\%N$
Parameter set 1						
$N = 500, \text{density} = 0.1$	0.61 (<0.01)	0.65 (0.01)	0.77 (0.01)	0.57 (<0.01)	0.59 (0.01)	0.65 (0.01)
$N = 500, \text{density} = 0.5$	0.61 (<0.01)	0.64 (<0.01)	0.76 (<0.01)	0.57 (<0.01)	0.59 (<0.01)	0.65 (<0.01)
$N = 500, \text{density} = 1^*$	0.61	0.64	0.76	0.57	0.59	0.65
$N = 800, \text{density} = 0.1$	0.59 (0.04)	0.63 (0.04)	0.75 (0.04)	0.56 (0.04)	0.58 (0.04)	0.65 (0.04)
$N = 800, \text{density} = 0.5$	0.48 (0.07)	0.51 (0.07)	0.63 (0.07)	0.44 (0.07)	0.47 (0.07)	0.54 (0.07)
$N = 800, \text{density} = 1^*$	0.34	0.37	0.49	0.30	0.33	0.40
Parameter set 2						
$N = 500, \text{density} = 0.1$	0.60 (0.01)	0.64 (0.01)	0.77 (0.01)	0.56 (0.01)	0.58 (0.01)	0.65 (0.01)
$N = 500, \text{density} = 0.5$	0.60 (<0.01)	0.64 (<0.01)	0.76 (<0.01)	0.56 (<0.01)	0.58 (<0.01)	0.65 (<0.01)
$N = 500, \text{density} = 1^*$	0.60	0.64	0.76	0.56	0.58	0.65
$N = 800, \text{density} = 0.1$	0.58 (0.05)	0.62 (0.05)	0.74 (0.05)	0.55 (0.05)	0.57 (0.05)	0.64 (0.05)
$N = 800, \text{density} = 0.5$	0.44 (0.08)	0.48 (0.08)	0.60 (0.08)	0.41 (0.08)	0.43 (0.08)	0.50 (0.08)
$N = 800, \text{density} = 1^*$	0.26	0.30	0.42	0.23	0.26	0.33

Table 1.4: The value of welfare (the sum of probabilities of being healthy in the second period) averaged over 100 random networks (standard errors in parentheses). We use the Greedy Algorithm or the TWNI allocation to determine who in each network should be vaccinated.

* When the density of Erdős-Renyi network equals 1, standard errors are 0 since there is no randomness in the network structure.

Usually, in the literature on treatment assignment, researchers use observational data or experimental data without network structure information to study the optimal policy. As a result, the allocation regime assigns the treatment without considering spillover effects, which could lead to a sub-optimal result. We call this kind of regime Targeting Without Network Information (TWN¹⁰). In this simulation, we want to learn the welfare loss from using TWNI versus our method.

Generally, TWNI assigns treatment based on personal characteristics. In this study, we

¹⁰We assign the vaccine only to the second group (i.e., only older people receive the vaccine).

only have one covariate: *age*. This means either the old group receives the vaccine or the young group receives the vaccine. Under the previous setting (i.e., older people are more likely to be infected and to die), group 2 will consume the entire vaccine allocation. Given different capacity constraints, this assignment rule selects units to be vaccinated from group 2 until the upper bound is reached. Table 1.4 indicates the results for TWNI allocation are similar to those for random allocation. In addition, despite the outcome value varying with the SIR parameters, the sizable improvement from using network information to allocate vaccination is quite robust to variations in the size and density of network.




Figure 1.2: Comparison between Greedy Algorithm and Targeting Without Network Information

Our numerical study shows that if the number of available vaccine doses is small, the loss from ignoring network information is relatively small too (around 4%). This loss increases dramatically, however, with the number of available vaccines. In addition, the performance gap between our greedy algorithm and the other two allocation methods decreases with the network complexity (i.e., the number of edges). Under what might be described as a lockdown policy, the density of the network is maintained at a relatively low level, which raises the cost of ignoring spillovers. This cost also increases with the number of units in the population, which is a problem in a more realistic setting. The performance improvement from considering network information is robust to variation of the SIR parameters, and an

allocation rule which ignores spillovers waste a sizeable proportion of a scarce resource.

In Table 1.5, we illustrate the impact on the percentage of vaccinated younger units by varying the weight choice g_i (In this simulation exercise, we choose equal weight for the units in same group). If we assign weight $g_1 = 1.5$ for G_1 , we find all the vaccines are consumed by younger units. Comparing with the equal weight case, this number changes dramatically. Moreover, we find our greedy algorithm offers more vaccines to younger units in the case of parameter set 2 than parameter set 1, i.e., when the transmission rate parameters are higher within and across the groups.

Weight Choice Capacity Constraint	Weight $g_1 = 1, g_2 = 1$		Weight $g_1 = 1.1, g_2 = 1$		Weight $g_1 = 1.5, g_2 = 1$	
	$d = 7\%N$	$d = 20\%N$	$d = 7\%N$	$d = 20\%N$	$d = 7\%N$	$d = 20\%N$
Parameter set 1						
$N = 500, \text{density} = 0.1$	9%	17%	80%	61%	100%	100%
$N = 500, \text{density} = 0.5$	0%	1%	100%	94%	100%	100%
$N = 500, \text{density} = 1$	9%	17%	80%	61%	100%	100%
$N = 800, \text{density} = 0.1$	11%	14%	84%	69%	100%	100%
$N = 800, \text{density} = 0.5$	0%	1%	100%	95%	100%	100%
$N = 800, \text{density} = 1$	0%	0%	100%	100%	100%	100%
Parameter set 2						
$N = 500, \text{density} = 0.1$	31%	27%	89%	64%	100%	100%
$N = 500, \text{density} = 0.5$	3%	13%	100%	95%	100%	100%
$N = 500, \text{density} = 1$	31%	27%	89%	64%	100%	100%
$N = 800, \text{density} = 0.1$	23%	26%	91%	71%	100%	100%
$N = 800, \text{density} = 0.5$	5%	10%	100%	98%	100%	100%
$N = 800, \text{density} = 1$	0%	0%	100%	100%	100%	100%

Table 1.5: The percentage of vaccinated younger units in the second period under the vaccine allocation policies obtained by Greedy Algorithm, averaged over 100 random networks. We choose three different sets of weights in this comparison

1.7 Conclusion

In this work, we have introduced a novel method to estimate individualized vaccine allocation rules under network interference. We introduce the heterogeneous-interacted-SIR model to specify the spillover effects of infectious disease. We show that the welfare objective function of the vaccine allocation problem is non-decreasing and submodular, and so is its empirical

analogue formed by plugging in the estimates of the SIR parameters. Based on this specific diminishing returns property, we provide a greedy algorithm with performance guarantee under two different exogenous constraints, which can easily accommodate various targets that policymakers commonly face in reality. Moreover, we show that this algorithm implies an upper bound for regret that converges uniformly at $\mathcal{O}(n^{-1/2})$. Using simulation, we point out the importance of considering network information in the allocation problem.

Several open questions and extensions are worth considering in future work. First, this paper considered a one-time vaccine allocation. We did not consider if there are multiple allocation periods, and how to decide the allocation dynamically. A relevant important question is how to jointly optimize allocations and timing of first- and second-doses of vaccines, as recently discussed for Covid-19 vaccines in [Maier et al. \(2021\)](#), [Tuite et al. \(2021\)](#), and [Wang et al. \(2021\)](#)). Moreover, we do not study how the vaccine allocation rule impacts on the outcome variables after multiple periods. As discussed in [Bu et al. \(2020\)](#), changes to the network structure should be considered in a dynamic setting. Second, we only compare the greedy algorithm with the brute-force optimum in a small network. Other than the universal bounds of Theorem 1.5.1, we do not know the performance of our method relative to the optimal solution in the large network data setting. Third, we did not impose any other constraints than the capacity constraint and the targeting constraint. For interpretability and fairness, we may want to additionally restrict the policy rule as a simple function of observed covariates. We regard these as interesting questions that are worthy of consideration.

Supplement to “ Who Should Get Vaccinated? Individualized Allocation of Vaccines Over SIR Network”

A.1 The Transmission Term

Consider a susceptible individual i with κ_s contacts which depends on his own characteristics at each period. Of these contacts, a fraction $\sum_{j \in N_i} I_j(1 - v_j)a_j/|N_i|$ are contacts with infected neighbors from group 1, and a fraction $\sum_{j \in N_i} I_j(1 - v_j)b_j/|N_i|$ are contacts with infected neighbors from group 2. If we define c_{ij} as the probability of successful disease transmission at each contact, then $1 - c_{sk}$ is the probability that transmission between group s and group k does not take place. Therefore, we have the probability that a unit i is not infected in one time period:

$$1 - q_i = (1 - c_{11})^{\frac{\kappa_1 \sum_{j \in N_i} I_j(1-v_j)a_j a_i}{|N_i|}} \cdot (1 - c_{12})^{\frac{\kappa_1 \sum_{j \in N_i} I_j(1-v_j)b_j a_i}{|N_i|}} \\ \cdot (1 - c_{21})^{\frac{\kappa_2 \sum_{j \in N_i} I_j(1-v_j)a_j b_i}{|N_i|}} \cdot (1 - c_{22})^{\frac{\kappa_2 \sum_{j \in N_i} I_j(1-v_j)b_j b_i}{|N_i|}}.$$

We now define $\beta_{sk} = -\kappa_s \ln(1 - c_{sk})$ and plug it into the expression for $1 - q_i$, which allows us to rewrite the above equation as:

$$q_i = 1 - e^{-z},$$

where

$$\begin{aligned} z = & \frac{\beta_{11}}{|N_i|} \sum_{j \in N_i} I_j(1 - v_j) a_j a_i + \frac{\beta_{12}}{|N_i|} \sum_{j \in N_i} I_j(1 - v_j) b_j a_i \\ & + \frac{\beta_{21}}{|N_i|} \sum_{j \in N_i} I_j(1 - v_j) a_j b_i + \frac{\beta_{22}}{|N_i|} \sum_{j \in N_i} I_j(1 - v_j) b_j b_i. \end{aligned}$$

Recalling that $e^x = 1 + x + \frac{x^2}{2!} + \frac{x^3}{3!} + \dots$, we now have the probability of infection at each time period is

$$q_i \simeq z.$$

A.2 Lemmas

A.2.1 Preliminary Lemma

In this section we collect a set of lemmas from past literature that we use in our proofs.

Lemma A.2.1 (Proposition 6.3 [Bach \(2011\)](#)). *Let $Q \in \mathbb{R}^{p \times p}$, $q \in \mathbb{R}^p$, and $\mathcal{N} = \{1, 2, \dots, p\}$. For $A \in 2^{\mathcal{N}}$, define $1_A = (1_{1 \in A}, \dots, 1_{p \in A})'$. The function $F: A \mapsto q^\top 1_A + \frac{1}{2} 1_A^\top Q 1_A$ is submodular if and only if all off-diagonal elements of Q are non-positive.*

Lemma A.2.2 (Theorem 2.2 [Cunningham \(1985\)](#)). *Function F is a cut function if and only if: For any three disjoint subsets A, B, C of S ,*

$$F(A \cup B \cup C) = F(A \cup B) + F(A \cup C) + F(B \cup C) - F(A) - F(B) - F(C) + F(\emptyset).$$

The following lemmas are some common techniques that are often used in the statistical learning literature, as reviewed in [Lugosi \(2002\)](#).

Lemma A.2.3 (Hoeffding's inequality [Hoeffding \(1963\)](#)). *Let X_1, \dots, X_n be independent bounded random variables such that X_i falls in the interval $[a_i, b_i]$ with probability one. Denote their sum by $S_n = \sum_{i=1}^n X_i$. Then for any $\epsilon > 0$ we have*

$$\mathbb{P}\{S_n - \mathbb{E} S_n \geq \epsilon\} \leq e^{-2\epsilon^2 / \sum_{i=1}^n (b_i - a_i)^2},$$

and

$$\mathbb{P}\{S_n - \mathbb{E} S_n \leq -\epsilon\} \leq e^{-2\epsilon^2 / \sum_{i=1}^n (b_i - a_i)^2}.$$

A.2.2 Proof of Lemma 1.4.1

Let $\hat{W} \in \mathbb{R}^{N \times N}$ and $\hat{C} \in \mathbb{R}^N$. Then the function $F_n: V \mapsto \mathbf{v}^\top \hat{W} \mathbf{v} + \hat{C}^\top \mathbf{v} - \mathbf{1}_{N \times 1}^\top \hat{W} \mathbf{v} - \mathbf{v}^\top \hat{W} \mathbf{1}_{N \times 1}$ is submodular if and only if $\hat{w}_{ij} \leq 0 \forall i \neq j$.

Proof. The first step is to show our objective function is a cut function based on Lemma [A.2.2](#).

In our case, simply consider three arbitrary disjoint sets $A, B, C \subseteq \mathcal{N}$.

$$F_n(A) = \sum_{j \in A} \{\hat{w}_{jj} + \hat{c}_j\} + \sum_{i \neq j \in A} \{\hat{w}_{ij}\} - \sum_{j \in A} \sum_{m=1}^N \{\hat{w}_{mj} + \hat{w}_{jm}\},$$

$$F_n(B) = \sum_{j \in B} \{\hat{w}_{jj} + \hat{c}_j\} + \sum_{i \neq j \in B} \{\hat{w}_{ij}\} - \sum_{j \in B} \sum_{m=1}^N \{\hat{w}_{mj} + \hat{w}_{jm}\},$$

$$\begin{aligned}
F_n(\{A \cup B\}) &= \sum_{j \in \{A \cup B\}} \{\hat{w}_{jj} + \hat{c}_j\} + \sum_{i \neq j \in \{A \cup B\}} \{\hat{w}_{ij}\} - \sum_{j \in \{A \cup B\}} \sum_{m=1}^N \{\hat{w}_{mj} + \hat{w}_{jm}\} \\
&= \sum_{j \in A} \{\hat{w}_{jj} + \hat{c}_j\} + \sum_{i \neq j \in A} \{\hat{w}_{ij}\} - \sum_{j \in A} \sum_{m=1}^N \{\hat{w}_{mj} + \hat{w}_{jm}\} \\
&\quad + \sum_{j \in B} \{\hat{w}_{jj} + \hat{c}_j\} + \sum_{i \neq j \in B} \{\hat{w}_{ij}\} - \sum_{j \in B} \sum_{m=1}^N \{\hat{w}_{mj} + \hat{w}_{jm}\} \\
&= F_n(A) + F_n(B).
\end{aligned}$$

Therefore, we have:

$$\begin{aligned}
F_n(\{A \cup B\} \cup C) &= F_n(\{A \cup B\}) + F_n(C) \\
&= F_n(A) + F_n(B) + F_n(C)
\end{aligned}$$

Combining the previous results, we get:

$$\begin{aligned}
F_n(\{A \cup B\} \cup C) &= F_n(A \cup B) + F_n(A \cup C) + F_n(B \cup C) \\
&\quad - F_n(A) - F_n(B) - F_n(C) - F_n(\emptyset),
\end{aligned}$$

since $F_n(\emptyset) = 0$.

Now, we have shown that $F_n(V)$ is a cut function. The next step is to find the sufficient and necessary conditions for submodularity of the cut function. Lemma A.2.1 indicates, for any cut function which can be written as a quadratic function plus a linear part, submodularity holds if and only if all off-diagonal elements of the weighting matrix are non-positive. That requires $\hat{w}_{ij} \leq 0, \forall i \neq j$. \square

A.2.3 Proof of Lemma 1.5.1

Under Assumption 1, 2, and 3, we have

$$\mathbb{E}_{P^n} |\hat{w}_{ij} - w_{ij}| \leq \sqrt{\frac{1 + \ln(2)}{2n}} \frac{A_{ij} g_i}{N}, \quad \mathbb{E}_{P^n} |\hat{c}_i - c_i| \leq \sqrt{\frac{1 + \ln(2)}{2n}} \frac{I_i g_i}{N}.$$

Proof. We first prove the upper bound of $\mathbb{E}_{P^n} |\hat{w}_{ij} - w_{ij}|$.

$$\hat{w}_{ij} - w_{ij} = \frac{S_i g_i A_{ij} I_j}{|N_i| N} \left[(\beta_{11} - \hat{\beta}_{11}) a_i a_j + (\beta_{12} - \hat{\beta}_{12}) a_i b_j + (\beta_{21} - \hat{\beta}_{21}) b_i a_j + (\beta_{22} - \hat{\beta}_{22}) b_i b_j \right].$$

If we take the absolute value and expectation of each side, by the triangle inequality, we get

$$\begin{aligned} \mathbb{E}_{P^n} |\hat{w}_{ij} - w_{ij}| &= \mathbb{E}_{P^n} \left| \frac{S_i g_i A_{ij} I_j}{|N_i| N} \left[(\hat{\beta}_{11} - \beta_{11}) a_i a_j + (\hat{\beta}_{12} - \beta_{12}) a_i b_j \right. \right. \\ &\quad \left. \left. + (\hat{\beta}_{21} - \beta_{21}) b_i a_j + (\hat{\beta}_{22} - \beta_{22}) b_i b_j \right] \right| \\ &\leq \frac{S_i g_i A_{ij} I_j a_i a_j}{|N_i| N} \mathbb{E}_{P^n} |\hat{\beta}_{11} - \beta_{11}| + \frac{S_i g_i A_{ij} I_j a_i b_j}{|N_i| N} \mathbb{E}_{P^n} |\hat{\beta}_{12} - \beta_{12}| \\ &\quad + \frac{S_i g_i A_{ij} I_j b_i a_j}{|N_i| N} \mathbb{E}_{P^n} |\hat{\beta}_{21} - \beta_{21}| + \frac{S_i g_i A_{ij} I_j b_i b_j}{|N_i| N} \mathbb{E}_{P^n} |\hat{\beta}_{22} - \beta_{22}|. \end{aligned} \tag{A.11}$$

Since β_{sk} is the effective contact rate of the disease between group s and k , it is naturally bounded in $[0, 1]$. We can apply Lemma A.2.3 to get the upper bound of each component:

$$\mathbb{P} \left\{ |\hat{\beta}_{sk} - \beta_{sk}| \geq \epsilon \right\} \leq 2e^{-2n\epsilon^2} \quad \forall s, k = 1, 2. \tag{A.12}$$

Now we can bound $\mathbb{E}(|\hat{\beta} - \beta|)$. Recall that for any nonnegative random variable Y , $\mathbb{E}(Y) =$

$\int_0^\infty \mathbb{P}(Y \geq t)dt$. Hence, for any $a > 0$,

$$\begin{aligned}\mathbb{E}(|\hat{\beta} - \beta|^2) &= \int_0^\infty \mathbb{P}(|\hat{\beta} - \beta|^2 \geq t)dt \\ &= \int_0^a \mathbb{P}(|\hat{\beta} - \beta|^2 \geq t)dt + \int_a^\infty \mathbb{P}(|\hat{\beta} - \beta|^2 \geq t)dt \\ &\leq a + \int_a^\infty \mathbb{P}(|\hat{\beta} - \beta|^2 \geq t)dt.\end{aligned}$$

Equation (A.12) implies that $\mathbb{P}(|\hat{\beta} - \beta| \geq \sqrt{t}) \leq 2e^{-2nt}$. Hence,

$$\begin{aligned}\mathbb{E}(|\hat{\beta} - \beta|^2) &\leq a + \int_a^\infty \mathbb{P}(|\hat{\beta} - \beta|^2 \geq t)dt \\ &= a + \int_a^\infty \mathbb{P}(|\hat{\beta} - \beta| \geq \sqrt{t})dt \\ &\leq a + 2 \int_a^\infty e^{-2nt}dt \\ &= a + \frac{e^{-2na}}{n}.\end{aligned}$$

Set $a = \ln(2)/(2n)$ and we have

$$\mathbb{E}(|\hat{\beta} - \beta|^2) \leq \frac{\ln(2)}{2n} + \frac{1}{2n} = \frac{1 + \ln(2)}{2n}.$$

Therefore, we have

$$\mathbb{E}(|\hat{\beta} - \beta|) \leq \sqrt{\mathbb{E}(|\hat{\beta} - \beta|^2)} \leq \sqrt{\frac{1 + \ln(2)}{2n}}.$$

Plugging this upper bound back to equation (A.11), we get

$$\begin{aligned}\mathbb{E}|\hat{w}_{ij} - w_{ij}| &\leq \sqrt{\frac{1 + \ln(2)}{2n}} (a_i S_i A_{ij} a_j I_j + a_i S_i A_{ij} b_j I_j + b_i S_i A_{ij} a_j I_j + b_i S_i A_{ij} b_j I_j) \frac{g_i}{|N_i| N} \\ &\quad (\because |N_i| \geq 1 \text{ and by treating 0 neighbor as equal to 1}) \\ &\leq \sqrt{\frac{1 + \ln(2)}{2n}} \frac{A_{ij} g_i}{N}.\end{aligned}$$

The steps to prove the upper bound for $E_{P^n}|\hat{c}_i - c_i|$ are exactly the same,

$$\hat{c}_i - c_i = -(\hat{\gamma}_1 - \gamma_1) \frac{a_i I_i g_i}{N} - (\hat{\gamma}_2 - \gamma_2) \frac{b_i I_i g_i}{N}.$$

Take the absolute value and expectation of both sides,

$$\begin{aligned} \mathbb{E}|\hat{c}_i - c_i| &= \mathbb{E} \left| (\hat{\gamma}_1 - \gamma_1) \frac{a_i I_i g_i}{N} + (\hat{\gamma}_2 - \gamma_2) \frac{b_i I_i g_i}{N} \right| \\ &\leq \mathbb{E}|\hat{\gamma}_1 - \gamma_1| \frac{a_i I_i g_i}{N} + \mathbb{E}|\hat{\gamma}_2 - \gamma_2| \frac{b_i I_i g_i}{N}. \end{aligned} \tag{A.13}$$

With the same idea as for β , γ is also bounded in $[0, 1]$. By using Lemma A.2.3, we get

$$\mathbb{E}|\hat{\gamma} - \gamma| \leq \sqrt{\frac{1 + \ln(2)}{2n}} \quad \forall \hat{\gamma} = \hat{\gamma}_1, \hat{\gamma}_2.$$

Plugging this upper bound back into equation A.13, we get

$$\begin{aligned} \mathbb{E}|\hat{c}_i - c_i| &\leq \sqrt{\frac{1 + \ln(2)}{2n}} (a_i + b_i) \frac{I_i g_i}{N} \\ &= \sqrt{\frac{1 + \ln(2)}{2n}} \frac{I_i g_i}{N}. \end{aligned}$$

□

A.3 Proofs for Theorems

A.3.1 Proof of Theorem 1.4.1

The objective function $F_n(V)$ is a non-decreasing submodular function for any adjacency matrix, covariate values, and parameter estimates.

Proof. Recall

$$F_n(V) = \mathbf{v}^\top \hat{W} \mathbf{v} + \hat{C}^\top \mathbf{v} - \mathbf{1}_{N \times 1}^\top \hat{W} \mathbf{v} - \mathbf{v}^\top \hat{W} \mathbf{1}_{N \times 1}.$$

Here, \mathbf{v} is a vector of integers. Let us first, instead, look at this function in the continuous case. Imagine now we have a vector $\tilde{\mathbf{v}}$ with continuous elements. Then, this function becomes:

$$\tilde{F}_n(V) = \tilde{\mathbf{v}}^\top \hat{W} \tilde{\mathbf{v}} + \hat{C}^\top \tilde{\mathbf{v}} - \mathbf{1}_{N \times 1}^\top \hat{W} \tilde{\mathbf{v}} - \tilde{\mathbf{v}}^\top \hat{W} \mathbf{1}_{N \times 1}.$$

We can write the derivative of $F_n(V)$ with respect to $\tilde{\mathbf{v}}$:

$$\begin{aligned} \frac{\partial \tilde{F}_n(V)}{\partial \tilde{\mathbf{v}}} &= \tilde{\mathbf{v}}^\top \hat{W}^\top + \tilde{\mathbf{v}}^\top \hat{W} + \hat{C}^\top - \mathbf{1}_{N \times 1}^\top \hat{W} - \mathbf{1}_{N \times 1}^\top \hat{W}^\top \\ &= \underbrace{(\tilde{\mathbf{v}}^\top - \mathbf{1}_{N \times 1}^\top)}_{\leq 0} \underbrace{\hat{W}^\top}_{\leq 0} + \underbrace{(\tilde{\mathbf{v}}^\top - \mathbf{1}_{N \times 1}^\top)}_{\leq 0} \underbrace{\hat{W}}_{\leq 0} + \underbrace{\hat{C}^\top}_{\geq 0} \geq 0. \end{aligned}$$

Given all the elements in $\frac{\partial \tilde{F}_n(V)}{\partial \tilde{\mathbf{v}}}$ are non-negative, this non-decreasing property also holds under the integer increment in every element of \mathbf{v} . Therefore, $F_n(V)$ is a non-decreasing set function. Combining this with Lemma 1.4.1, we complete the proof. \square

A.3.2 Proof of Theorem 1.5.1

Let $N_M = \max_{i \in \mathcal{N}} |N_i|$, N_I be the total number of infected units, and $g = \max_{i \in \mathcal{N}} g_i$. Under Assumptions 1, 2, and 3, we have

$$\mathbb{E}_{P^n} \left[F(V^*) - F(\hat{V}) \right] \leq \bar{C} \cdot \frac{g[d \min\{N_M, d\} + 2dN_M + \min\{N_I, d\}]}{N} \sqrt{\frac{1}{n} + \frac{1}{e}} F(V^*),$$

where \bar{C} is a universal constant and d is the number of available vaccine doses.

Proof.

$$\begin{aligned}
\mathbb{E}_{P^n} \left[\sup_{V \in \mathcal{V}_d} |F_n(V) - F(V)| \right] &\leq \mathbb{E}_{P^n} \left[\sup_{V \in \mathcal{V}_d} \mathbf{v}^\top |\hat{W} - W| \mathbf{v} \right] + \mathbb{E}_{P^n} \left[\sup_{V \in \mathcal{V}_d} |\hat{C}^\top - C^\top| \mathbf{v} \right] \\
&\quad + \mathbb{E}_{P^n} \left[\sup_{V \in \mathcal{V}_d} \mathbf{1}_{N \times 1}^\top |\hat{W} - W| \mathbf{v} \right] + \mathbb{E}_{P^n} \left[\sup_{V \in \mathcal{V}_d} \mathbf{v}^\top |\hat{W} - W| \mathbf{1}_{N \times 1} \right] \\
&= \sup_{V \in \mathcal{V}_d} \mathbf{v}^\top \mathbb{E}_{P^n} |\hat{W} - W| \mathbf{v} + \sup_{V \in \mathcal{V}_d} \mathbb{E}_{P^n} |\hat{C}^\top - C^\top| \mathbf{v} \\
&\quad + \sup_{V \in \mathcal{V}_d} \mathbf{1}_{N \times 1}^\top \mathbb{E}_{P^n} |\hat{W} - W| \mathbf{v} + \sup_{V \in \mathcal{V}_d} \mathbf{v}^\top \mathbb{E}_{P^n} |\hat{W} - W| \mathbf{1}_{N \times 1}.
\end{aligned} \tag{A.14}$$

From equation (A.14), $\mathbb{E}_{P^n} \left[\sup_{V \in \mathcal{V}_d} |F_n(V) - F(V)| \right]$ can be decomposed into four components. Since \mathbf{v} only contains $\{0, 1\}$ and the absolute value must be non-negative, V that maximizes each component under capacity constraint must select units with a greater number of edges, as compare to those that are not selected. We define the maximum number of edges for each unit as N_M . Hence, the number of edges for selected units must be lower or equal to N_M . Next, we look at each term in equation A.14 separately. Using Lemma 1.5.1, the first term is bounded as:

$$\begin{aligned}
\sup_{V \in \mathcal{V}_d} \mathbf{v}^\top \mathbb{E}_{P^n} |\hat{W} - W| \mathbf{v} &\leq \sup_{V \in \mathcal{V}_d} \sqrt{\frac{1 + \ln(2)}{2n}} \sum_{i \in V} \sum_{j \in V} \frac{A_{ij} g_i}{N} \\
&\leq \sqrt{\frac{1 + \ln(2)}{2n}} \sum_{i \in V} \frac{g_i \min\{N_M, d\}}{N} \\
&\quad \left(\because \sum_{j \in V} A_{ij} \leq \min\{N_M, d\} \quad \forall i \in \mathcal{N} \right) \\
&\leq \sqrt{\frac{1 + \ln(2)}{2n}} \frac{dg \cdot \min\{N_M, d\}}{N}.
\end{aligned}$$

The second term is bounded as:

$$\sup_{V \in \mathcal{V}_d} \mathbb{E}_{P^n} |\hat{C}^\top - C^\top| \mathbf{v} \leq \sup_{V \in \mathcal{V}_d} \sqrt{\frac{1 + \ln(2)}{2n}} \sum_{i \in V} \frac{I_i g_i}{N} \leq \sqrt{\frac{1 + \ln(2)}{2n}} \frac{g \min\{N_I, d\}}{N}.$$

The third term is bounded as:

$$\begin{aligned}
\sup_{V \in \mathcal{V}_d} \mathbf{1}_{N \times 1}^\top \mathbb{E}_{P^n} \left| \hat{W} - W \right| \mathbf{v} &\leq \sup_{V \in \mathcal{V}_d} \sqrt{\frac{1 + \ln(2)}{2n}} \sum_{i \in \mathcal{N}} \sum_{j \in V} \frac{A_{ij} g_i}{N} \\
&\leq \sqrt{\frac{1 + \ln(2)}{2n}} \frac{\sum_{j \in V} g N_M}{N} \\
&= \sqrt{\frac{1 + \ln(2)}{2n}} \frac{dg N_M}{N}.
\end{aligned}$$

The fourth term is bounded as:

$$\begin{aligned}
\sup_{V \in \mathcal{V}_d} \mathbf{v}^\top \mathbb{E}_{P^n} \left| \hat{W} - W \right| \mathbf{1}_{N \times 1} &\leq \sup_{V \in \mathcal{V}_d} \sqrt{\frac{1 + \ln(2)}{2n}} \sum_{i \in V} \sum_{j \in \mathcal{N}} \frac{A_{ij} g_i}{N} \\
&\leq \sqrt{\frac{1 + \ln(2)}{2n}} \frac{\sum_{i \in V} g N_M}{N} \\
&= \sqrt{\frac{1 + \ln(2)}{2n}} \frac{dg N_M}{N}.
\end{aligned}$$

Combining the bounds of the four terms, we get

$$\begin{aligned}
\mathbb{E}_{P^n} \left[\sup_{V \in \mathcal{V}_d} |F_n(V) - F(V)| \right] &\leq \sqrt{\frac{1 + \ln(2)}{2n}} \frac{dg \min\{N_M, d\}}{N} + 2\sqrt{\frac{1 + \ln(2)}{2n}} \frac{dg N_M}{N} \\
&\quad + \sqrt{\frac{1 + \ln(2)}{2n}} \frac{g \min\{N_I, d\}}{N} \\
&= \frac{dg \min\{N_M, d\} + 2dg N_M + g \min\{N_I, d\}}{N} \sqrt{\frac{1 + \ln(2)}{2n}}.
\end{aligned}$$

Therefore, we have from equation (1.9)

$$\begin{aligned}
\mathbb{E}_{P^n} [F(V^*) - F(\hat{V})] &\leq \left(2 + \frac{1}{e}\right) \frac{dg \min\{N_M, d\} + 2dg N_M + g \min\{N_I, d\}}{N} \sqrt{\frac{1 + \ln(2)}{2n}} \\
&\quad + \frac{1}{e} F(V^*) \\
&= \left(2 + \frac{1}{e}\right) \frac{g[d \min\{N_M, d\} + 2d N_M + \min\{N_I, d\}]}{N} \sqrt{\frac{1 + \ln(2)}{2n}} \\
&\quad + \frac{1}{e} F(V^*)
\end{aligned}$$

Setting $\bar{C} = (2 + 1/e)\sqrt{\frac{1+\ln(2)}{2}}$ completes the proof.

□

Chapter 2

Individualized Treatment Allocation in Sequential Network Games

2.1 Introduction

The question of how best to allocate treatment to units interacting in a network is relevant to many policy areas, including the provision of local public goods ([Bramoullé and Kranton, 2007](#)), the diffusion of microfinance ([Banerjee et al. \(2013\)](#); and [Akbarpour et al. \(2020b\)](#)), and strategic immunization ([Galeotti and Rogers \(2013\)](#); and [Kitagawa and Wang \(2023b\)](#)). Obtaining an optimal individualized allocation, however, is often infeasible due to analytical and computational challenges. As a consequence, practical counterfactual policy analysis in the presence of network spillovers is limited to simulating and comparing outcome distributions or welfare values across a few benchmark candidate policies. This leaves the magnitude of the potential welfare gains of an optimal individualized assignment policy unknown.

Focusing on a class of social network models in which interacting agents play sequential decision games ([Jackson and Watts \(2002\)](#); [Nakajima \(2007\)](#); [Mele \(2017\)](#); and [Christakis et al. \(2020\)](#)), this paper develops a method to obtain optimal treatment assignment rules that maximize a social welfare criterion. We consider an individualized allocation of binary

treatments over agents who are heterogeneous in terms of their own observable characteristics, their network configurations, and their neighbors' observable characteristics. Each agent chooses a binary outcome so as to maximize their own utility. This choice depends upon the agent's own characteristics and treatment as well as their neighbors' characteristics, treatments and choices. The sequential decisions of randomly ordered agents induce a unique stationary distribution of choices (Nakajima, 2007; Mele, 2017). We specify the planner's welfare criterion to be the mean of the aggregate outcomes (i.e., the sum of the means of binary outcomes over all agents in the network) at the stationary state that is associated with a given treatment allocation. We aim to maximize the welfare evaluated at the stationary outcome distribution with respect to the individualized allocation of treatments.

There are analytical and computational challenges to solving the maximization problem for optimal targeting. First, fixing an allocation of treatments, the sequential decision games induce a Markov random field (MRF) and the stationary outcome distribution has a Gibbs distribution representation. The analytical properties of the mean of the aggregate outcomes, however, are difficult to characterize. To approximate the joint distribution of outcomes, the literature on MRFs performs numerical methods such as Markov Chain Monte Carlo (MCMC) (Geman and Geman, 1984). If the size of the network is moderate to large though, MCMC can be slow to converge. It is, therefore, infeasible to perform MCMC to evaluate the welfare at every candidate treatment assignment policy. Second, obtaining an optimal individualized assignment is a combinatorial optimization problem with respect to a binary vector whose cardinality is equal to the number of agents in the network. A brute force search quickly becomes infeasible as the size of network expands.

We tackle these challenges by proposing methods for approximately solving the combinatorial optimization problem for individualized assignment. Our proposal is to perform a variational approximation of the stationary distribution of outcomes and to optimize the approximated equilibrium welfare with respect to an assignment vector by using a greedy optimization algorithm. The variational approximation step lessens the computational bur-

den of running MCMC at each candidate policy. The greedy optimization step reduces the computational burden of the combinatorial optimization by assigning treatment sequentially to the agent who generates the largest welfare gain given previous assignments. Since our proposal involves approximation of the objective function and the heuristic method of greedy optimization, our proposal is not guaranteed to lead to a global optimum. A novel contribution of this paper is that we derive a performance guarantee for our proposed method in terms of an analytical upper bound on the welfare loss relative to the globally optimal assignment. The upper bound on the welfare regret consists of two terms: the welfare loss due to variational approximation and the welfare loss due to greedy optimization. We show that, once scaled such that it can be interpreted as the per-person welfare loss, the first term of the upper bound on the welfare loss (originating from the variational approximation) vanishes asymptotically as the number of agents in the network increases. On the other hand, we show that the second term of the upper bound on the welfare loss (originating from the greedy optimization) does not generally converge to zero.

To highlight this paper’s unique contributions to the literature, we abstract from estimation of the structural parameters underlying the sequential decision game and assume that they are known. See [Geyer and Thompson \(1992\)](#), [Snijders et al. \(2002\)](#), [Wainwright et al. \(2008\)](#), [Chatterjee and Diaconis \(2013\)](#), [Mele \(2017\)](#), [Boucher and Mourifié \(2017\)](#), and [Mele and Zhu \(2022\)](#) for identification and estimation of these parameters. In practical terms, our proposed method is useful for computing an optimal assignment of treatment, with point estimates plugged-in in place of structural parameters.

To assess the performance of our proposal, we perform extensive numerical studies. Given that the number of possible configurations increases exponentially with the size of the network, we are only able to apply the brute force method to search for the optimal allocation rule in a small network setting. We find that our proposed method leads to a globally optimal solution in a small network setting under our assumptions. In a large network setting, we first examine the performance of using our method compared with *No treatment* and with *ran-*

dom allocation. We attain a welfare improvement of around 50% relative to a No treatment rule, and an improvement of around 10% relative to a random allocation rule. In addition, we evaluate the welfare performance gap of using variational approximation compared with MCMC to approximate the stationary distribution of the outcome. Under our assumptions, the variational approximation performs as well as MCMC for all of the treatment allocation rules that we consider.

We augment these numerical studies with an empirical application that illustrates the implementation and welfare performance of our method. Specifically, we apply our procedure to Indian microfinance data that has been previously analyzed by [Banerjee et al. \(2013\)](#). The data contains information about households in a number of villages, their relation to other households in their village, and whether they chose to purchase a microfinance product. For each village in the sample, we estimate the structural parameters of an assumed utility function using the method outlined in [Snijders et al. \(2002\)](#). Plugging in the parameter estimates, we obtain an individualized treatment allocation rule using our greedy algorithm. We compare the village welfare attained by our greedy algorithm with the welfare achieved by an NGO called Bharatha Swamukti Samsthe (BSS). For all 43 villages in the sample, our method is associated with a higher welfare-level, with the magnitude of improvement varying from 9.82% to 137.46% (average improvement is 40.69%). The magnitude of the welfare gain is remarkable and demonstrates the benefits of individualized targeting under interference.

The remainder of this paper is organized as follows. We first review the relevant literature in the remainder of this section. Section [2.2](#) details the sequential decision process and the stationary distribution of the outcome variable. Section [2.3](#) contains theoretical results relating to the implementation of a variational approximation and to the maximization of the variationally approximated outcome. Simulation results are shown in Section [2.4](#). We apply our proposed method to the Indian microfinance data, which is studied by [Banerjee et al. \(2013\)](#), and demonstrate its performance in Section [2.5](#). Section [2.6](#) concludes. All proofs and derivations are shown in Appendix [A.1](#) to [A.2](#).

2.1.1 Literature Review

This paper intersects with several literatures in economics and econometrics, including graphical game analysis, Markov random fields and variational approximation, discrete optimization of non-submodular functions, and statistical treatment rules.

Graphical game analysis has a long history in economics, see [Rosenthal \(1973\)](#), [Kakade et al. \(2003\)](#), [Ballester et al. \(2006\)](#), [Roughgarden \(2010\)](#), [Kearns et al. \(2013\)](#), [Babichenko and Tamuz \(2016\)](#), [De Paula et al. \(2018\)](#), [Leung \(2020\)](#) and [Parise and Asuman \(2023\)](#). The most relevant paper to our work are [Mele \(2017\)](#) and [Christakis et al. \(2020\)](#), which study strategic network formation. Despite both of them consider the sequential network formation process, there exists some variations in their meeting technology. [Mele \(2017\)](#) assumes units can revise their actions frequently and can meet multiple times, while [Christakis et al. \(2020\)](#) assumes units can only meet once and their actions are permanent. Under their meeting technology, [Mele \(2017\)](#) formulates the network formation game as a potential game ([Monderer and Shapley, 1996](#)), and characterizes the stationary distribution of the network as an exponential distribution. We also formulate our game as a potential game and adopt a similar sequential decision process ([Blume, 1993](#)) to study the stationary distribution for our game. [Jackson and Watts \(2002\)](#) indicates that this sequential decision process is a specific equilibrium selection mechanism. [Kashaev et al. \(2023\)](#) introduces a similar sequential structure into a discrete choice model with peer effect and estimates the model with panel data. [Kline et al. \(2021\)](#) discusses the difficulties and potential solutions of analyzing counterfactuals with multiple equilibria. [Lee and Pakes \(2009\)](#) suggests using the best response dynamics learning model (our sequential decision process) to perform counterfactual analysis with multiple equilibria. [De Paula \(2013\)](#) reviews the recent literature on the econometric analysis of games with multiple equilibria. [Badev \(2021\)](#) extends the setting in [Mele \(2017\)](#) to study how behavioral choices change the network formation. [De Paula \(2020\)](#) reviews recent works on network formation. Whilst [Mele \(2017\)](#) devotes considerable attention to characterizing the stationary distribution of a network formation game, the main focus of

this paper is to develop a method for approximating optimal targeting in network games. [Ballester et al. \(2006\)](#) and [Galeotti et al. \(2020\)](#) also study targeting an intervention in a network setting, but the utility specification, the objective function, and the action space in their paper differ from those in our setting. Those differences lead to a different equilibrium and a different identification strategy for the optimal intervention compared with our setting.

This paper is also relevant to the literature on Markov random fields (MRF) and variational approximation. MRF offer a way to represent the joint distribution of random variables as a collection of conditional distributions. MRF has been used in a wide variety of fields, including in statistical physics (e.g., Ising model; [Ising, 1925](#)), and in image processing ([Wang et al., 2013](#)). We model an individual’s choice of outcomes as the maximization of a latent payoff function that depends upon a treatment allocation and their neighbors’ choices, and derive an MRF representation of the joint distribution of outcomes. We use variational approximation as a computationally tractable approximation of the stationary outcome distribution. See [Wainwright et al. \(2008\)](#) for a comprehensive survey on variational approximation and MRF. [Chatterjee and Dembo \(2016\)](#) provides an approximation error bound to variational approximation applied to MRF of binary outcomes. [Mele and Zhu \(2022\)](#) apply this method to estimate parameters in a network formation model. However, none of the aforementioned papers have considered how changing parameter values affects the mean value or the distribution of MRF. This literature has not, to the best of our knowledge, studied how to obtain an optimal intervention in terms of a criterion function defined on the joint distribution of outcomes characterized through a MRF.

We build a connection to the literature on providing a theoretical performance guarantee to using a greedy algorithm. [Nemhauser et al. \(1978\)](#) provides a performance guarantee for a general greedy algorithm solving submodular maximization problems with a cardinality constraint. Many optimization problems, however, are not submodular ([Krause et al., 2008](#)) and greedy algorithms usually still exhibit good empirical performance ([Das and Kempe, 2011](#)). Given this, there is considerable interest in solving non-submodular optimization

problems using greedy algorithms amongst researchers. See [Das and Kempe \(2011\)](#), [Bian et al. \(2017\)](#), [El Halabi and Jegelka \(2020\)](#), and [Jagalur-Mohan and Marzouk \(2021\)](#). We use the result from [Bian et al. \(2017\)](#) to produce a performance guarantee for our treatment allocation problem by clarifying sufficient conditions for obtaining non-trivial bounds on the submodularity ratio and the curvature of our objective function. We are unaware of these recent advances in the literature on discrete optimization of non-submodular functions being applied elsewhere to the problem of optimal targeting in the presence of network spillovers.

Although it does not introduce sampling uncertainty, this paper shares some motivation with the literature on statistical treatment rules, which was first introduced into econometrics by [Manski \(2004\)](#). Following the pioneering works of [Savage \(1951\)](#), and [Hannan \(1957\)](#), researchers often characterize the performance of decision rules using regret.¹ See [Dehejia \(2005\)](#), [Stoye \(2009, 2012\)](#), [Hirano and Porter \(2009, 2020\)](#), [Chamberlain \(2011, 2020\)](#), [Tetenov \(2012\)](#), and [Christensen et al. \(2022\)](#) for decision theoretic analyses of statistical treatment rules. There is also a growing literature learning on studying individualized treatment assignments including [Kitagawa and Tetenov \(2018\)](#), [Athey and Wager \(2021\)](#), [Kasy and Sautmann \(2021\)](#), [Kitagawa et al. \(2021\)](#), [Mbakop and Tabord-Meehan \(2021\)](#), [Sun \(2021\)](#), and [Adjaho and Christensen \(2022\)](#), among others. These works do not consider settings that allow for the network spillovers of treatments. There are some recent works that introduce network spillovers into statistical treatment choice, such as [Viviano \(2019, 2020\)](#), [Ananth \(2020b\)](#), [Kitagawa and Wang \(2023b\)](#), and [Munro et al. \(2023\)](#). [Viviano \(2019\)](#) and [Ananth \(2020b\)](#) assume the availability of network data from a randomised control trial (RCT) experiment. They do not model the behavior of units from a structural perspective. [Viviano \(2020\)](#) studies how to assign treatments over the social network in an experiment design setting. [Munro et al. \(2023\)](#) studies targeting analysis taking into account spillovers through the market equilibrium. [Kitagawa and Wang \(2023b\)](#) considers the allocation of

¹In econometrics and machine learning, regret typically arises due to uncertainty surrounding the value of underlying parameters (i.e., estimation). In this work, regret arises from our use of variational approximation and of a greedy algorithm (i.e., identification).

vaccines over an epidemiological network model (a Susceptible-Infected-Recovered, or SIR, network). [Kitagawa and Wang \(2023b\)](#) considers a simple two-period transition model and does not consider the long-run stationary distribution of health status over the network. In contrast, in this paper, we consider sequential decision games and aim to optimize the long-run equilibrium welfare by exploiting its MRF representation. As an application of variational approximation to treatment choice in a different context, [Kitagawa et al. \(2022a\)](#) applies variational approximation to a quasi-posterior distribution for individualized treatment assignment policies and studies welfare regret performances when assignment policies are drawn randomly from the variationally approximated posterior.

2.2 Model

2.2.1 Setup

Let $\mathcal{N} = \{1, 2, \dots, N\}$ be the population. Each unit has a K -dimensional vector of observable characteristics that we denote by X_i , $i \in \mathcal{N}$. Assuming that the support of X_i is bounded, we normalize the measurements of X_i to be nonnegative, such that $X_i \in \mathbb{R}_+^K$. Let $\mathcal{X} = \{X_1, \dots, X_N\}$ be a matrix that collects the characteristics of units in the population, and let \mathcal{X}^N denote the set of all possible matrices \mathcal{X} . Let $D = \{d_1, \dots, d_N\}$ denote a vector of treatment allocation, where $d_i \in \{0, 1\}$, $i \in \mathcal{N}$, indicates whether unit i is treated ($d_i = 1$) or untreated ($d_i = 0$).

The social network is represented by an $N \times N$ binary matrix that we denote by $G = \{G_{ij}\}_{i,j \in \mathcal{N}}$, and that is fixed and exogenous in this work. $G_{ij} = 1$ indicates that units i and j are connected in the social network, whilst $G_{ij} = 0$ indicates that they are not. Let \mathcal{N}_i indicate the set of neighbors of unit i . \bar{N} denotes the maximum number of edges for one unit in the network (i.e., $\bar{N} = \max_i |\mathcal{N}_i|$); \underline{N} denotes the minimum number of edges for one unit in the network (i.e., $\underline{N} = \min_i |\mathcal{N}_i|$). As a convention, we assume there are no self-links

(i.e., $G_{ii} = 0, \forall i \in \mathcal{N}$). We further assume that the following property holds for the network structure:

Assumption 4. (Undirected Link) *The adjacency matrix G is undirected. i.e., $G_{ij} = G_{ji}$.*

The symmetric property of interaction in Assumption 4 is a necessary condition for our interacted sequential decision game to be a proper potential game (Definition 2.2.1 below) that can yield a unique stationary outcome distribution. The size of the spillover between units i and j depends not only upon G_{ij} but also upon the treatment allocation and upon covariates, which are allowed to be asymmetric. We, accordingly, have a directed weighted network structure for the spillovers.

As we have previously mentioned, we consider a sequential decision game setting to derive the unique stationary outcome distribution. We now introduce the notation for our sequential decision game. Let $Y_i^t \in \mathcal{Y} = \{0, 1\}$ be unit i 's choice made at time t , which we refer to as i 's outcome. Let Y^t be the collection of outcome variables $\{Y_1^t, \dots, Y_N^t\} \in \mathcal{Y}^N$ at time t . We consider a discrete-time infinite-horizon setting. For each time period t in the decision process, we denote the realization of Y^t by $y^t \in \{0, 1\}^N$, and the realization of unit i 's outcome by y_i^t . The outcome set that includes all of the current outcomes but y_i^t , that is, $y^t \setminus y_i^t$, is denoted by y_{-i}^t . Let $Y = \{Y_i\}_{i=1}^N \in \mathcal{Y}^N$ denote the collection of the outcome variables in equilibrium, which follows the stationary outcome distribution.

The game, which we denote by \mathcal{G} , comprises:

- the aforementioned set of individuals that we label \mathcal{N} , a social planner, and nature;
- a set of actions Y^t that records the binary choice that is made by each individual in every time period t in which they are selected (by nature) to move, and a treatment choice D for each individual that is made by the social planner in the initial period upon observing \mathcal{X} and G but before Y^1 is chosen;
- a player function that selects a single individual to be active in each time period based upon whom nature indicates;

- a sequence of histories over an infinite-horizon that is summarised by an initial treatment allocation and by the identity of the individual that is selected by nature in each time period alongside their corresponding action;
- the preferences (utilities) of individuals $\{U_i(y^t, \mathcal{X}, D, G; \theta)\}_{i=1}^N$, which depend upon both their own and others' actions (i.e., upon each individual's initial treatment allocation and the binary choices that they subsequently make whenever they are selected to do so by nature) and by the social planner's actions (i.e., upon the treatment choice that the social planner makes in the initial period), and that we imbue with certain properties specified in Section 2.2.5;
- the individual selected by nature in each time period receives a pair of preference shocks (one for each of their two choices) before they make a decision. Each individual maximizes their utility at each time period in which they are selected by nature. The social planner chooses the initial treatment allocation to maximize an objective function, which we call the *planner's welfare*.

2.2.2 Potential Game

We consider pure-strategy Nash equilibrium as the solution concept of our game. Recall that the definition of a pure-strategy Nash equilibrium is a set of actions $y^* = \{y_1^*, \dots, y_N^*\}$ such that

$$U_i(y_i^*, y_{-i}^*, \mathcal{X}, D, G; \theta) \geq U_i(y_i', y_{-i}^*, \mathcal{X}, D, G; \theta)$$

for any $y_i' \in \mathcal{Y}$ and for all $i \in \mathcal{N}$. This requires that no individual has a profitable deviation from her current decision when she is randomly selected by nature. To analyze the Nash equilibrium of our game, we characterize our game as a potential game. The concept of a potential game has been used to study strategic interaction since [Rosenthal \(1973\)](#). It provides a tool to analyze the Nash equilibria of (non) cooperative games in various settings

(e.g., [Jackson, 2010](#), and [Bramoullé et al., 2014](#)). We now formally define the potential game.

Definition 2.2.1. (Potential Game ([Monderer and Shapley, 1996](#))) \mathcal{G} is a potential game if there exists a potential function $\Phi : \mathcal{Y}^N \rightarrow \mathbb{R}$ such that for all $i \in \mathcal{N}$ and for all $y_i, y'_i \in \mathcal{Y}$

$$U_i(y_i, y_{-i}) - U_i(y'_i, y_{-i}) = \Phi(y_i, y_{-i}) - \Phi(y'_i, y_{-i}). \quad (2.1)$$

The change in potentials from any player's unilateral deviation matches the change in their payoffs. Nash equilibria, therefore, must be the local maximizers of potential. [Monderer and Shapley \(1996, §Theorem 4.5\)](#) states that

$$\frac{\partial U_i}{\partial y_i \partial y_j} = \frac{\partial U_j}{\partial y_j \partial y_i} \quad (2.2)$$

is a necessary and sufficient condition for a game featuring a twice continuously differentiable utility function to be a potential game. For the discrete outcome case, a condition² – that we refer to as the *symmetry property* – analogous to Eq.2.2 is a necessary and sufficient condition for the existence of a potential function. [Chandrasekhar and Jackson \(2014\)](#), and [Mele \(2017\)](#) also use a potential game framework to analyze Nash equilibria in a network game. We restrict our analysis to potential games equipped with a potential function $\Phi(y, \mathcal{X}, D, G; \theta)$. We later specify a functional form for the utility function that satisfies the symmetry property and provide an explicit functional form for the potential function in Section 2.2.5. In assuming that our game is a potential game, we guarantee that at least one pure strategy Nash equilibrium exists, as per [Monderer and Shapley \(1996\)](#).

²Replacing the second-order derivative in Eq.2.2 with second-order differences. See [Monderer and Shapley \(1996, §Corollary 2.9\)](#) for further details.

2.2.3 Sequential Decision Process

The details of the sequential decision process are as follows. In the initial period, the social planner observes the connections in the social network and individuals' attributes, and decides the treatment allocation so as to maximize the planner's welfare. Then, at the beginning of every period t , an individual i is randomly chosen from \mathcal{N} by nature. Unit i chooses an action (outcome) y_i^t . The *selection process* is a stochastic sequence $O = (O^t)_{t=1}^\infty$ with support \mathcal{N} . Realizations of O^t indicate the unit that makes a decision in period t ; all other units maintain the same choice as in the last period.

The probability of unit i being randomly chosen from \mathcal{N} at time t is given by:

$$\Pr(O^t = i | y^{t-1}, \mathcal{X}, D, G) = \rho_i(y_{-i}^{t-1}),$$

where $\sum_{i=1}^N \rho_i(y_{-i}^{t-1}) = 1$ for all $y \in \{0, 1\}^N$. In the simplest case, $\rho_i(y_{-i}^{t-1}) = 1/N$ for all t . The idea here is that only previously-made choices (outcome) factor into the decision of the unit that is selected by nature in period t . Without this, it is not possible to provide a closed-form expression for the joint distribution of the outcome. We require that any individual can be selected and that this selection depends upon y_{-i}^{t-1} rather than upon y^{t-1} .

Assumption 5. (Decision Process) *The probability of unit i being selected at time t does not depend upon y_i^{t-1} , and each action has a positive probability of occurring:*

$$\rho_i(y_{-i}^{t-1}) = \Pr(O^t = i | y_{-i}^{t-1}, \mathcal{X}, D, G) > 0 \quad \forall i \in \mathcal{N}.$$

Once unit i has been selected in period t , they choose action y_i^t so to maximize their current utility. We assume that there is *complete information*, such that unit i can observe the attributes and treatment status of their neighbors. Before making their decision, unit i

receives an idiosyncratic shock ε . Then, unit i chooses $Y_i^t = 1$ if and only if:

$$U_i(1, y_{-i}^{t-1}, \mathcal{X}, D, G; \boldsymbol{\theta}) + \varepsilon_{1t} \geq U_i(0, y_{-i}^{t-1}, \mathcal{X}, D, G; \boldsymbol{\theta}) + \varepsilon_{0t}.$$

Following the discrete choice literature (e.g., [Train et al., 1987](#); [Brock and Durlauf, 2001](#)), [Mele \(2017\)](#), and [Christakis et al. \(2020\)](#), we put the following assumption about the idiosyncratic shock.

Assumption 6. (Preference Shock) ε_{1t} and ε_{0t} follow a Type 1 extreme value distribution and are independent and identically distributed among units and across time.

Under Assumption 6, the conditional probability of unit i choosing $Y_i^t = 1$ is given by:

$$P(Y_i^t = 1 | Y_{-i}^{t-1} = y_{-i}^{t-1}, \mathcal{X}, D, G; \boldsymbol{\theta}) = \frac{\exp[U_i(1, y_{-i}^{t-1}, \mathcal{X}, D, G; \boldsymbol{\theta})]}{\sum_{y_i \in \{0,1\}} \exp[U_i(y_i, y_{-i}^{t-1}, \mathcal{X}, D, G; \boldsymbol{\theta})]}.$$

Therefore, the sequence $[y^0, y^1, \dots, y^t]$ evolves as a Markov chain such that:

$$y_i^t = \begin{cases} y_i^{t-1} & \text{w/p } 1 - \rho_i(y_{-i}^{t-1}) \\ y & \text{w/p } \rho_i(y_{-i}^{t-1}) \cdot P(Y_i^t = y | Y_{-i}^{t-1} = y_{-i}^{t-1}), \end{cases} \quad \forall i \in \mathcal{N},$$

where $y \in \{0, 1\}$. Under Assumption 4 to 6, this Markov chain is *irreducible* and *aperiodic*,³ which has a unique stationary distribution. Note that in the special case when there is no idiosyncratic shock, the sequence will stay in one Nash equilibrium in the long run.

The individual decision process is a stochastic best response dynamic process ([Blume, 1993](#)). This sequential decision process generates a Markov Chain of decisions. [Jackson and Watts \(2002\)](#) shows that the sequential decision process plays the role of a stochastic equilibrium selection mechanism. Without this sequential structure, the model would be an incomplete model. [Lee and Pakes \(2009\)](#) performs counterfactual predictions of policy

³It is irreducible since every configuration could happen in a finite time given our assumption on the selection process. It is aperiodic since the selected unit has a positive probability to choose the same choice as in the last period.

interventions in the presence of multiple equilibria, with best response dynamics playing the role of an equilibrium selection mechanism.

2.2.4 Stationary Distribution

Following [Mele \(2017, §Theorem 1\)](#), the stationary joint distribution of the outcomes in our sequential decision game is given by:

Theorem 2.2.1. Unique Stationary Distribution ([Nakajima, 2007](#); [Mele, 2017](#)): *Under Assumption 4 to 6, the interacted decision game has a unique stationary distribution:*

$$P[Y = y | \mathcal{X}, D, G; \theta] = \frac{\exp[\Phi(y, \mathcal{X}, D, G; \theta)]}{\sum_{\delta \in \{0,1\}^N} \exp[\Phi(\delta, \mathcal{X}, D, G; \theta)]}. \quad (2.3)$$

[Mele \(2017\)](#) discusses the relationship between Nash equilibria and this stationary distribution. The set of Nash equilibria is the set of local maxima of the potential function. We also know that the probability of a given configuration increases with the value of the potential. Nash equilibria of the game must, therefore, be visited more often in the long run. Given this, modes of the stationary distribution will correspond to Nash equilibria.

Theorem 2.2.1 shows that, given the parametric specification of the distribution of unobservables (Assumption 6), the joint distribution of the outcomes is given by a Gibbs distribution characterized by the potentials. This result has a close connection to the MRF literature. Specifically, we can view the joint distribution of the outcomes in the stationary as a Markov random field (see, e.g., [Brémaud, 2013](#)):

The random field $\{Y_i\}_{i=1}^N$ is a collection of random variables on the state space \mathcal{Y} . This random field is a Markov random field if for all $i \in \mathcal{N}$ and $y \in \mathcal{Y}^N$:

$$P(Y_i = y_i | Y_{-i} = y_{-i}) = P(Y_i = y_i | Y_{j \in \mathcal{N}_i} = y_{j \in \mathcal{N}_i}). \quad (2.4)$$

Given the specification of our utility function, the conditional distribution of Y_i satisfies

this Markov property. By connecting Y to MRF, the *Hammersley-Clifford Theorem* ([Hammersley and Clifford, 1971](#); [Besag, 1974](#)) establishes that the joint distribution of Y must follow a *Gibbs distribution*, which is consistent with the result of Theorem 2.2.1.

The stationary distribution of the outcomes shown in Theorem 2.2.1 is structural in the sense that the specification of the potential function in the Gibbs distribution relies on the functional form specification of the latent payoff function of agents. An advantage of the current structural approach is that we are transparent about the assumptions that we impose on the behavior of agents, on the structure of social interaction, and on the equilibrium concept. The structural approach, accordingly, disciplines the class of joint distributions of observed outcomes to be analyzed. As an alternative to the structural approach, we can consider a reduced-form approach where we model the conditional distribution of the observed outcomes given the treatment vector. Maintaining the family of Gibbs distributions, the reduced-form approach corresponds to introducing a more flexible functional form for the potential functions without guaranteeing that it is supported as a Nash equilibrium of the potential game. Despite this potential issue, our approach of variational approximation and greedy optimization can be used to obtain an optimal targeting rule for a broad class of potential functions.

2.2.5 Preference

As in [Nakajima \(2007\)](#), [Mele \(2017\)](#), and [Sheng \(2020\)](#), we specify the individual utility function as a quadratic function of own and neighbors' choices. This implies that we allow the spillover effects to be at the first degree but not at a higher degree, i.e., the spillover effect is bilateral within each pair of individuals. The deterministic component of the utility of player i of choosing y_i relative to $y_i = 0$ is given by:

$$U_i(y, \mathcal{X}, D, G; \alpha, \beta) = \alpha_i y_i + \sum_{j \neq i}^N \beta_{ij} y_i y_j.$$

Given a network G , covariates $\mathcal{X} = (X_1, \dots, X_N)$, and a treatment allocation $D = (d_1, \dots, d_N)$, the coefficient α_i on unit i 's choice depends upon their own covariates and treatment status as well as those of all of their neighbors; the coefficient β_{ij} on the quadratic term $y_i y_j$ depends upon their own covariates and treatment status as well as those of their neighbor unit j . Allowing for α_i and β_{ij} to be unconstrained, this specification of the utility function is without loss of generality since choice is binary. The condition for the existence of a potential function (Eq.2.2), however, requires that $\beta_{ij} = \beta_{ji}$ for all $i \neq j \in \mathcal{N}$.⁴ This symmetry assumption on β_{ij} restricts the spillover effect of unit i 's choice on unit j . The approach that is proposed in this paper to obtain an optimal treatment allocation can be implemented for any utility function specification as long as this symmetry condition is imposed. Nevertheless, to obtain a specific welfare performance guarantee for our method, we consider the following parametric specification of the utility functions in the remaining sections.

$$U_i(y, \mathcal{X}, D, G; \theta) = \left[\theta_0 + \theta_1 d_i + X_i' \theta_2 + X_i' \theta_3 d_i + A_N \sum_{j \in \mathcal{N}_i} \theta_4 m_{ij} d_j \right] y_i + A_N \sum_{j \in \mathcal{N}_i} m_{ij} (\theta_5 + \theta_6 d_i d_j) y_i y_j, \quad (2.5)$$

where $m_{ij} = m(X_i, X_j)$ is a (bounded) real-valued function of personal characteristics. In the absence of binary treatments, this specification appears in [Mele \(2017\)](#), and [Sheng \(2020\)](#). m_{ij} measures the distance between unit i 's characteristics and unit j 's characteristics; the spillover effect is weighted by how similar two units appear. A_N is a term that controls the magnitude of spillovers so that the unknown parameters $(\theta_4, \theta_5, \theta_6)$ can be independent of the size of the network. As A_N increases, unit i 's decision is more heavily influenced by their neighbor's choices and treatment status. To generate non-degenerate stochastic choice irrespective of the size and density of the network, we consider the following condition on A_N :

Assumption 7. *A positive sequence A_N satisfies:*

$$A_N \bar{N} \leq \mathcal{O}(1).$$

⁴For a potential function to exist, after eliminating zero terms, we require that $U_i(1, 0, y_{-ij}) - U_i(1, 1, y_{-ij}) + U_j(1, 1, y_{-ij}) - U_j(0, 1, y_{-ij}) = 0$. This implies that $-\beta_{ij} + \beta_{ji} = 0$.

A choice of A_N that satisfies Assumption 7 depends on the sparsity of the network structure. For the dense network (i.e., the maximum number of links a node can have increases linearly with N), one may choose $A_N = \frac{1}{N}$ to accommodate the increasing magnitude of the spillover effect. Similar choices of A_N have been used in many settings. For example, Sheng (2020) chooses $A_N = \frac{1}{N-2}$; Galeotti et al. (2020) chooses $A_N = 1$ but imposes an additional assumption on the coefficient. For the sparse network where \bar{N} is $\mathcal{O}(1)$, A_N can be set to 1.

The utility that unit i derives from an action is the sum of the net benefits that they accrue from their own actions and from those of their neighbors. In this work, we assume that only direct neighbors are valuable and units do not receive utility from one-link-away contacts. The total benefit of playing action $Y_i = 1$ has six components. When unit i chooses action $Y_i = 1$, they receive utility θ_0 from their own choice without treatment. They also receive additional utility $\theta_1 d_i$ depending upon their own treatment status. Their utility also has a heterogeneous treatment effect component $X_i'(\theta_2 + \theta_3 d_i)$, which depends upon their personal characteristics X_i . Units value treatment externalities; that is, treatment received by other units. Unit i receives additional utility $\theta_4 m_{ij}$ if their neighbor unit j receives treatment, no matter their own treatment status. Units value choice spillovers. When unit i is deciding whether to play action 1, they observe unit j 's choice and attributes. If unit j is a neighbor of unit i that chooses action 1 then this provides $\theta_5 m_{ij}$ additional utility to unit i . The final component corresponds to the choice spillovers from those neighbors who receive treatment. If both unit i and unit j receive treatment and both of them choose action 1, unit i receives additional utility $\theta_6 m_{ij}$ from the common treatment and choice.

Example 1. (Customer Purchase Decisions) Individual i makes a purchase decision Y_i (i.e., buy or not buy) for one product (e.g., Dropbox subscription, Orange from Sainsbury, iPhone). In this example, the social planner is the company that is trying to maximize the total number of customers that purchase its products. Individuals' purchase decisions sequentially depend upon the purchase decision of their friends or of their colleagues. The company observes individuals' friendships and then decides how to allocate discount offers to achieve its own targets. (e.g.,

Richardson and Domingos, 2002)

Example 2. (Criminal Network) In a criminal network, suspects are connected by a social network. Suspect i makes a decision whether to commit a crime, $Y_i = 0$, or not, $Y_i = 1$. The social planner in this example is the government or a police force that is trying to minimize the total number of crimes in the long run. The decision that a suspect makes about whether to commit a crime is based upon whether they and their friends have been arrested before ($d_i = 1$ denotes they have been arrested before and $d_i = 0$ denotes they have not been arrested in the past). The social planner observes the criminal network and decides which suspects to arrest. (e.g, *Lee et al., 2021*)

To ensure that our game is a potential game, we impose an additional assumption on m_{ij} . We assume that the following condition is satisfied.

Assumption 8. (Non-negative, Bounded and Symmetric Property) Function m_{ij} satisfies the following restrictions:

$$m(X_i, X_j) = m(X_j, X_i), \quad \forall i, j \in \mathcal{N}.$$

$$m_{ij} \in [\underline{m}, \overline{m}] \in \mathbb{R}_+, \quad \forall i, j \in \mathcal{N}.$$

Assumption 8 ensures that m_{ij} is symmetric hold for all $i, j \in \mathcal{N}$. Researchers can freely choose any m_{ij} which satisfies the above assumption. The following proposition indicates that our decision game is a potential game.

Proposition 2.2.1. (Potential Function) Under Assumption 4 and 8, the potential function $\Phi(y, \mathcal{X}, D, G; \theta)$ for $U_i(y, \mathcal{X}, D, G; \theta)$ specified in Eq.2.5 can be defined as:

$$\begin{aligned} \Phi(y, \mathcal{X}, D, G; \theta) = & \sum_{i=1}^N \left(\theta_0 + \theta_1 d_i + X_i'(\theta_2 + \theta_3 d_i) + A_N \sum_{j=1}^N \theta_4 m_{ij} G_{ij} d_j \right) y_i \\ & + \frac{A_N}{2} \sum_{i=1}^N \sum_{j=1}^N m_{ij} G_{ij} y_i y_j (\theta_5 + \theta_6 d_i d_j), \end{aligned}$$

and our interacted decision game is a potential game.

Proof of Proposition 2.2.1 is provided in Appendix A.1.2. We can, however, easily verify that this specification satisfies the definition of a potential function (i.e., Eq.2.1). Notice that the potential function is not the summation of the utility function across all units; summation of the utility function counts the interaction terms twice and violates Eq.2.1. By characterizing our game as a potential game, we can employ the stationary outcome distribution that we derived in Theorem 2.2.1 to evaluate the planner's expected welfare.

2.3 Treatment Allocation

The objective of the social planner is to select a treatment assignment $D^* \in \{0, 1\}^N$ that maximizes equilibrium mean outcomes subject to a capacity constraint that the number of individuals that are treated cannot exceed $\kappa > 0$:

$$D^* = \arg \max_{D \in \{0, 1\}^N} \sum_{i=1}^N \mathbb{E}_P[Y_i | \mathcal{X}, D, G; \theta], \quad (2.6)$$

$$s.t. \quad \sum_{i=1}^N d_i \leq \kappa.$$

From Theorem 2.2.1, the stationary joint distribution of Y depends on the treatment allocation D . Fixing the parameters θ , attributes \mathcal{X} , and network G , the social planner selects the joint distribution that maximizes equilibrium outcomes by manipulating treatment allocation rules.

In this work, we assume that the structural parameters θ underlying the sequential decision game are given and abstract from uncertainty in parameters estimation. In general, there are three common estimation strategies used in the MRF literature.

- *Markov chain Monte Carlo (MCMC; Metropolis and Ulam, 1949):* Geyer and Thompson (1992), Snijders et al. (2002), Mele (2017), Badev (2021).

- *Pseudo-likelihood*: [Besag \(1974\)](#), [Van Duijn et al. \(2009\)](#), [Boucher and Mourifié \(2017\)](#).
- *Variational approximation*: [Wainwright et al. \(2008\)](#), [Chatterjee and Diaconis \(2013\)](#), [Mele and Zhu \(2022\)](#).

MCMC involves sampling from a large class of joint distributions and scales well with the dimensionality of the sample space ([Bishop and Nasrabadi, 2006](#)). An issue, however, is that the mixing time of the Markov chain generated by Metropolis or Gibbs sampling takes exponential time ([Bhamidi et al., 2008](#); [Chatterjee and Diaconis, 2013](#)). Pseudo-likelihood focuses on the conditional probability—rather than the joint distribution—and is computationally fast. Its properties, however, are not well studied except in some specific cases ([Boucher and Mourifié, 2017](#)). Variational approximation, which is optimization-based rather than sampling-based, is an attractive alternative to MCMC if a fast optimization algorithm is available. To approximate the Gibbs distribution, a fast iteration algorithm for optimization is known ([Wainwright et al., 2008](#)) and is what we employ in a part of our algorithm (Algorithm 3 in Section 2.3.2).

2.3.1 Welfare Approximation

We cannot directly maximize the equilibrium welfare; instead, we seek to maximize the approximated welfare. We now discuss what prevents us from maximizing the equilibrium welfare. Recall that the objective function $W(D)$ from Eq. 2.6 is:

$$\begin{aligned}
W(D) &= \sum_{i=1}^N \mathbb{E}_P[Y_i | \mathcal{X}, D, G] \\
&= \sum_{i=1}^N \sum_{y \in \{0,1\}^N} y_i P(Y = y | \mathcal{X}, D, G) \\
&= \sum_{i=1}^N \sum_{y \in \{0,1\}^N} y_i \frac{\exp[\Phi(y, \mathcal{X}, D, G; \theta)]}{\sum_{\delta \in \{0,1\}^N} \exp[\Phi(\delta, \mathcal{X}, D, G; \theta)]} \\
&= \sum_{i=1}^N \sum_{y \in \{0,1\}^N} y_i \frac{\exp[w'_1 y + y' w_2 y]}{\sum_{\delta \in \{0,1\}^N} \exp[w'_1 \delta + \delta' w_2 \delta]},
\end{aligned} \tag{2.7}$$

where w_1 is a $N \times 1$ weighting vector and w_2 is a $N \times N$ weighting matrix. The i -th element in w_1 takes the value:

$$w_1^i = \theta_0 + \theta_1 d_i + (\theta_2 + \theta_3 d_i) X_i + A_N \sum_{j=1}^N \theta_4 m_{ij} G_{ij} d_j.$$

The i, j -th element in w_2 takes the value:

$$w_2^{ij} = \frac{A_N}{2} m_{ij} G_{ij} (\theta_5 + \theta_6 d_i d_j).$$

We define the denominator in Eq.2.7 – the *partition function* – as \mathcal{Z} :

$$\mathcal{Z} := \sum_{\delta \in \{0,1\}^N} \exp[w_1' \delta + \delta' w_2 \delta].$$

Since the partition function \mathcal{Z} sums all possible configurations (of which there are 2^N), it is infeasible to evaluate the expectation. When $N > 276$, there are more configurations than atoms in the observable universe (De Paula, 2020).

Given this well-known problem, we seek to approximate the distribution P using a tractable distribution Q . Defining $\mu_i^P := \mathbb{E}_P[Y_i | \mathcal{X}, D, G]$ and $\mu_i^Q := \mathbb{E}_Q[Y_i | \mathcal{X}, D, G]$, the objective function can be bounded from above by:

$$W(D) = \sum_{i=1}^N \mu_i^P \leq \sum_{i=1}^N |\mu_i^P - \mu_i^Q| + \sum_{i=1}^N \mu_i^Q. \quad (2.8)$$

The approximation error is, therefore, bounded by:

$$\sum_{i=1}^N \mu_i^P - \sum_{i=1}^N \mu_i^Q \leq \sum_{i=1}^N |\mu_i^P - \mu_i^Q|. \quad (2.9)$$

To characterize the approximation error, we introduce the Wasserstein 1-distance.

Definition 2.3.1 (Wasserstein 1-Distance). Let P and Q be two probability distributions

over $\{0, 1\}^N$. Define $\Omega(P, Q)$ as the set of all couplings of P and Q , i.e., all joint distributions of (Y, Y') such that $Y \sim P$ and $Y' \sim Q$. Let the Hamming distance be

$$d_H(Y, Y') = \sum_{i=1}^N \mathbb{1}\{Y_i \neq Y'_i\}.$$

Then the Wasserstein 1-distance equipped with Hamming distance is given by

$$W_1(P, Q) = \inf_{\omega \in \Omega(P, Q)} \mathbb{E}_\omega[d_H(Y, Y')].$$

We show that the approximation error in Eq. 2.9 is upper bounded by $W_1(P, Q)$ in the following lemma.

Lemma 2.3.1. Error Bound: *Let P and Q be two probability distributions over $\{0, 1\}^N$. We have:*

$$\sum_{i=1}^N |\mu_i^P - \mu_i^Q| \leq W_1(P, Q).$$

Proof of Lemma 2.3.1 is provided in Appendix A.1.3. To uniformly bound the Wasserstein 1-distance, we show our unique stationary distribution P satisfies the Talagrand transportation inequality⁵ under the following assumption.

Assumption 9. Spillover Effect: *The coefficients of strategic interaction satisfy:*

$$|\theta_5| + |\theta_6| < 4 \left(A_N \max_{i=1, \dots, N} \sum_{j=1}^N m_{ij} G_{ij} \right)^{-1}.$$

Or equivalently, there exists a constant $\alpha \in (0, 1)$ such that:

$$A_N (|\theta_5| + |\theta_6|) \max_{i=1, \dots, N} \sum_{j=1}^N m_{ij} G_{ij} \leq 4(1 - \alpha).$$

Assumption 9 is a common condition in the Markov random field literature, known as

⁵A distribution P satisfies the Talagrand transportation inequality if

Dobrushin's condition (Dobrushin, 1970), and it guarantees the uniqueness of the Gibbs distribution (i.e., given a system of conditional distributions of Y_i given Y_{-i} , for all $i \in \mathcal{N}$, as shown in Eq.2.4, there exists a unique Gibbs distribution as the joint distribution of Y). Assumption 9 restricts the magnitude of the spillover effects through concordant actions with neighbors, $y_i = y_j = 1$. Furthermore, Götze et al. (2019) shows that the distribution P is sub-Gaussian under Dobrushin's condition. Given the knowledge of parameter values, we can directly check if Assumption 9 holds or not in the given application. In the simulation exercise given below, we examine the performance of our greedy algorithm when Assumption 9 is relaxed.

Proposition 2.3.1. Transportation Inequality: *Under Assumption 4 to 9, there is a universal constant C_t such that the stationary distribution P , defined in Eq.2.3, satisfies:*

$$W_1(P, Q) \leq C_t \sqrt{\mathbb{KL}(Q||P)},$$

for all probability measures Q on $\{0, 1\}^N$ with a finite first moment.

Proof of Proposition 2.3.1 is provided in Appendix A.1.4. As a consequence, the approximation error is bounded by the $C\sqrt{\mathbb{KL}(Q||P)}$. It is natural to choose a distribution Q that minimizes the upper bound $\mathbb{KL}(Q||P)$ so as to reduce the approximation error. It is not, however, feasible to search over all tractable distributions to find Q ; We choose to work with an *independent Bernoulli distribution*.

Remark 2.3.1. Some social planners may target maximizing the expected utilitarian welfare (i.e., the summation of individual utilities) when choosing the optimal treatment allocation,

in which case the objective function becomes:

$$\begin{aligned}
W_U(D) &= \sum_{i=1}^N \mathbb{E}_P[U_i(y, \mathcal{X}, D, G; \boldsymbol{\theta}) | \mathcal{X}, D, G] \\
&= \sum_{i=1}^N \theta_{ij}^1 \mathbb{E}_P[y_i | \mathcal{X}, D, G] + \sum_{i=1}^N \sum_{j=1}^N \theta_{ij}^2 \mathbb{E}_P[y_i y_j | \mathcal{X}, D, G] \\
&= \sum_{i=1}^N \theta_{ij}^1 \mu_i^P + \sum_{i=1}^N \sum_{j=1}^N \theta_{ij}^2 \mu_{ij}^P,
\end{aligned}$$

where $\theta_{ij}^1 = \theta_0 + \theta_1 d_i + X_i' \theta_2 + X_i' \theta_3 d_i + A_N \sum_{j \in \mathcal{N}_i} \theta_4 m_{ij} d_j$, $\theta_{ij}^2 = A_N \sum_{j \in \mathcal{N}_i} (\theta_5 + \theta_6 d_i d_j) m_{ij}$, and $\mu_{ij}^P = \mathbb{E}_P[y_i y_j | \mathcal{X}, D, G]$. This μ_{ij}^P term leads to the bound on the objective function differing substantially from the one in Eq.2.8. Standard variational approximation does not apply in this setting. We leave analysis of this problem for future research.

2.3.2 Mean Field Method

Using an independent Bernoulli distribution to approximate the target distribution is called *naive mean field approximation* (Wainwright et al., 2008). This method can be viewed as a specific method in the general approach of *variational approximation*, which approximates a complicated probability distribution by a distribution belonging to a class of analytically tractable parametric distributions. In Eq.2.8, P corresponds to the target distribution to be approximated and Q corresponds to a simple parametric distribution approximating P . We consider the class of independent Bernoulli distributions as a parametric family for Q , since it delivers a feasible and fast optimization algorithm and the magnitude of its approximation error is already established in the literature.

The probability mass function of an independent Bernoulli distribution Q is expressed as:

$$Q(Y = y) = \prod_{i=1}^N (\mu_i^Q)^{y_i} (1 - \mu_i^Q)^{1-y_i}.$$

Let μ^Q be an $N \times 1$ vector that collects $\{\mu_i^Q\}_{i=1}^N$. The Kullback–Leibler divergence between Q and P equals:

$$\begin{aligned}
\mathbb{KL}(Q||P) &= \mathbb{E}_Q \left[\log \frac{Q(y)}{P(y)} \right] \\
&= \mathbb{E}_Q \left[\log \frac{Q(y)}{\exp[w_1' y + y' w_2 y - \log \mathcal{Z}]} \right] \\
&= \mathbb{E}_Q [\log Q(y) - w_1' y - y' w_2 y + \log \mathcal{Z}] \\
&= \log \mathcal{Z} - \left[w_1' \mu^Q + (\mu^Q)' w_2 \mu^Q - \sum_{i=1}^N [\mu_i^Q \log(\mu_i^Q) + (1 - \mu_i^Q) \log(1 - \mu_i^Q)] \right].
\end{aligned}$$

The last line holds since the diagonal entries of w_2 are zero and

$$\begin{aligned}
\mathbb{E}_Q[y' w_2 y] &= \mathbb{E}_Q \left[\sum_{i=1}^N \sum_{j \neq i}^N w_2^{ij} y_i y_j \right] = \sum_{i=1}^N \sum_{j \neq i}^N w_2^{ij} \mathbb{E}_Q[y_i y_j] = \sum_{i=1}^N \sum_{j \neq i}^N w_2^{ij} \mathbb{E}_Q[y_i] \mathbb{E}_Q[y_j] \\
&= (\mu^Q)' w_2 \mu^Q.
\end{aligned}$$

Recall \mathcal{Z} in Eq.2.7 sums over all possible configurations. \mathcal{Z} is, therefore, independent of Y (i.e., it is constant). We define $\mathcal{A}(\mu^Q, \mathcal{X}, D, G)$ as:

$$\mathcal{A}(\mu^Q, \mathcal{X}, D, G) := w_1' \mu^Q + (\mu^Q)' w_2 \mu^Q - \sum_{i=1}^N [\mu_i^Q \log(\mu_i^Q) + (1 - \mu_i^Q) \log(1 - \mu_i^Q)].$$

As such, minimizing $\mathbb{KL}(Q||P)$ is equivalent to maximizing $\mathcal{A}(\mu^Q, \mathcal{X}, D, G)$. We denote by $\tilde{\mu}$ the result of the following optimization:

$$\begin{aligned}
\tilde{\mu} &= \arg \sup_{\mu^Q} \mathcal{A}(\mu^Q, \mathcal{X}, D, G) \\
&= \arg \sup_{\mu^Q} w_1' \mu^Q + (\mu^Q)' w_2 \mu^Q - \sum_{i=1}^N [\mu_i^Q \log(\mu_i^Q) + (1 - \mu_i^Q) \log(1 - \mu_i^Q)].
\end{aligned} \tag{2.10}$$

Then the approximated distribution Q^* is expressed as:

$$Q^*(Y = y) = \prod_{i=1}^N (\tilde{\mu}_i)^{y_i} (1 - \tilde{\mu}_i)^{1-y_i}.$$

The first order condition of Eq.2.10 is:

$$\begin{aligned} \tilde{\mu}_i &= \frac{1}{1 + \exp[-(\theta_0 + \theta_1 d_i + X'_i(\theta_2 + \theta_3 d_i) + A_N \theta_4 \sum_{j \neq i} m_{ij} G_{ij} d_j + A_N \sum_{j \neq i} m_{ij} G_{ij} (\theta_5 + \theta_6 d_i d_j) \tilde{\mu}_j)]]} \\ &= \Lambda \left[\theta_0 + \theta_1 d_i + X'_i(\theta_2 + \theta_3 d_i) + A_N \theta_4 \sum_{j \neq i} m_{ij} G_{ij} d_j + A_N \sum_{j \neq i} m_{ij} G_{ij} (\theta_5 + \theta_6 d_i d_j) \tilde{\mu}_j \right]. \end{aligned} \quad (2.11)$$

Given that the above objective function (Eq.2.10) is non-concave, there may exist multiple maximizers. In the following proposition, we show that this optimization problem does have a unique maximizer.

Proposition 2.3.2. Unique Maximizer: *Under Assumptions 4 to 8, the optimization problem defining $\tilde{\mu}$ has a unique maximizer and the iteration procedure of Algorithm 3 converges to it.*

Proof of Proposition 2.3.2 is provided in Appendix A.1.5. To obtain the global optimum, it is sufficient to solve the first-order condition (Eq.2.11). Finding a root of the first-order conditions is feasible and there exists a fast off-the-shelf iterative method to compute $\tilde{\mu}$ (see Algorithm 3). Convergence of this algorithm has been extensively studied in the literature on variational approximation (Wainwright et al., 2008). Iteration in Algorithm 3 amounts to coordinate ascent of the mean field variational problem (Eq.2.10). Given that Eq.2.10 is a strictly concave function of μ_i when all other coordinates μ_{-i} are held fixed (Wainwright et al., 2008, §Chapter 5.3), the maximum is uniquely attained at every coordinate update. Bertsekas (2016, §Chapter 1.8) guarantees that $\{\tilde{\mu}^0, \tilde{\mu}^1, \dots\}$ converges to a local optimum. If, in addition, we constrain the magnitude of the spillovers of the neighbor's actions by $A_N \overline{m}(|\theta_5| + |\theta_6|) \overline{N} \leq 4$, we can show that the above iteration process forms a contraction mapping, implying that it converges to a unique global optimum (see Appendix A.1.6 for a

proof).

In approximating an interacted joint distribution by a fully independent distribution it is conceivable that there should be some information loss. The following theorem shows, however, that the information loss due to variational approximation (measured in terms of the Kullback-Leibler divergence) converges to zero as the size of the network grows to infinity.

Theorem 2.3.1. Approximation Error Bound: *Let Q^* denote the independent Bernoulli distribution solving Eq.2.10. Under Assumptions 4 to 8, the Kullback–Leibler divergence of Q^* from P is bounded from above by:*

$$\mathbb{KL}(Q^*||P) \leq C_1 A_N \bar{N} + C_2 N + \mathcal{O}\left(\sqrt{A_N^2 \bar{N}^2 N}\right) + \mathcal{O}\left(\sqrt{A_N^3 \bar{N}^2 N^2}\right) + o(N), \quad (2.12)$$

where C_1, C_2 are known constants that depend only upon θ and \bar{m} .

This theorem follows as a direct corollary of Chatterjee and Dembo (2016, §Theorem 1.6). Proof of Theorem 2.3.1 is provided in Appendix A.2.1. Theorem 2.3.1 shows that the upper bound on the approximation error depends upon the size of the network N and the magnitude of the spillover effect $A_N \bar{N}$. Given that Assumption 7 holds with non-increasing sequence of A_N , Theorem 2.3.1 clarifies that the leading term in the approximation error bound Eq.2.12 grows at $\mathcal{O}(N)$.

Recall from Eq.2.9 and Eq.2.12 that the error due to approximating the welfare at P by the welfare at Q^* can be bounded by $\sqrt{\mathbb{KL}(Q^*||P)} \leq \mathcal{O}(N^{1/2})$. If our objective is to maximize $\frac{1}{N} \sum_{i=1}^N \mu_i^P$, Theorem 2.3.1 implies that this term can be bounded from above by $\frac{1}{N} \sqrt{\mathbb{KL}(Q^*||P)} \leq \mathcal{O}(N^{-1/2})$, which converges to zero as N becomes large. This means that, as the size of the network becomes large, spillover effects become less important.

Remark 2.3.2. We derive the result in Theorem 2.3.1 by assuming that the structural parameters θ are independent of N . Recently, Joseph and Mark (2022) discusses the potential issue of using variational approximation when structural parameters depend upon the network size N . Joseph and Mark (2022) shows numerically that estimation of the parameters using

variational approximation can deliver estimates that are far from true parameter values if the network is more transitive. If the transitivity of a network depends upon its size, then the approximation error may not shrink to zero. The analogous problem for us is when the parameters that measure the spillover effects depend upon N , our approximation error bound may not converge to zero. We emphasize several things. First, the estimation method that we adopted in the empirical application section is MCMC-MLE, which is known to converge to the true parameter value. That is also supported by [Joseph and Mark \(2022\)](#). Second, in our empirical application, we show for networks of different sizes that our parameter estimates are quite similar to each other. This is indicative of our parameters not depending upon N . In addition, the empirical results illustrate that our proposed method performs well across all network sizes in the dataset.

2.3.3 Implementation

In the last section, we discussed how to approximate the mean value of the outcome variable using the mean field method. In this section, we propose an algorithm to allocate treatment so as to maximize the approximated welfare and discuss its implementation.

Suppose that the set of feasible allocations is subject to a capacity constraint, $\sum_{i=1}^N d_i \leq \kappa$, where $\kappa \in \mathbb{N}_+$ specifies the maximum number of units that can be treated. We denote the set of feasible allocations by $\mathcal{D}_\kappa \equiv \{D \in \{0, 1\}^N : \sum_{i=1}^N d_i \leq \kappa\}$, and the approximated welfare by:

$$\tilde{W}(D) = \sum_{i=1}^N \tilde{\mu}_i.$$

We seek to maximize the approximated welfare:

$$\tilde{D} = \arg \max_{D \in \mathcal{D}_\kappa} \tilde{W}(D). \quad (2.13)$$

As shown in the [Eq. 2.11](#), $\{\tilde{\mu}_i\}_{i=1}^N$ is a large non-linear simultaneous equation system. The ap-

proximated mean value $\tilde{\mu}_i$ of each unit i depends non-linearly upon the approximated mean value $\tilde{\mu}_j$ and the treatment assignment d_j of her neighbor, unit j . Hence, the optimization problem (Eq.2.13) becomes a complicated combinatorial optimization. We propose a greedy algorithm (Algorithm 4) to solve this problem heuristically.

The idea of our greedy algorithm is to assign treatment to the unit that contributes most to the welfare objective, repeating this until the capacity constraint binds. Specifically, in each round, Algorithm 3 computes the marginal gain of receiving treatment for each untreated unit. We refer to the unit whose treatment induces the largest increase in the approximated welfare as the most influential unit in that round. We provide a theoretical performance guarantee for our greedy algorithm in Section 2.3.4. We also numerically examine the performance of our method in Section 2.4.

In Algorithm 4, we use a variational approximation method to compute $\tilde{\mu}$ for each assignment rule and for each round (i.e., there are $\mathcal{O}(N)$ operations in each round). Alternatively, MCMC can be used to simulate the mean value μ of the unique stationary distribution (Eq.2.3) instead of computing the variationally approximated $\tilde{\mu}$. Since MCMC may require exponential time for convergence (Chatterjee and Diaconis, 2013) though, simulating μ is infeasible for a large network (i.e., MCMC needs to be run $\mathcal{O}(\kappa N)$ times). In Section 2.4.2, we compare the welfare computed using these two methods for various treatment allocation rules.

Algorithm 3: Computing $\tilde{\mu}$

Input: Weighted adjacency matrix G , treatment allocation D , covariates \mathcal{X} , parameters θ , and threshold ϱ

Initialization: Draw $\tilde{\mu}_i^0 \sim U[0, 1], \forall i \in \mathcal{N}; \quad t = 1$

while $\mathcal{A}(\tilde{\mu}^t, \mathcal{X}, D, G) - \mathcal{A}(\tilde{\mu}^{t-1}, \mathcal{X}, D, G) > \varrho$ **do**

$t \leftarrow t + 1$

for $i \leftarrow 1, \dots, N$ **do**

$\tilde{\mu}_i^t \leftarrow \Lambda \left[\theta_0 + \theta_1 d_i + X_i'(\theta_2 + \theta_3 d_i) + A_N \theta_4 \sum_{j \neq i} m_{ij} G_{ij} d_j + A_N \sum_{j \neq i} m_{ij} G_{ij} (\theta_5 + \theta_6 d_i d_j) \tilde{\mu}_j^{t-1} \right]$

end

end

Return $\tilde{\mu} \leftarrow \tilde{\mu}^t$

Algorithm 4: Maximizing Over Treatment Allocation Rules

Input: Weighted adjacency matrix G , covariates \mathcal{X} , parameters θ , capacity constraint

κ

Output: Treatment allocation regime \tilde{D}

Initialization: $D \leftarrow 0_{N \times 1}$

if $\sum_{i=1}^N d_i < \kappa$ **then**

for i with $d_i = 0$ **do**

$d_i \leftarrow 1$, denote new treatment vector as D'

$\tilde{\mu} \leftarrow \text{Run Algorithm 3 with } D'$

$\Delta_i \leftarrow \tilde{W}(D') - \tilde{W}(D)$

end

$i^* \leftarrow \arg \max_i \Delta_i$

$d_{i^*} \leftarrow 1$

else

$\tilde{D} \leftarrow D$

end

2.3.4 Theoretical Analysis

In this section, we analyze the regret of the treatment allocation rule computed using our greedy algorithm. Given $D^* = \arg \max_{D \in \mathcal{D}_\kappa} W(D)$ is the maximizer of $W(D)$, then $W(D^*)$ denotes the maximum value of $W(D)$. *Regret* is the gap between the maximal equilibrium (oracle) welfare $W(D^*)$ and the equilibrium welfare attained at the treatment allocation rule computed using our greedy algorithm $W(D_G)$. We decompose regret into four terms:

$$\begin{aligned} W(D^*) - W(D_G) &= \underbrace{W(D^*) - \tilde{W}(D^*)}_{\leq \sqrt{2\text{KL}(Q^*||P)}} + \underbrace{\tilde{W}(D^*) - \tilde{W}(\tilde{D})}_{\leq 0} + \underbrace{\tilde{W}(\tilde{D}) - \tilde{W}(D_G)}_{\text{Regret from greedy}} + \underbrace{\tilde{W}(D_G) - W(D_G)}_{\leq \sqrt{2\text{KL}(Q^*||P)}} \\ &\leq \sqrt{8\text{KL}(Q^*||P)} + \tilde{W}(\tilde{D}) - \tilde{W}(D_G). \end{aligned} \tag{2.14}$$

The first term corresponds to the approximation error of using variational approximation; the second term comes from using the maximizer of the approximated equilibrium welfare \tilde{D} ; the third term comes from using our greedy algorithm instead of using the maximizer of the variationally approximated welfare; and the last component is again introduced by using the approximated equilibrium welfare $\tilde{W}(D)$.

Theorem 2.3.1 provides an upper bound on the approximation error $\sqrt{8\text{KL}(Q^*||P)}$:

$$\sqrt{8\text{KL}(Q^*||P)} \leq \sqrt{8 \left[C_1 A_N \bar{N} + C_2 N + \mathcal{O} \left(\sqrt{A_N^2 \bar{N}^2 N} \right) + \mathcal{O} \left(\sqrt{A_N^3 \bar{N}^2 N^2} \right) \right]} + o(N^{1/2}). \quad (2.15)$$

In Eq.2.15, the convergence rate of the upper bound on approximation error depends upon the network size N , sparsity of the network \bar{N} , and the choice of normalization A_N . Given Assumption 7, the convergence rate for the upper bound on the variational approximation error simplifies to:

$$\sqrt{8\text{KL}(Q^*||P)} \leq \mathcal{O}(N^{1/2}). \quad (2.16)$$

For a general objective function mapping $\{0, 1\}^N$ to \mathbb{R} , however, there is no theoretical performance guarantee for the greedy algorithm, i.e., it is not known how much worse the greedy optimizer can be than the global optimum in terms of the value of the objective function. For a class of non-decreasing submodular functions on $\mathcal{D}_\kappa \subset \{0, 1\}^N$, Nemhauser et al. (1978) shows the existence of performance guarantees $(1 - 1/e)$.

Unfortunately, *submodularity* does not generally hold for our problem (Eq.2.13). Other applications have faced the same issue. Relaxing the requirement of submodularity, Conforti and Cornuéjols (1984) introduces the concept of *curvature* to characterize a constant factor in the performance guarantee. Das and Kempe (2011) introduces the *submodularity ratio* to define the closeness of a set function to submodularity. Bian et al. (2017) combines these two concepts (curvature and the submodularity ratio) to obtain a performance guarantee for the greedy algorithm for a large class of non-submodular functions. In what follows, we apply these techniques to the variationally approximated welfare.

The definitions of submodularity, the submodularity ratio, and the curvature of a set function f are as follows.

Definition 2.3.2. (Submodularity): A set function is a submodular function if:

$$\sum_{k \in R \setminus S} [f(S \cup \{k\}) - f(S)] \geq f(S \cup R) - f(S), \quad \forall S, R \subseteq \mathcal{N}.$$

Definition 2.3.3. (Submodularity Ratio) The submodularity ratio of a non-negative set function $f(\cdot)$ is the largest γ such that

$$\sum_{k \in R \setminus S} [f(S \cup \{k\}) - f(S)] \geq \gamma [f(S \cup R) - f(S)], \quad \forall S, R \subseteq \mathcal{N}.$$

Definition 2.3.4. (Curvature) The curvature of a non-negative set function $f(\cdot)$ is the smallest value of ξ such that

$$f(R \cup \{k\}) - f(R) \geq (1 - \xi)[f(S \cup \{k\}) - f(S)], \quad \forall S \subseteq R \subseteq \mathcal{N}, \forall k \in \mathcal{N} \setminus R.$$

The submodularity of a set function is analogous to concavity of a real function and implies that the function has diminishing returns. The marginal increase in the probability of choosing action 1 decreases with the number of treated units. The submodularity ratio captures how much greater the probability of choosing action 1 is from providing treatment to a group of units versus the combined benefit of treating each unit individually. Curvature can be interpreted as how close a set function is to being additive.

We associate the set function $f(\cdot)$ in the above definitions with the variationally approximated welfare $\tilde{W}(\cdot)$, which we view as a real-valued mapping of treatment allocation sets $\mathcal{D} \subset \mathcal{N}$ (i.e., $\mathcal{D} = \{i \in \mathcal{N} : d_i = 1\}$):

$$\begin{aligned}\tilde{W}(\mathcal{D}) = & \sum_{i \in \mathcal{D}} \Lambda[\theta_0 + \theta_1 + X'_i(\theta_2 + \theta_3) + A_N \theta_5 \sum_{\substack{j \neq i \\ j \in \mathcal{N}}} m_{ij} G_{ij} \tilde{\mu}_j + A_N \sum_{\substack{j \neq i \\ j \in \mathcal{D}}} m_{ij} G_{ij} (\theta_4 + \theta_6 \tilde{\mu}_j)] \\ & + \sum_{k \in \mathcal{N} \setminus \mathcal{D}} \Lambda[\theta_0 + X'_k \theta_2 + A_N \theta_4 \sum_{\ell \in \mathcal{D}} m_{k\ell} G_{k\ell} + A_N \theta_5 \sum_{\substack{\ell \neq k \\ \ell \in \mathcal{N}}} m_{k\ell} G_{k\ell} \tilde{\mu}_\ell].\end{aligned}$$

We characterize the submodularity ratio and curvature of $\tilde{W}(\cdot)$ to obtain an analytical performance guarantee for our greedy algorithm. In addition, we restrict our analysis to settings of positive treatment and spillover effects by imposing the following assumption.

Assumption 10. (Positivity and Monotonicity) We assume that $\tilde{W}(\mathcal{D})$ is a non-negative and non-decreasing set function (i.e., $\tilde{W}(\mathcal{D}) \geq \tilde{W}(\tilde{\mathcal{D}})$ if $\tilde{\mathcal{D}} \subseteq \mathcal{D}$).

Assumption 10 restricts the signs of the overall treatment effects, which includes direct and indirect treatment effects. One sufficient condition of Assumption 10 is that parameters $\theta_1, \theta_3, \theta_4, \theta_5, \theta_6 \geq 0$. This assumption works in many applications, such as allocating vaccinations to increase social health, providing discounts to encourage purchase, and assigning tax auditing to encourage paying tax.

By showing that a set function is non-decreasing, its curvature ξ and its submodularity ratio γ must belong to $[0, 1]$ (Bian et al., 2017). Having $\xi \in [0, 1]$ and $\gamma \in [0, 1]$ is not, however, enough to attain a nontrivial performance guarantee. For instance, if $\gamma = 0$, the lower bound in Theorem 2.3.2 equals 0, which is a trivial lower bound; if $\xi = 0$, then the lower bound equals γ , which could be 0. To rule out these trivial cases, we impose the following assumption, which gives a sufficient condition to bound the submodularity ratio and curvature away from 0 and 1.

Assumption 11. (Lower Bound on N) We assume that the sample size satisfies:

$$N \geq (A_N \theta_4 \underline{N} \cdot \underline{m} + \theta_1)/4.$$

Assumption 11 restricts the sample size of the network. Even in a dense network, \underline{N} (the minimum number of edges for one unit in the network) can be small and, accordingly, the requirement on the sample size. Assumption 11 easily holds for the medium or large size of networks. We are now able to provide a performance guarantee for our greedy algorithm.

Theorem 2.3.2. *Performance Guarantee for greedy Algorithm:* *Under Assumptions 4 to 11, the curvature ξ of $\tilde{W}(\mathcal{D})$ and the submodularity ratio γ of $\tilde{W}(\mathcal{D})$ are in $(0, 1)$. The greedy algorithm enjoys the following approximation guarantee for the problem in Eq. 2.13:*

$$\tilde{W}(D_G) \geq \frac{1}{\xi}(1 - e^{-\xi\gamma})\tilde{W}(\tilde{D}),$$

where D_G is the treatment assignment rule that is obtained by Algorithm 4.

The second part of Theorem 2.3.2 is taken from (Bian et al., 2017, §Theorem 1). Proof of Theorem 2.3.2 is provided in Appendix A.2.2. Theorem 2.3.2 indicates that there exists a performance guarantee that depends upon the unknown curvature and upon the submodularity ratio. The first part of Theorem 2.3.2 dictates that the performance guarantee is a non-trivial bound. It is infeasible to determine ξ and γ for $\tilde{W}(\mathcal{D})$; it is, however, possible to derive an upper bound for ξ and a lower bound for γ , which combined with Assumption 11, excludes triviality. Combining all of the previous results, we are able to use Bian et al. (2017, §Theorem 1) to provide a non-trivial performance guarantee on $\tilde{W}(\mathcal{D})$. We emphasize that if $\xi = 1$ and $\gamma = 1$, the performance guarantee in Theorem 2.3.2 coincides with the well-known performance guarantee constant of the greedy algorithm for submodular functions (i.e., $1 - 1/e$ Nemhauser et al., 1978). If $\xi < 1$ or $\gamma < 1$, the performance guarantee is worse than $1 - 1/e$.

Via Theorem 2.3.2 we can obtain an upper bound on the regret from using our greedy algorithm:

$$\tilde{W}(\tilde{D}) - \tilde{W}(D_G) \leq \left[1 - \frac{1}{\xi}(1 - e^{-\xi\gamma})\right]\tilde{W}(\tilde{D}). \quad (2.17)$$

Combining Eq.2.16, Eq.2.17, and Eq.2.14, and noting that $\tilde{W}(\tilde{D})$ is at most $\mathcal{O}(N)$, we obtain the next theorem:

Theorem 2.3.3. Regret Bound: *Let D^* denote the maximizer of $\tilde{W}(\mathcal{D})$ and D_G be the assignment vector obtained by Algorithm 4. Under Assumptions 4 to 11, given curvature ξ and submodularity ratio γ , the regret is bounded from above by:*

$$W(D^*) - W(D_G) \leq \mathcal{O}(N^{1/2}) + \mathcal{O}(N) \left[1 - \frac{1}{\xi} (1 - e^{-\xi\gamma}) \right].$$

Theorem 2.3.3 is our key result. It characterizes the convergence rate of overall regret, showing its dependence on the network complexity and the network size. The dependence upon the parameters in the utility function is shown implicitly via the terms C_1 and C_2 in Theorem 2.3.1. If we examine the average equilibrium welfare, then the regret bound becomes:

$$\frac{1}{N} (W(D^*) - W(D_G)) \leq \mathcal{O}(N^{-1/2}) + \frac{1}{N} \tilde{W}(\tilde{D}) \left[1 - \frac{1}{\xi} (1 - e^{-\xi\gamma}) \right].$$

The first term is the approximation error and shrinks to zero as N goes to infinity. Given that $\tilde{W}(\tilde{D})$ can be a function of N , the regret that is associated with our greedy algorithm can converge to a constant.

2.4 Simulation Exercises

In this section, we evaluate the performance of our greedy algorithm in simulation exercises. We use an Erdős-Renyi model to generate random social networks. For each choice of N , we generate 100 networks with fixed density (i.e., 0.3 and 0.6)⁶ and use the average of the equilibrium welfare over these 100 networks to assess the performance of our method. For personal covariates \mathcal{X} , we choose a binary variable that is generated from a Bernoulli distribution $B(0.5)$. We specify $m(X_i, X_j)$ as $m_{ij} = |X_i - X_j|$. We report the equilibrium welfare as

⁶Number of edges = density $\times \frac{N(N-1)}{2}$.

the per-person equilibrium average, $\max_{D \in \mathcal{D}_\kappa} 1/N \sum_{i=1}^N \mathbb{E}[Y_i | \mathcal{X}, D, G]$. In addition, we specify the tolerance level ϱ of Algorithm 3 as $1.0\text{E} - 9$. The capacity constraint that we choose is $\kappa = 30\%N$. To evaluate the impact of the unique fixed point of the iteration on the performance of our greedy algorithm, we choose two parameter sets. The first set of parameters guarantees the unique fixed point whilst the second set of parameters violates this condition. Table 2.1 summarizes the values of the parameters in our simulation.

<i>Parameters</i>	θ_0	θ_1	θ_2	θ_3	θ_4	θ_5	θ_6
<i>Set 1</i>	-2	0.5	0.1	0.6	0.7	0.8	0.9
<i>Set 2</i>	-2	0.5	0.1	0.6	0.7	7	7

Table 2.1: Summary of the parameter values

In the following sections, we compare our greedy algorithm with random allocation in a small network setting and in a large network setting. Random allocation assigns treatment to a fraction κ of units independently of personal characteristics and network structure. In the small network setting, we are able to compute the distribution of outcomes at every possible assignment vector, and use a brute force method to find an optimal treatment allocation. Using the welfare level at the optimal assignment as a benchmark, we can calculate the regret of our greedy algorithm. Since the number of possible assignment vectors grows rapidly with the number of units, we cannot compute the regret in the large network analysis of Section 2.4.2 in this way. We instead assess the welfare performance of the greedy targeting rule in comparison to the welfare level of the No treatment rule.

2.4.1 Small Network

We consider $N = 5, 7, 9, 11, 13$ or 15 to be a small network setting in our simulation exercise. First, we consider all possible treatment allocations subject to the capacity constraint and perform brute force search to find an optimal assignment. For instance, when $N = 15$, the number of feasible assignment vectors is 32,768. We compute the joint distribution of

outcomes at each possible treatment allocation by applying the joint probability mass function of the Gibbs distribution (Eq.2.3). Second, to assess the welfare loss from implementing the variational approximation, we evaluate the regret of a treatment assignment rule that is obtained by maximizing the variationally approximated welfare over every feasible treatment allocation meeting the capacity constraint (without greedy optimization). We label this method of obtaining the optimal treatment assignment as *brute force with variational approximation* (BFVA). Table 2.2 records the main differences between the two aforementioned methods and the greedy targeting rule in terms of in-sample average welfare.

From Table 2.2, we find that our greedy algorithm performs as well as the brute force method in a small network setting except when $N = 5$ (1% gap for $N = 5$). This indicates a good performance of our method. We find that the regret when $N = 5$ mainly comes from the approximation error of using a variational approximation. As we have shown in Theorem 2.3.1, the upper bound on the Kullback–Leibler divergence can be large when the sample size is small. This coincides with the empirical result. Our greedy algorithm can, however, achieve the same performance as BFVA, which means that using our greedy algorithm has a negligible effect upon regret.

In Figure 2.1, we compare the regret from using our greedy algorithm to random allocation for parameter set 1.⁷ Here, random allocation means that we randomly draw 50 allocation rules that satisfy the capacity constraint and average the welfare that they generate. The left-hand graph presents this comparison for density equal to 0.3; the right-hand graph presents this comparison for density equal to 0.6. From Figure 2.1, we find that the performance gap between our greedy targeting rule and random allocation in terms of regret ranges from 7% to 14%.

Figure 2.2 indicates the results from using parameter set 2. Parameter set 2 assumes much larger spillover effects that create highly dependent Gibbs distribution. Regret is greater than for parameter set 1, both when using our greedy algorithm and using random allocation. This

⁷Parameter set 1 satisfies the sufficient condition for convergence to global optimum shown in Appendix A.1.6.

indicates that the stronger dependence in Gibbs distribution deteriorates the quality of the variational approximation by independent Bernoullis. We note, however, that for the size of networks being $N = 7, 9, 11, 13, 15$, the regret from using our greedy algorithm stays within 10% of the maximal welfare and it dominates the performance of random allocation. These numerical evidence suggests that the advantage to using our greedy method is maintained for a wide range of the size of the spillover effects.




Figure 2.1: Comparison between the greedy algorithm and random allocation for the parameter set 1 (Left: density = 0.3 and Right: density = 0.6)




Figure 2.2: Comparison between the greedy algorithm and random allocation for the parameter set 2 (Left: density = 0.3 and Right: density = 0.6)

2.4.2 Large Network

We now extend our simulation exercise to large network settings where $N = 50, 100$ or 150 . As previously mentioned, we can neither search over all possible allocation vectors nor

compute the joint distribution over all possible vectors in a large network setting. To deal with these two problems, we first choose a baseline assignment rule – the No treatment rule – with which to compare the allocation rules that we compute. We evaluate the additional average welfare that we gain by providing treatment relative to the No treatment rule, rather than relative to the optimal assignment rule as we did for the small network setting. In Table 2.3, we summarize the average welfare for treatment assignment rules corresponding to greedy targeting, random allocation, and No treatment. Second, we use Gibbs sampling to approximate the joint distribution (Eq.2.3), iterating 10,000 times (burning period equal to 5,000) for each class of treatment rule.

Using Gibbs sampling, however, is not necessarily a feasible method to evaluate random allocation (and more generally) in a large network given its slow convergence. In the exercise, we use 10 random networks and 10 random draws, which takes approximately 30 hours to compute a result for random allocation.⁸ In contrast, it takes only 20 seconds to obtain a result for random allocation using variational approximation.

In Table 2.3, we compare the welfare delivered by Gibbs sampling with that delivered by variational approximation for the three aforementioned classes of treatment assignment rules. All the results in Table 2.3 are computed across 100 random networks, using the average of 10 random draws for random allocation, and with the capacity constraint set at $0.3N$. Table 2.4 indicates that variational approximation constitutes a good approximation of the Gibbs distribution (Eq.2.3), providing strong evidence in favour of using the variational approximation in our algorithm even for strongly dependent Gibbs distributions.

Table 2.3 indicates that using our greedy algorithm leads to an increase in welfare of approximately 10% as compared with random allocation. Relative to No treatment, our greedy algorithm performs 37% \sim 55% better than the random allocation. This result is robust to the network density. This suggests that the welfare gain from using our greedy algorithm carries over to the large network setting.

⁸We use parallel processing on a computer with an 8 core Intel i7-10700 CPU and 32GB RAM.

<i>Allocation Rule</i>	<i>Average Welfare with MCMC</i>			<i>Average Welfare with VA</i>		
	<i>N = 50</i>	<i>N = 100</i>	<i>N = 150</i>	<i>N = 50</i>	<i>N = 100</i>	<i>N = 150</i>
<i>Density = 0.3</i>						
greedy algorithm	0.186 (< 0.01)	0.186 (< 0.01)	0.186 (< 0.01)	0.186 (< 0.01)	0.186 (< 0.01)	0.186 (< 0.01)
Random allocation	0.166 (< 0.01)	0.170 (< 0.01)	0.170 (< 0.01)	0.164 (< 0.01)	0.170 (< 0.01)	0.169 (< 0.01)
No treatment rule	0.126 (< 0.01)	0.127 (< 0.01)	0.127 (< 0.01)	0.126 (< 0.01)	0.127 (< 0.01)	0.127 (< 0.01)
<i>Density = 0.6</i>						
greedy algorithm	0.194 (< 0.01)	0.193 (< 0.01)	0.193 (< 0.01)	0.194 (< 0.01)	0.193 (< 0.01)	0.193 (< 0.01)
Random allocation	0.173 (< 0.01)	0.178 (< 0.01)	0.178 (< 0.01)	0.172 (< 0.01)	0.178 (< 0.01)	0.178 (< 0.01)
No treatment rule	0.128 (< 0.01)	0.129 (< 0.01)	0.129 (< 0.01)	0.127 (< 0.01)	0.129 (< 0.01)	0.129 (< 0.01)

Table 2.3: Comparison between the average welfare computed using Gibbs sampling and variational approximation for parameter set 1

<i>Allocation Rule</i>	<i>Average Welfare with MCMC</i>			<i>Average Welfare with VA</i>		
	<i>N = 50</i>	<i>N = 100</i>	<i>N = 150</i>	<i>N = 50</i>	<i>N = 100</i>	<i>N = 150</i>
<i>Density = 0.3</i>						
greedy algorithm	0.227 (< 0.01)	0.218 (< 0.01)	0.215 (< 0.01)	0.237 (< 0.01)	0.228 (< 0.01)	0.225 (< 0.01)
Random allocation	0.201 (< 0.01)	0.203 (< 0.01)	0.203 (< 0.01)	0.209 (< 0.01)	0.214 (< 0.01)	0.213 (< 0.01)
No treatment rule	0.143 (< 0.01)	0.143 (< 0.01)	0.143 (< 0.01)	0.149 (< 0.01)	0.150 (< 0.01)	0.149 (< 0.01)
<i>Density = 0.6</i>						
greedy algorithm	0.317 (< 0.01)	0.305 (< 0.01)	0.299 (< 0.01)	0.346 (< 0.01)	0.343 (< 0.01)	0.339 (< 0.01)
Random allocation	0.287 (< 0.01)	0.293 (< 0.01)	0.292 (< 0.01)	0.321 (< 0.01)	0.334 (< 0.01)	0.333 (< 0.01)
No treatment rule	0.171 (< 0.01)	0.171 (< 0.01)	0.170 (< 0.01)	0.207 (< 0.01)	0.208 (< 0.01)	0.208 (< 0.01)

Table 2.4: Comparison between the average welfare computed using Gibbs sampling and variational approximation for parameter set 2

2.5 Empirical Application

We illustrate our proposed method using the dataset of [Banerjee et al. \(2013\)](#), which examines take-up of a microfinance initiative in India.⁹ A detailed description of the study can be

⁹The dataset is available at <https://doi.org/10.7910/DVN/U3BIHX>.

found in the original paper. This study features 43 villages in Karnataka that participated in a newly available microfinance loan program. Bharatha Swamukti Samsthe (BSS)—an Indian non-governmental microfinance institution administering the initiative—provided information about the availability of microfinance and program details (the treatment) to individuals that they identified as ‘leaders’ (e.g., teachers, shopkeepers, savings group leaders, etc.) so as to maximize the number of households that chose to adopt the microfinance product. The data provide network information at the household level (network data is available across 12 dimensions, including financial and medical links, social activity, and known family members) for each village. We use all available households characteristics that are available in the dataset (quality of access to electricity, quality of latrines, number of beds, number of rooms, the number of beds per capita, and the number of rooms per capita) as covariates. The program started in 2007, and the survey for microfinance adoption was finished in early 2011. We treat each household’s choice about whether to purchase microfinance or not as observations drawn from a stationary distribution of the sequential game.

The most common occupations in these villages are in agriculture, sericulture, and dairy production ([Banerjee et al., 2013](#)). In addition, these villages had almost no exposure to microfinance institutions and other types of credit before this program. Our target in this application is to maximize the participation rate of microfinance (4 years after program assignment) given a capacity constraint on treatment (i.e., Eq.2.6); we set our capacity constraint equal to the number of households that BSS contacted in the original study. In the previous sections, we have assumed that the parameters θ are given; here, we must estimate them. We allow the parameters of our utility function to be different across villages, and estimate them for each village using Markov Chain Monte Carlo Maximum Likelihood ([Snijders et al., 2002](#)). For each iteration of the procedure, we set the number of draws in the Gibbs sampling procedure equal to $200N^2$. In addition, we choose $m(X_i, X_j) = \frac{1}{1+|X_i - X_j|}$, which is a monotonically decreasing function in the metric between X_i and X_j . We also note that each household is connected to approximately 10 others on average across all of the 43 villages.

Comparing this with the total number of households in each village (there are between 107 and 341 households in each village), we find that the household network in each village is a sparse network. We, therefore, choose $A_N = 1$ in this application.

Table 2.5 shows the average probability of taking up microfinance observed in each village (column Sample Avg.) and the prediction of village-level take-up probability obtained from our MCMC-MLE estimates (column Welfare under Original). We refer to the former as the *Sample Average* and the latter as the *Welfare under Original Allocation*. We provide standard errors for the Sample Average, which we calculate using network HAC estimation (Leung (2019); and Kojevnikov et al. (2021)). We compute Welfare under Original Allocation by substituting the estimated parameters and the original treatment allocation (used by BSS) into our model. To further evaluate the performance of our proposed method, we randomly draw 100 treatment allocations that satisfy the capacity constraint in each village, and calculate the probability of purchasing microfinance for each allocation. We refer to the average probability over these draws as the *Welfare under Random Allocation*. We then implement our proposed method with the estimated parameters to find the optimal treatment allocation rule. We refer to the share of households adopting microfinance according to the optimal treatment allocation and our model as the *Welfare under Greedy Allocation*. Table 2.5 records these statistics for the 43 villages in the dataset, with the final column comparing the Welfare under our Greedy Allocation with the Welfare under Original Allocation. It also contains bootstrap standard errors for welfare based on 100 bootstrap samples for each village. The bootstrap samples are obtained by drawing outcomes from the MCMC stationary distribution simulated under the structural parameters estimated with the original sample. To obtain the standard errors for the Welfare under Greedy Allocation, for each bootstrap sample, we estimate the structural parameters and simulate welfare under a greedy optimal allocation.¹⁰

First, we note that the estimated average share of households who adopt microfinance

¹⁰If an optimal welfare is differentiable in the structural parameters and the greedy optimal allocation under the true structural parameter value is unique, we expect the bootstrap delta method applies to yield asymptotically valid standard errors.

under the MCMC-MLE estimates fits the data well for all 43 villages. We also find that the HAC standard error for the empirical average tends to be large, since there are relatively few households in each village, and this noise is a possible explanation for the large differences between the estimated share of households purchasing microfinance and the empirical average in some villages. Second, we find that random allocation delivers a comparable level of welfare to the original treatment allocation. This result is indicative that the leaders that BSS selected were not particularly effective in encouraging take-up by other households. Third, we find that our proposed method compares favourably to the method that is implemented in [Banerjee et al. \(2013\)](#), yielding a treatment allocation that attains a higher welfare-level. As shown in Table 2.5, the welfare gain is positive for all 43 villages (exceeding 100% in some villages). This indicates that if the specification of the sequential network game is correct in the context of the current application, individualized treatment allocation that takes into account network spillovers can generate large welfare gains—something that indicates that the network structure matters. Existing empirical work around social networks has not quantified the welfare gain from individualized treatment allocation under spillovers due to a lack of feasible procedures to obtain an optimal individualized assignment policy. In contrast, we uncover evidence of the welfare gains that can be realised by exploiting network spillovers.

We note that [Akbarpour et al. \(2020b\)](#) argues that the optimal treatment allocation rule under a capacity constraint may lose any advantage that it enjoys over random allocation if the capacity constraint is even slightly relaxed by a few additional households. The objective in that paper, however, is to maximize information diffusion in the context of a large network asymptotic. This is distinct from our target. We focus on the proportion of households that purchase microfinance in equilibrium. Spillover effects for product purchase are distinct from those for information diffusion both in mechanism and in intuition. In particular, we emphasize that in our setting it matters who households receive their information from and how this affects their purchase decision, rather than simply whether they are informed. We additionally note that the condition on the spillover effects that is maintained in the analysis of

[Akbarpour et al. \(2020b\)](#) becomes more restrictive when spillover effects are heterogeneous.

2.6 Conclusion

In this work, we have introduced a novel method to obtain individualized treatment allocation rules that maximize the equilibrium welfare in sequential network games. We have considered settings where the stationary joint distribution of outcomes follows a Gibbs distribution. To handle the analytical and computational challenge of analyzing the Gibbs distribution, we use variational approximation and maximize the approximated welfare criterion using a greedy maximization algorithm over treatment allocations. We have obtained bounds on the approximation error of the variational approximation and of the greedy maximization in terms of the equilibrium welfare. Moreover, we derive an upper bound on the convergence rate of the welfare regret bound. Using simulation, we have shown that our greedy algorithm performs as well as the globally optimal treatment allocation in a small network setting. In a large network setting with a given specification of parameter values, our greedy algorithm dominates random allocation and leads to a welfare improvement of around 50% compared with No treatment. We then apply our method to the Indian microfinance data ([Banerjee et al., 2013](#)). We find our method outperforms both the original allocation and random allocation across all the villages. The average welfare gain is around 40%.

We suggest that several questions remain open and that there are several ways in which our work can be extended. First, we have not considered parameter estimation in this work. A relevant question is how to incorporate the uncertainty from parameter estimation into our analysis of regret. In addition, we may want to perform inference for the welfare at the obtained assignment rule, taking into account the uncertainty of parameter estimates and a potential winner’s bias ([Andrews et al., 2020](#)). Second, to validate the iteration method for computing the variational approximation, we rely on assumptions on the spillover effect to guarantee convergence to an optimal variational approximation. Relaxing this assumption to

allow for unconstrained parameter values remains a topic for future research. Third, we have used a naive mean field method in this work. As is mentioned in [Wainwright et al. \(2008\)](#), using a structural mean field method can improve the performance of an approximation and can lead to better welfare performance.

<i>Allocation Rule</i>	<i>Network Size</i>					
	$N = 5$	$N = 7$	$N = 9$	$N = 11$	$N = 13$	$N = 15$
<i>Density = 0.3</i>						
Brute force	0.189 (< 0.01)	0.185 (< 0.01)	0.192 (< 0.01)	0.181 (< 0.01)	0.189 (< 0.01)	0.194 (< 0.01)
Brute force with var. approx.	0.187 (< 0.01)	0.185 (< 0.01)	0.192 (< 0.01)	0.181 (< 0.01)	0.189 (< 0.01)	0.194 (< 0.01)
greedy with var. approx.	0.187 (< 0.01)	0.185 (< 0.01)	0.192 (< 0.01)	0.181 (< 0.01)	0.189 (< 0.01)	0.194 (< 0.01)
<i>Density = 0.6</i>						
Brute force	0.197 (< 0.01)	0.193 (< 0.01)	0.201 (< 0.01)	0.183 (< 0.01)	0.195 (< 0.01)	0.201 (< 0.01)
Brute force with var. approx.	0.196 (< 0.01)	0.193 (< 0.01)	0.201 (< 0.01)	0.183 (< 0.01)	0.195 (< 0.01)	0.201 (< 0.01)
greedy with var. approx.	0.196 (< 0.01)	0.193 (< 0.01)	0.201 (< 0.01)	0.183 (< 0.01)	0.195 (< 0.01)	0.201 (< 0.01)

Table 2.2: Comparison for parameter set 1

Village	Sample Avg. (s.e.)	Original	Welfare under Random	Greedy (s.e.)	Welfare Gain* Level (s.e.)	Percentage
1	0.240 (0.025)	0.249 (0.030)	0.248 (0.029)	0.296 (0.034)	0.047 (0.014)	18.75%
2	0.146 (0.021)	0.188 (0.025)	0.195 (0.024)	0.216 (0.036)	0.032 (0.020)	16.81%
3	0.140 (0.023)	0.163 (0.020)	0.166 (0.022)	0.282 (0.051)	0.118 (0.041)	75.52%
4	0.077 (0.024)	0.101 (0.027)	0.106 (0.027)	0.173 (0.027)	0.073 (0.016)	72.01%
5	0.234 (0.028)	0.242 (0.035)	0.232 (0.036)	0.305 (0.044)	0.064 (0.020)	26.18%
6	0.184 (0.025)	0.156 (0.021)	0.162 (0.022)	0.279 (0.030)	0.123 (0.017)	78.57%
7	0.295 (0.038)	0.432 (0.037)	0.488 (0.046)	0.610 (0.058)	0.178 (0.040)	41.25%
8	0.118 (0.026)	0.130 (0.030)	0.146 (0.033)	0.216 (0.049)	0.086 (0.033)	66.68%
9	0.205 (0.030)	0.199 (0.024)	0.215 (0.025)	0.297 (0.034)	0.099 (0.025)	49.81%
10	0.356 (0.042)	0.364 (0.041)	0.405 (0.046)	0.482 (0.056)	0.117 (0.039)	32.18%
11	0.451 (0.046)	0.467 (0.037)	0.475 (0.038)	0.515 (0.036)	0.048 (0.014)	10.30%
12	0.153 (0.024)	0.153 (0.020)	0.150 (0.018)	0.176 (0.028)	0.023 (0.015)	14.66%
13	0.190 (0.021)	0.192 (0.029)	0.194 (0.033)	0.223 (0.040)	0.031 (0.021)	16.01%
14	0.169 (0.030)	0.188 (0.028)	0.202 (0.028)	0.277 (0.031)	0.090 (0.014)	47.82%
15	0.267 (0.020)	0.283 (0.023)	0.297 (0.024)	0.325 (0.027)	0.042 (0.013)	14.89%
16	0.354 (0.041)	0.346 (0.045)	0.343 (0.045)	0.414 (0.051)	0.068 (0.018)	19.82%
17	0.185 (0.029)	0.281 (0.022)	0.283 (0.023)	0.314 (0.024)	0.034 (0.012)	12.09%
18	0.186 (0.027)	0.203 (0.039)	0.201 (0.037)	0.245 (0.059)	0.043 (0.033)	20.82%
19	0.078 (0.017)	0.088 (0.021)	0.088 (0.021)	0.097 (0.040)	0.009 (0.033)	9.82%
20	0.193 (0.022)	0.200 (0.023)	0.190 (0.023)	0.328 (0.040)	0.128 (0.031)	64.10%
21	0.347 (0.038)	0.377 (0.038)	0.369 (0.039)	0.473 (0.038)	0.096 (0.028)	25.45%
22	0.245 (0.042)	0.245 (0.044)	0.239 (0.040)	0.282 (0.053)	0.037 (0.016)	15.09%
23	0.209 (0.027)	0.258 (0.029)	0.271 (0.029)	0.382 (0.035)	0.124 (0.019)	47.89%
24	0.235 (0.030)	0.250 (0.027)	0.235 (0.027)	0.299 (0.030)	0.053 (0.011)	21.64%
25	0.229 (0.024)	0.213 (0.037)	0.231 (0.036)	0.312 (0.046)	0.099 (0.033)	46.31%
26	0.187 (0.036)	0.218 (0.028)	0.225 (0.027)	0.314 (0.030)	0.096 (0.009)	43.95%
27	0.094 (0.019)	0.088 (0.016)	0.088 (0.017)	0.129 (0.018)	0.040 (0.013)	14.89%
28	0.123 (0.029)	0.104 (0.021)	0.112 (0.022)	0.186 (0.035)	0.082 (0.019)	14.89%
29	0.101 (0.016)	0.117 (0.024)	0.113 (0.023)	0.175 (0.045)	0.058 (0.033)	49.57%
30	0.110 (0.046)	0.174 (0.036)	0.179 (0.037)	0.203 (0.043)	0.029 (0.016)	16.75%
31	0.153 (0.023)	0.196 (0.024)	0.187 (0.024)	0.304 (0.031)	0.108 (0.016)	54.79%
32	0.084 (0.015)	0.091 (0.013)	0.086 (0.013)	0.216 (0.044)	0.125 (0.040)	137.46%
33	0.150 (0.024)	0.143 (0.019)	0.135 (0.019)	0.203 (0.037)	0.060 (0.021)	41.61%
34	0.177 (0.035)	0.230 (0.027)	0.229 (0.026)	0.274 (0.036)	0.044 (0.020)	19.16%
35	0.109 (0.021)	0.115 (0.019)	0.112 (0.020)	0.150 (0.026)	0.035 (0.016)	30.67%
36	0.167 (0.020)	0.166 (0.020)	0.177 (0.023)	0.238 (0.039)	0.072 (0.026)	43.57%
37	0.298 (0.039)	0.346 (0.026)	0.346 (0.027)	0.410 (0.027)	0.065 (0.015)	18.62%
38	0.147 (0.028)	0.215 (0.040)	0.205 (0.039)	0.251 (0.056)	0.036 (0.032)	16.44%
39	0.211 (0.025)	0.222 (0.028)	0.210 (0.026)	0.267 (0.039)	0.045 (0.026)	20.34%
40	0.164 (0.033)	0.170 (0.020)	0.179 (0.020)	0.310 (0.044)	0.140 (0.036)	81.97%
41	0.160 (0.035)	0.154 (0.030)	0.116 (0.031)	0.217 (0.031)	0.063 (0.017)	40.99%
42	0.177 (0.029)	0.167 (0.029)	0.194 (0.032)	0.228 (0.043)	0.111 (0.025)	66.38%
43	0.238 (0.036)	0.181 (0.030)	0.202 (0.034)	0.326 (0.060)	0.146 (0.038)	80.82%

Table 2.5: Comparison using 43 Indian villages microfinance data from [Banerjee et al. \(2013\)](#) * Welfare Gain compares the equilibrium welfare simulated under Greedy Allocation and the equilibrium welfare simulated under Original Allocation implemented by BSS.

Supplement to “Individualized Treatment Allocations in Sequential Network Games”

A.1 Lemma and Proposition

A.1.1 Preliminary Lemma

In this section, we collect various lemmas that we use to prove our main results. We first denote the matrix norms induced by vector norms as $\|A\|_{a,b} := \sup\{\|Ax\|_b : \|x\|_a \leq 1\}$. Let δ be a measure with full support on \mathcal{Y}^N , we denote the conditional distribution of Y_i as $\delta_i(Y_i|Y_{-i}) := \frac{\delta(Y)}{\sum_{Y_i \in \mathcal{Y}^N} \delta(Y_i, Y_{-i})}$ given the choices other units Y_{-i} . We also introduce the coupling matrix here. A matrix $A = (a_{ij})_{i,j \leq N}$ is a coupling matrix if it satisfies $a_{ii} = 0$ for all i and for $i \neq j$

$$\|\delta_i(\cdot | Y_{-i}) - \delta_i(\cdot | Y'_{-i})\|_{\text{TV}} \leq a_{ij},$$

whenever $Y, Y' \in \mathcal{Y}^N$ differ only at the j -th coordinate.

Lemma A.1.1. *Let P be the stationary distribution P , defined in Eq.2.3. Then, the coupling matrix is given by J , with each element $J_{ij} = \frac{A_N}{4} m_{ij} G_{ij}(|\theta_5| + |\theta_6|)$. Under Assumption 9,*

$$\|J\|_{2,2} \leq \|J\|_{\infty,\infty} \leq 1 - \alpha \quad \text{holds for } P. \quad (\text{A.18})$$

Moreover,

$$\delta_i(\cdot \mid Y_{-i}) \in (1 - C_\alpha, C_\alpha) \text{ for some } C_\alpha \text{ depending only on } \alpha \text{ and ,}$$

uniformly in i, N , and Y_{-i} .

Proof.

$$\begin{aligned} \delta_i(1 \mid Y_{-i}) &= \frac{\pi(1, Y_{-i})}{\pi(1, Y_{-i}) + \pi(0, Y_{-i})}, \\ \delta_i(0 \mid Y_{-i}) &= \frac{\pi(0, Y_{-i})}{\pi(1, Y_{-i}) + \pi(0, Y_{-i})}, \end{aligned}$$

where

$$\begin{aligned} \pi(1, Y_{-i}) &= \exp \left[\alpha_i + \sum_{j \neq i} \alpha_j Y_j + \frac{A_N}{2} \sum_{\ell=1}^N \sum_{j=1}^N m_{\ell j} G_{\ell j} Y_\ell Y_j (\theta_5 + \theta_6 d_\ell d_j) \right], \\ \pi(0, Y_{-i}) &= \exp \left[\sum_{j \neq i} \alpha_j Y_j + \frac{A_N}{2} \sum_{\ell \neq i} \sum_{j \neq i} m_{\ell j} G_{\ell j} Y_\ell Y_j (\theta_5 + \theta_6 d_\ell d_j) \right], \end{aligned}$$

where $\alpha_i = \theta_0 + \theta_1 d_i + X'_i(\theta_2 + \theta_3 d_i) + A_N \sum_{j=1}^N \theta_4 m_{ij} G_{ij} d_j$. Denote $\gamma(x) = \frac{\exp(x)}{1 + \exp(x)}$ as a sigmoid function of $x \in \mathbb{R}$, then we have:

$$\begin{aligned} \delta_i(1 \mid \bar{Y}_i) &= \gamma \left(\alpha_i + A_N \sum_{j=1}^N m_{ij} G_{ij} Y_j (\theta_5 + \theta_6 d_i d_j) \right). \\ \delta_i(0 \mid \bar{Y}_i) &= 1 - \delta_i(1 \mid \bar{Y}_i). \end{aligned} \tag{A.19}$$

Therefore,

$$d_{TV}(\delta_i(\cdot \mid \bar{z}_i), \delta_i(\cdot \mid \bar{y}_i)) = |\delta_i(1 \mid \bar{z}_i) - \delta_i(1 \mid \bar{y}_i)| = |\gamma(g_i(Y)) - \gamma(g_i(T_k Y))|,$$

where $g_i(Y) = \alpha_i + A_N \sum_{j=1}^N m_{ij} G_{ij} Y_j (\theta_5 + \theta_6 d_i d_j)$. Therefore

$$d_{TV}(\delta_i(\cdot \mid \bar{z}_i), \delta_i(\cdot \mid \bar{y}_i)) \leq \frac{1}{4} |g_i(y) - g_i(T_k y)| \leq \frac{1}{4} |A_N m_{ik} G_{ik} (\theta_5 + \theta_6)|.$$

Thus, the coupling matrix is given by J , with each element $J_{ij} = \frac{A_N}{4} m_{ij} G_{ij} (|\theta_5| + |\theta_6|)$. First inequality in Eq.A.18 holds since:

$$\|J\|_{2,2} \leq \sqrt{\|J\|_{\infty,\infty} \|J^T\|_{\infty,\infty}} = \|J\|_{\infty,\infty},$$

where the eigenvalue $|\lambda_i(JJ^T)| \leq \|JJ^T\| \leq \|J\| \|J^T\|$ for any operator norm, and J is symmetric. Under Assumption 9,

$$\|J\|_{\infty,\infty} = \frac{A_N}{4} (|\theta_5| + |\theta_6|) \max_{i=1,\dots,N} \sum_{j=1}^N m_{ij} G_{ij} \leq 1 - \alpha.$$

The second statement follows by using Eq.A.19:

$$\delta_i(1|Y_{-i}) \leq \gamma(|\theta_0| + |\theta_1| + \bar{X}'|\theta_2| + \bar{X}'|\theta_3| + A_N \bar{N} \bar{m} |\theta_4| + 4(1 - \alpha)) =: C_\alpha.$$

Therefore, $\delta_i(0 | Y_{-i}) \leq 1 - C_\alpha$. Without loss of generality, we choose $1 - C_\alpha \leq C_\alpha$, and finish the proof. \square

Lemma A.1.2. *For any $\tau > 0$, there is a finite set of $N \times 1$ vectors $\mathcal{W}(\tau)$ such that*

$$|\mathcal{W}(\tau)| \leq 2^N,$$

and for any $N \times 1$ vector M with entries in $\{0, 1\}$, there exists a $W \in \mathcal{W}(\tau)$ such that

$$\sum_i (M_i - W_i)^2 \leq \tau^2.$$

Proof. Since all the entries of M in $\{0, 1\}$, M must be a vertex of the N -dimensional unit hypercube $[0, 1]^N$. As such, we let \mathcal{W} be the collection of all vertices of the N -dimensional

unit hypercube. For any M , we can always find an element in \mathcal{W} such that:

$$\sum_i (M_i - W_i)^2 = 0 \leq \tau^2.$$

Then $|\mathcal{W}(\tau)| = 2^N$. □

Lemma A.1.3. *Suppose Assumptions 4 to 9 hold. Let $R \subset \mathcal{N}$ be a treatment allocation set, $R = \{i \in \mathcal{N} : d_i = 1\}$ such that $\mathcal{N} \setminus R \neq \emptyset$. Given $k \in \mathcal{N} \setminus R$, let $\tilde{\mu}_i$, $i = 1, \dots, N$, be a solution of the first-order conditions of Eq.2.11 when the treatment allocation set is $R \cup \{k\}$ and $\tilde{\mu}'_i$, $i = 1, \dots, N$, be a solution of Eq.2.11 when the treatment allocation set is R , i.e.,*

$$\tilde{\mu}_i = \Lambda(\theta_0 + \theta_1 d_i + X'_i(\theta_2 + \theta_3 d_i) + A_N \theta_5 \sum_{j \neq i} m_{ij} G_{ij} \tilde{\mu}_j + A_N \sum_{\substack{j \neq i \\ j \in R}} m_{ij} G_{ij} (\theta_4 + \theta_6 d_i \tilde{\mu}_j) + \mathcal{M}_{ik}(\tilde{\mu})),$$

$$\tilde{\mu}'_i = \Lambda(\theta_0 + \theta_1 d_i + X'_i(\theta_2 + \theta_3 d_i) + A_N \theta_5 \sum_{j \neq i} m_{ij} G_{ij} \tilde{\mu}'_j + A_N \sum_{\substack{j \neq i \\ j \in R}} m_{ij} G_{ij} (\theta_4 + \theta_6 d_i \tilde{\mu}'_j)),$$

where $\mathcal{M}_{ik} : [0, 1]^N \rightarrow \mathbb{R}_+$ is defined by

$$\mathcal{M}_{ik}(\mu) = A_N m_{ik} G_{ik} (\theta_4 + \theta_6 d_i \mu_k) + \mathbb{1}\{i = k\} (X'_i \theta_3 + A_N \sum_{j \in R} \theta_6 m_{ij} G_{ij} \mu_j).$$

Then, $\tilde{\mu}_i \geq \tilde{\mu}'_i$ holds for all $i \in \mathcal{N}$ at any $\mathcal{X} \in \mathbb{R}^{N \times k}$ and $G \in \{0, 1\}^{N \times N}$.

Proof. Let us define $\tilde{\mu}$ as a vector with elements $\{\tilde{\mu}_i\}_{i=1}^N$ and $\tilde{\mu}'$ as a vector with elements $\{\tilde{\mu}'_i\}_{i=1}^N$. By the assumption that $\theta_3, \theta_4, \theta_6 \geq 0$, \mathcal{M}_{ik} is non-negative for all $i \in \mathcal{N}$. We define

$$\tilde{\mu}_i^1 = \Lambda(\theta_0 + \theta_1 d_i + X'_i(\theta_2 + \theta_3 d_i) + A_N \theta_5 \sum_{j \neq i} m_{ij} G_{ij} \tilde{\mu}'_j + A_N \sum_{\substack{j \neq i \\ j \in R}} m_{ij} G_{ij} (\theta_4 + \theta_6 d_i \tilde{\mu}'_j) + \mathcal{M}_{ik}(\tilde{\mu}')),$$

for all $i \in \mathcal{N}$, where $\tilde{\mu}_i^1 \geq \tilde{\mu}'_i$. Then, we use $\tilde{\mu}^1$ to generate $\tilde{\mu}^2$:

$$\tilde{\mu}_i^2 = \Lambda(\theta_0 + \theta_1 d_i + X'_i(\theta_2 + \theta_3 d_i) + A_N \theta_5 \sum_{j \neq i} m_{ij} G_{ij} \tilde{\mu}_j^1 + A_N \sum_{\substack{j \neq i \\ j \in R}} m_{ij} G_{ij} (\theta_4 + \theta_6 d_i \tilde{\mu}_j^1) + \mathcal{M}_{ik}(\tilde{\mu}^1)),$$

for all $i \in \mathcal{N}$, where $\tilde{\mu}_i^2 \geq \tilde{\mu}_i^1$. We iterate the above process until it converges. If this iteration is a contraction mapping, we can guarantee convergence to $\tilde{\mu}$, which is a unique equilibrium (by Theorem A.3.1). We denote this iteration process as $\{\tilde{\mu}^t\}_{t=1}^T$ and show that the above mapping $T : [0, 1]^N \rightarrow [0, 1]^N$ is a contraction mapping; we equate the distance metric with the ℓ_1 -distance. For any $t \geq 1$,

$$d(T(\tilde{\mu}^t), T(\tilde{\mu}^s)) = \sum_{i=1}^N |\tilde{\mu}_i^{t+1} - \tilde{\mu}_i^{s+1}|,$$

$$d(\tilde{\mu}^t, \tilde{\mu}^s) = \sum_{i=1}^N |\tilde{\mu}_i^t - \tilde{\mu}_i^s|.$$

First, we know that $\Lambda(\cdot)$ is a sigmoid function. Hence, its steepest slope is 0.25.

$$\tilde{\mu}_i^{t+1} - \tilde{\mu}_i^{s+1} \leq 0.25 A_N \sum_{j \neq i} (|\theta_5| + |\theta_6| d_i) m_{ij} G_{ij} (\tilde{\mu}_j^t - \tilde{\mu}_j^s)$$

(By Multivariate Mean Value Theorem)

$$\leq 0.25 A_N (|\theta_5| + |\theta_6|) \sum_{j \neq i} m_{ij} G_{ij} (\tilde{\mu}_j^t - \tilde{\mu}_j^s)$$

$$\leq 0.25 A_N (|\theta_5| + |\theta_6|) \sum_{j \neq i} m_{ij} G_{ij} |\tilde{\mu}_j^t - \tilde{\mu}_j^s|.$$

Therefore,

$$|\tilde{\mu}_i^{t+1} - \tilde{\mu}_i^{s+1}| \leq \frac{A_N (|\theta_5| + |\theta_6|)}{4} \sum_{j \neq i} m_{ij} G_{ij} |\tilde{\mu}_j^t - \tilde{\mu}_j^s|.$$

Hence,

$$\begin{aligned} \sum_{i=1}^N |\tilde{\mu}_i^{t+1} - \tilde{\mu}_i^{s+1}| &\leq \frac{A_N(|\theta_5| + |\theta_6|)}{4} \sum_{i=1}^N \sum_{j \neq i} m_{ij} G_{ij} |\tilde{\mu}_j^t - \tilde{\mu}_j^s| \\ &\leq \frac{A_N(|\theta_5| + |\theta_6|)}{4} \sum_{i=1}^N |\tilde{\mu}_i^t - \tilde{\mu}_i^s| \max_{i \in \mathcal{N}} \sum_{j \neq i} m_{ij} G_{ij}. \end{aligned}$$

Under Assumption 9, we can show that T is a contraction mapping. In addition, since $\tilde{\mu}^t \in [0, 1]$ for all $t \geq 1$, the metric space $(\tilde{\mu}, d)$ is a complete metric space. By Theorem A.3.1, $\{\tilde{\mu}', \tilde{\mu}^1, \dots\}$ converges to $\tilde{\mu}$. Since $\mathcal{M}_{ik} \in [0, +\infty)$ for all $i \in \mathcal{N}$, this result applies for any \mathcal{M}_{ik} . We conclude that $\tilde{\mu}_i \geq \tilde{\mu}_i^{T-1} \geq \dots \geq \tilde{\mu}_i', \forall i \in \mathcal{N} \setminus \{k\}$. \square

Lemma A.1.4. Under Assumptions 4 to 11, the upper bound on the curvature ξ^{up} and the lower bound on the submodularity ratio γ^{low} is given by:

$$\xi \leq \xi^{up} = 1 - \zeta \in (0, 1),$$

$$\gamma \geq \gamma^{low} = \zeta \in (0, 1),$$

where

$$\zeta = \frac{1}{N} \min\{\Lambda'(\theta_0 + \underline{X'\theta_2}), \Lambda'(\theta_0 + \theta_1 + \overline{X'\theta_2} + \overline{X'\theta_3} + A_N(\theta_4 + \theta_5 + \theta_6)\overline{mN})\}(A_N\theta_4\underline{N} \cdot \underline{m} + \theta_1),$$

and $\underline{X'\theta_2} = \min_{i \in \mathcal{N}} X'_i\theta_2$, $\overline{X'\theta_2} = \max_{i \in \mathcal{N}} X'_i\theta_2$, $\overline{X'\theta_3} = \arg \max_{i \in \mathcal{N}} X'_i\theta_3$.

Proof. Curvature:

The curvature is defined as the smallest value of ξ such that

$$\tilde{W}(R \cup \{k\}) - \tilde{W}(R) \geq (1 - \xi)[\tilde{W}(S \cup \{k\}) - \tilde{W}(S)] \quad \forall S \subseteq R \subseteq \mathcal{N}, \forall k \in \mathcal{N} \setminus R.$$

As a consequence,

$$\xi = \max_{S \subseteq R \subseteq \mathcal{N}, k \in \mathcal{N} \setminus R} 1 - \frac{\tilde{W}(R \cup \{k\}) - \tilde{W}(R)}{\tilde{W}(S \cup \{k\}) - \tilde{W}(S)}. \quad (\text{A.20})$$

We first denote $\tilde{W}_i(\mathcal{D})$ as:

$$\begin{aligned}\tilde{W}_i(\mathcal{D}) &:= \mathbb{1}_{\{i \in \mathcal{D}\}} \Lambda [\theta_0 + \theta_1 + X'_i(\theta_2 + \theta_3) + A_N \theta_5 \sum_{\substack{j \neq i \\ j \in \mathcal{N}}} m_{ij} G_{ij} \tilde{\mu}_j + A_N \sum_{\substack{j \neq i \\ j \in \mathcal{D}}} m_{ij} G_{ij} (\theta_4 + \theta_6 \tilde{\mu}_j)] \\ &\quad + \mathbb{1}_{\{k \in \mathcal{N} \setminus \mathcal{D}\}} \Lambda [\theta_0 + X'_k \theta_2 + A_N \theta_4 \sum_{\ell \in \mathcal{D}} m_{k\ell} G_{k\ell} + A_N \theta_5 \sum_{\substack{\ell \neq k \\ \ell \in \mathcal{N}}} m_{k\ell} G_{k\ell} \tilde{\mu}_\ell].\end{aligned}$$

We can upper bound the denominator in Eq.A.20 by:

$$\tilde{W}(S \cup \{k\}) - \tilde{W}(S) = \sum_{i=1}^N \tilde{W}_i(S \cup \{k\}) - \tilde{W}_i(S) \leq N,$$

A lower bound on the numerator in Eq.A.20 is:

$$\Delta F \doteq \tilde{W}(R \cup \{k\}) - \tilde{W}(R) \quad \forall k \in \mathcal{N}, R \subseteq \mathcal{N}.$$

Then we define ξ^{up} as:

$$\xi^{up} \equiv 1 - \frac{\Delta F}{N} \geq \xi.$$

Let $\tilde{\mu}$ denote the solution of Eq.2.10 with $d_i = 1$ for $i \in R \cup \{k\}$ and $\tilde{\mu}'$ denote the solution of Eq.2.10 with $d_i = 1$ for $i \in R$. We first rewrite $\tilde{W}(R \cup \{k\}) - \tilde{W}(R)$ with a view to deriving a lower bound. To do so, we define

$$\bar{\phi}_i(\mu) := \theta_0 + \theta_1 + X'_i(\theta_2 + \theta_3) + A_N \theta_5 \sum_{j \neq i} m_{ij} G_{ij} \mu_j + A_N \sum_{\substack{j \neq i \\ j \in R}} m_{ij} G_{ij} (\theta_4 + \theta_6 \mu_j) + A_N m_{ik} G_{ik} (\theta_4 + \theta_6 \mu_k);$$

$$\underline{\phi}_i(\mu) := \theta_0 + \theta_1 + X'_i(\theta_2 + \theta_3) + A_N \theta_5 \sum_{j \neq i} m_{ij} G_{ij} \mu_j + A_N \sum_{\substack{j \neq i \\ j \in R}} m_{ij} G_{ij} (\theta_4 + \theta_6 \mu_j);$$

$$\bar{\varphi}_c(\mu) := \theta_0 + X'_c \theta_2 + A_N \theta_5 \sum_{b \neq c} G_{cb} m_{bc} \mu_b + A_N \theta_4 \sum_{z \in R} m_{cz} G_{cz} + A_N \theta_4 m_{ck} G_{ck};$$

$$\underline{\varphi}_c(\mu) := \theta_0 + X'_c \theta_2 + A_N \theta_5 \sum_{b \neq c} G_{cb} m_{bc} \mu_b + A_N \theta_4 \sum_{z \in R} m_{cz} G_{cz};$$

$$\bar{\psi}_k(\mu) := \theta_0 + \theta_1 + X'_k(\theta_2 + \theta_3) + A_N \theta_5 \sum_{e \neq k} m_{ke} G_{ke} \mu_e + A_N \sum_{l \in R} m_{kl} G_{kl}(\theta_4 + \theta_6 \mu_l);$$

$$\underline{\psi}_k(\mu) := \theta_0 + X'_k \theta_2 + A_N \theta_5 \sum_{e \neq k} m_{ke} G_{ke} \mu_e + A_N \theta_4 \sum_{l \in R} m_{kl} G_{kl}$$

Then, we have

$$\begin{aligned} & \tilde{W}(R \cup \{k\}) - \tilde{W}(R) \\ &= \sum_{i \in R} \left[\Lambda(\bar{\phi}_i(\tilde{\mu})) - \Lambda(\underline{\phi}_i(\tilde{\mu}')) \right] + \sum_{c \in \mathcal{N} \setminus R \cup \{k\}} \left[\Lambda(\bar{\varphi}_c(\tilde{\mu})) - \Lambda(\underline{\varphi}_c(\tilde{\mu}')) \right] + \Lambda(\bar{\psi}_k(\tilde{\mu})) - \Lambda(\underline{\psi}_k(\tilde{\mu}')). \end{aligned} \quad (\text{A.21})$$

By Lemma A.1.3, we can bound the above equation from below by replacing $\tilde{\mu}'_i$ with $\tilde{\mu}_i$ for all $i \in \mathcal{N}$. Then, Eq.A.21 is bounded by:

$$\begin{aligned} & \tilde{W}(R \cup \{k\}) - \tilde{W}(R) \\ & \geq \sum_{i \in R} \left[\Lambda(\bar{\phi}_i(\tilde{\mu})) - \Lambda(\underline{\phi}_i(\tilde{\mu})) \right] + \sum_{c \in \mathcal{N} \setminus R \cup \{k\}} \left[\Lambda(\bar{\varphi}_c(\tilde{\mu})) - \Lambda(\underline{\varphi}_c(\tilde{\mu})) \right] + \Lambda(\bar{\psi}_k(\tilde{\mu})) - \Lambda(\underline{\psi}_k(\tilde{\mu})). \end{aligned}$$

Using the mean value theorem, and letting $\phi_i \in (\underline{\phi}_i(\tilde{\mu}), \bar{\phi}_i(\tilde{\mu}))$, $\varphi_c \in (\underline{\varphi}_c(\tilde{\mu}), \bar{\varphi}_c(\tilde{\mu}))$, $\psi_k \in$

$(\underline{\psi}_k(\tilde{\mu}), \overline{\psi}_k(\tilde{\mu}))$, we have:

$$\begin{aligned}
& \tilde{W}(R \cup \{k\}) - \tilde{W}(R) \\
& \geq \sum_{i \in R} \Lambda'(\phi_i) A_N (\theta_4 + \theta_6 \tilde{\mu}_k) m_{ik} G_{ik} + \sum_{c \in \mathcal{N} \setminus R \cup \{k\}} \Lambda'(\varphi_c) A_N \theta_4 m_{ck} G_{ck} \\
& \quad + \Lambda'(\psi_k) (\theta_1 + X'_k \theta_3 + \theta_6 A_N \sum_{l \in R} m_{kl} G_{kl} \tilde{\mu}_l) \\
& \geq \sum_{i \in R} \Lambda'(\phi_i) A_N \theta_4 m_{ik} G_{ik} + \sum_{c \in \mathcal{N} \setminus R \cup \{k\}} \Lambda'(\varphi_c) A_N \theta_4 m_{ck} G_{ck} + \Lambda'(\psi_k) \theta_1 \tag{A.22} \\
& \geq \sum_{i \in R} \underline{\Lambda'}(\phi_i) A_N \theta_4 m_{ik} G_{ik} + \sum_{c \in \mathcal{N} \setminus R \cup \{k\}} \underline{\Lambda'}(\varphi_c) A_N \theta_4 m_{ck} G_{ck} + \underline{\Lambda'}(\psi_k) \theta_1 \\
& \geq \underline{\Lambda'} \cdot \left(\sum_{i \neq k} A_N \theta_4 m_{ik} G_{ik} + \theta_1 \right) \\
& \geq \underline{\Lambda'} \cdot (A_N \theta_4 \underline{N} \cdot \underline{m} + \theta_1),
\end{aligned}$$

where

$$\underline{\Lambda'}(\phi_i) = \min\{\Lambda'(\underline{\phi}_i), \Lambda'(\overline{\phi}_i)\},$$

$$\underline{\Lambda'}(\varphi_c) = \min\{\Lambda'(\underline{\varphi}_c), \Lambda'(\overline{\varphi}_c)\},$$

$$\underline{\Lambda'}(\psi_k) = \min\{\Lambda'(\underline{\psi}_k), \Lambda'(\overline{\psi}_k)\},$$

$$\underline{\Lambda'} = \min\{\Lambda'(\theta_0 + \underline{X}'\theta_2), \Lambda'(\theta_0 + \theta_1 + \overline{X}'\theta_2 + \overline{X}'\theta_3 + A_N(\theta_4 + \theta_5 + \theta_6)\overline{m}\overline{N})\}.$$

Then,

$$\xi^{up} = 1 - \frac{1}{N} \underline{\Lambda'} \cdot (A_N \theta_4 \underline{N} \cdot \underline{m} + \theta_1).$$

Under Assumption 5 and 10

$$\theta_1, \theta_4, A_N, \underline{N}, \underline{m} > 0 \Rightarrow (A_N \theta_4 \underline{N} \cdot \underline{m} + \theta_1) > 0.$$

In addition, we know that $\Lambda(x)$ is a logistic function, and so

$$\begin{aligned}\Lambda'(\theta_0 + \underline{X}'\theta_2) &\in (0, 0.25], \\ \Lambda'(\theta_0 + \theta_1 + \overline{X}'\theta_2 + \overline{X}'\theta_3 + A_N(\theta_3 + \theta_4 + \theta_5)\overline{m}\overline{N}) &\in (0, 0.25].\end{aligned}$$

Hence,

$$\underline{\Lambda}' \in (0, 0.25).$$

By assuming that $(A_N\theta_4\underline{N} \cdot \underline{m} + \theta_1) \leq 4N$, we obtain

$$\underline{\Lambda}' \cdot (A_N\theta_4\underline{N} \cdot \underline{m} + \theta_1) < N.$$

We conclude that $\xi^{up} < 1$. From [Bian et al. \(2017\)](#), ξ only equal to 0 if $\tilde{W}(\mathcal{D})$ is supermodular. Since $\tilde{W}(\mathcal{D})$ is not supermodular, $\xi^{up} > \xi > 0$.

Submodularity Ratio

The submodularity ratio of a non-negative set function is the largest γ such that

$$\sum_{k \in R \setminus S} [\tilde{W}(S \cup \{k\}) - \tilde{W}(S)] \geq \gamma [\tilde{W}(S \cup R) - \tilde{W}(S)], \quad \forall S, R \subseteq \mathcal{N}.$$

As a consequence,

$$\begin{aligned}\gamma &= \min_{S \neq R} \frac{\sum_{k \in R \setminus S} [\tilde{W}(S \cup \{k\}) - \tilde{W}(S)]}{\tilde{W}(S \cup R) - \tilde{W}(S)} \\ &= \min_{S \neq R} \frac{\sum_{k \in R \setminus S} \sum_i [\tilde{W}_i(S \cup \{k\}) - \tilde{W}_i(S)]}{\sum_i [\tilde{W}_i(S \cup R) - \tilde{W}_i(S)]}\end{aligned} \tag{A.23}$$

We can upper bound the denominator in Eq.A.23 by:

$$\sum_{i=1}^N [\tilde{W}_i(S \cup R) - \tilde{W}_i(S)] \leq N.$$

We first rewrite $\sum_{k \in R \setminus S} [\tilde{W}(S \cup \{k\}) - \tilde{W}(S)]$ with a view to deriving a lower bound.

$$\begin{aligned} & \sum_{k \in R \setminus S} \sum_i [\tilde{W}_i(S \cup \{k\}) - \tilde{W}_i(S)] \\ = & \sum_{k \in R \setminus S} \left[\sum_{i \in S} \left[\Lambda(\theta_0 + \theta_1 + X'_i(\theta_2 + \theta_3) + A_N \theta_5 \sum_{j \neq i} m_{ij} G_{ij} \tilde{\mu}_j + A_N \sum_{\substack{j \neq i \\ j \in S}} m_{ij} G_{ij} (\theta_4 + \theta_6 \tilde{\mu}_j) \right. \right. \\ & \left. \left. + A_N m_{ik} G_{ik} (\theta_4 + \theta_6 \tilde{\mu}_k) \right) - \Lambda(\theta_0 + \theta_1 + X'_i(\theta_2 + \theta_3) + A_N \theta_5 \sum_{j \neq i} m_{ij} G_{ij} \tilde{\mu}'_j + A_N \sum_{\substack{j \neq i \\ j \in S}} m_{ij} G_{ij} (\theta_4 + \theta_6 \tilde{\mu}'_j)) \right] \\ & + \sum_{c \in \mathcal{N} \setminus S \cup \{k\}} \left[\Lambda(\theta_0 + X'_c \theta_2 + A_N \theta_5 \sum_{b \neq c} G_{cb} m_{bc} \tilde{\mu}_b + A_N \theta_4 \sum_{z \in S} m_{cz} G_{cz} + A_N \theta_4 m_{ck} G_{ck}) \right. \\ & \left. - \Lambda(\theta_0 + X'_c \theta_2 + A_N \theta_5 \sum_{b \neq c} G_{cb} m_{bc} \tilde{\mu}'_b + A_N \theta_4 \sum_{z \in S} m_{cz} G_{cz}) \right] \\ & + \Lambda(\theta_0 + \theta_1 + X'_k(\theta_2 + \theta_3) + A_N \theta_5 \sum_{e \neq k} m_{ke} G_{ke} \tilde{\mu}_e + A_N \sum_{l \in S} m_{kl} G_{kl} (\theta_4 + \theta_6 \tilde{\mu}_l)) \\ & \left. - \Lambda(\theta_0 + X'_k \theta_2 + A_N \theta_5 \sum_{e \neq k} m_{ke} G_{ke} \tilde{\mu}'_e + A_N \theta_4 \sum_{l \in S} m_{kl} G_{kl}) \right]. \end{aligned}$$

We can lower bound the numerator in Eq.A.23 by:

$$\begin{aligned} \sum_{k \in R \setminus S} \sum_i [\tilde{W}_i(S \cup \{k\}) - \tilde{W}_i(S)] & \geq \sum_{k \in R \setminus S} \underline{\Lambda}' \cdot (A_N \theta_4 \underline{N} \cdot \underline{m} + \theta_1) \\ & \quad (\text{By Eq.A.22}) \\ & \geq \underline{\Lambda}' \cdot (A_N \theta_4 \underline{N} \cdot \underline{m} + \theta_1) \end{aligned}$$

Defining:

$$\gamma^{low} = \frac{1}{N} \underline{\Lambda}' \cdot (A_N \theta_4 \underline{N} \cdot \underline{m} + \theta_1),$$

the submodularity ratio can be bounded from below by:

$$\gamma \geq \gamma^{low} = 1 - \xi^{up} \in (0, 1).$$

□

A.1.2 Proof of Proposition 2.2.1

Under Assumptions 4 and 8, the potential function $\Phi(y, \mathcal{X}, D, G; \theta)$ for $U_i(y, \mathcal{X}, D, G; \theta)$ can be defined as:

$$\begin{aligned} \Phi(y, \mathcal{X}, D, G; \theta) = & \sum_{i=1}^N \left(\theta_0 + \theta_1 d_i + X'_i(\theta_2 + \theta_3 d_i) + A_N \sum_{j=1}^N \theta_4 m_{ij} G_{ij} d_j \right) y_i \\ & + \frac{A_N}{2} \sum_{i=1}^N \sum_{j=1}^N m_{ij} G_{ij} y_i y_j (\theta_5 + \theta_6 d_i d_j), \end{aligned}$$

and our interacted decision game is a potential game.

Proof. A potential function is a function $\Phi : Y \rightarrow \mathbb{R}$ such that:

$$\Phi(y_i = 1, y_{-i}, \mathcal{X}, D, G) - \Phi(y_i = 0, y_{-i}, \mathcal{X}, D, G) = U_i(y_i = 1, y_{-i}, \mathcal{X}, D, G) - U_i(y_i = 0, y_{-i}, \mathcal{X}, D, G).$$

Simple computation shows that for any $i \in \mathcal{N}$,

$$\begin{aligned} & \Phi(y_i = 1, y_{-i}, \mathcal{X}, D, G) - \Phi(y_i = 0, y_{-i}, \mathcal{X}, D, G) \\ &= \theta_0 + \theta_1 d_i + X'_i(\theta_2 + \theta_3 d_i) + A_N \sum_{j=1}^N \theta_4 m_{ij} G_{ij} d_j + A_N \sum_{j=1}^N m_{ij} G_{ij} y_j (\theta_5 + \theta_6 d_i d_j) \\ &= U_i(y_i = 1, y_{-i}, \mathcal{X}, D, G) - U_i(y_i = 0, y_{-i}, \mathcal{X}, D, G). \end{aligned}$$

Therefore, Φ is the potential of our interacted decision game.

□

A.1.3 Proof of Lemma 2.3.1

Lemma A.1.5. Error Bound: Let P and Q be two probability distributions over $\{0, 1\}^N$. We have:

$$\sum_{i=1}^N |\mu_i^P - \mu_i^Q| \leq W_1^H(P, Q).$$

Proof. For any coupling $\omega \in \Omega(P, Q)$ (i.e. a joint distribution on (Y, Y') with marginals P and Q), define the Hamming cost

$$\mathbb{E}_\omega \left[\sum_{i=1}^N \mathbf{1}\{Y_i \neq Y'_i\} \right] = \sum_{i=1}^N \mathbb{E}_\omega [\mathbf{1}\{Y_i \neq Y'_i\}] = \sum_{i=1}^N Pr_\omega(Y_i \neq Y'_i).$$

For each coordinate i , given Y_i is Bernoulli random variable, we have

$$Pr_\omega(Y_i \neq Y'_i) \geq |P(Y_i = 1) - Q(Y_i = 1)|.$$

Summing over $i = 1, \dots, N$, we get

$$\mathbb{E}_\omega \left[\sum_{i=1}^N \mathbf{1}\{Y_i \neq Y'_i\} \right] \geq \sum_{i=1}^N |P(Y_i = 1) - Q(Y_i = 1)|.$$

Since $W_1^H(P, Q)$ is the infimum of the above expected Hamming cost over all couplings $\gamma \in \Omega(P, Q)$, the result follows:

$$W_1^H(P, Q) = \inf_{\omega \in \Omega(P, Q)} \mathbb{E}_\omega \left[\sum_{i=1}^N \mathbf{1}\{Y_i \neq Y'_i\} \right] \geq \sum_{i=1}^N |P(Y_i = 1) - Q(Y_i = 1)|.$$

□

A.1.4 Proof of Proposition 2.3.1

Proposition A.1.1. *Transportation Inequality:* Under Assumption 9, there is a universal constant C_t such that the stationary distribution P , defined in Eq. 2.3, satisfies:

$$W_1(P, Q) \leq C_t \sqrt{\mathbb{KL}(Q \| P)},$$

for all probability measures Q on $\{0, 1\}^N$ with a finite first moment.

Proof. We first introduce the definition of approximate tensorization property:

Definition A.1.1 (Approximate tensorization of entropy). We say that a measure μ has the *approximate tensorization property* with constant C (abbreviated $AT(C)$) if for every function $f : \mathcal{Y} \rightarrow [0, \infty)$,

$$\text{Ent}_\mu(f) \leq C \mathbb{E}_\mu \left[\sum_{i=1}^n \text{Ent}_{\mu(\cdot | Y_{-i})}(f) \right],$$

where $\text{Ent}_\mu(f) := \mathbb{E}_\mu[f \log f] - \mathbb{E}_\mu[f] \log(\mathbb{E}_\mu[f]) \in [0, \infty)$ is the entropy functional of nonnegative function f .

We apply Lemma A.1.1 to see that by Lemma A.3.1 we have for some $\chi = \chi(\theta, X, G, \alpha)$ such that the stationary distribution P satisfies the approximate tensorization property $AT(\frac{1}{\chi\alpha^2})$. i.e.,

$$\text{Ent}_P(f) \leq \frac{1}{\chi\alpha^2} \mathbb{E}_P \left[\sum_{i=1}^n \text{Ent}_{P(\cdot | Y_{-i})}(f) \right].$$

In addition, the second statement of Lemma A.1.1 guarantees that $\chi \geq (1 - C_\alpha)$. Therefore, P satisfies the approximate tensorization property $AT(\frac{1}{(1 - C_\alpha)\alpha^2})$. Lemma A.3.2 guarantees the stationary distribution P satisfies the dimension-free convex concentration inequality with constant K depending only on $1/(1 - C_\alpha)\alpha^2$, where the definition of convex concentration property is given by

Definition A.1.2. Convex Concentration Property: a random vector X in \mathbb{R}^N has the convex concentration property with constant K if for any L -Lipschitz convex function $f : \mathbb{R}^N \rightarrow \mathbb{R}$,

and any $t > 0$,

$$\Pr(|f(X) - \text{Med } f(X)| \geq t) \leq 2 \exp\left(-\frac{t^2}{K^2 L^2}\right). \quad (5.1)$$

Apply Lemma A.3.3, we conclude there exists constant C_T such that P satisfies inequality \bar{T}_f with $f(x) = C_T \|x\|_1^2$, where the definition of \bar{T}_f is given by:

Definition A.1.3. Weak transport cost: Let μ and ν be probability measures on \mathbb{R}^N , and assume $\nu \in \mathbb{R}^N$ with finite first moment. define the weak transport cost between μ and ν as

$$\bar{T}_f(Q | P) = \inf_{\pi} \int_{\{0,1\}^N} f(y - \int_{\{0,1\}} y' p_y(dy')) P(dy).$$

where $f : \mathbb{R}^N \rightarrow [0, +\infty]$ is a lower-semicontinuous convex function, and the infimum is taken over all couplings π of μ and ν (i.e. measures on $\mathbb{R}^N \times \mathbb{R}^N$ with marginals μ and ν). For each $x \in \mathbb{R}^N$, $p_x(\cdot)$ denotes the conditional measure given by $\pi(dx dy) = p_x(dy) \mu(dx)$ (μ -almost surely). We will say that μ satisfies the inequality \bar{T}_f if for every probability measure $\nu \in \mathbb{R}^N$ with finite first moment,

$$\max(\bar{T}_f(\nu | \mu), \bar{T}_f(\mu | \nu)) \leq \mathbb{KL}(\nu | \mu).$$

Applying $f(x) = C_T \|x\|_1^2$, there exists constant C_T such that

$$\bar{T}_f(Q | P) = \inf_{\pi} \int_{\{0,1\}^N} C_T \left\| y - \int_{\{0,1\}} y' p_y(dy') \right\|_1^2 P(dy). \quad (A.24)$$

Therefore, we have

$$\sqrt{\frac{1}{C_T} \bar{T}_f(Q | P)} \leq \sqrt{\frac{1}{C_T} \mathbb{KL}(Q | P)}. \quad (A.25)$$

Now I claim: Suppose we have probability measures P, Q on $\{0, 1\}^N$. Given the Wasserstein 1-distance equipped with Hamming distance, defined in Definition 2.3.1, and the weak

transport cost is defined in Eq.A.24, we have

$$W_1(P, Q) \leq \sqrt{\frac{1}{C_T} \bar{T}_f(Q | P)}.$$

Combining with Eq.A.25, and denote $c_t = \frac{1}{\sqrt{C_T}}$, we finish the proof of the statement. Now, I will begin proofing the above claim. Let P, Q be probability measures on $\{0, 1\}^N$. Define the Wasserstein 1-distance (with Hamming distance) as:

$$W_1(P, Q) = \inf_{\pi \in \Omega(P, Q)} \mathbb{E}_\pi[\|Y - Y'\|_1],$$

where $\|y - y'\|_1 = \sum_{i=1}^N |y_i - y'_i|$ is the Hamming distance. Define the weak transport cost as:

$$\bar{T}_f(Q | P) = \inf_{\pi \in \Omega(P, Q)} \int_{\{0, 1\}^N} C_T \|y - \mathbb{E}_\pi[Y' | Y = y]\|_1^2 P(dy).$$

Define another weak transport cost as:

$$\bar{T}'(Q | P) = \inf_{\pi \in \Omega(P, Q)} \int_{\{0, 1\}^N} C_T^{1/2} \|y - \mathbb{E}_\pi[Y' | Y = y]\|_1 P(dy).$$

By Jensen's inequality and the square root function is concave, we have

$$\begin{aligned} \bar{T}'(Q | P) &= \inf_{\pi \in \Omega(P, Q)} \int_{\{0, 1\}^N} \sqrt{C_T \|y - \mathbb{E}_\pi[Y' | Y = y]\|_1^2} P(dy) \\ &\leq \inf_{\pi \in \Omega(P, Q)} \left[\int_{\{0, 1\}^N} C_T \|y - \mathbb{E}_\pi[Y' | Y = y]\|_1^2 P(dy) \right]^{1/2} \\ &= \sqrt{\bar{T}_f(Q | P)} \end{aligned} \tag{A.26}$$

We are going to show:

$$W_1(P, Q) = \frac{1}{\sqrt{C_T}} \bar{T}'(Q | P). \tag{A.27}$$

Consider each component in $\|y - \mathbb{E}_\pi[Y' | Y = y]\|_1$ separately. Given coupling π and condition-

ing on $Y = y$, we have:

$$\mathbb{E}_\pi[Y'_i|Y = y] = \mathbb{P}_\pi(Y'_i = 1|Y = y) := p_i(y).$$

Thus,

$$|y_i - p_i(y)| = \mathbb{P}_\pi(Y'_i \neq y_i|Y = y).$$

Summing over units, we get:

$$\|y - \mathbb{E}_\pi[Y'|Y = y]\|_1 = \sum_{i=1}^N \mathbb{P}_\pi(Y'_i \neq y_i|Y = y) = \mathbb{E}_\pi[\|y - Y'\|_1|Y = y].$$

Integrating w.r.t. $P(dy)$, we have exact equality:

$$\int_{\{0,1\}^N} \|y - \mathbb{E}_\pi[Y'|Y = y]\|_1 P(dy) = \mathbb{E}_\pi[\|Y - Y'\|_1].$$

Since this holds for any coupling π , we have equality of objective functions for each coupling. Thus, the infimum over all couplings must coincide. This establishes the desired equivalence. Combining Eq.A.26 and Eq.A.27, we finish the proof of our claim.

□

A.1.5 Proof of Proposition 2.3.2

Under Assumptions 4 to 8, the optimization for $\tilde{\mu}$ has a unique maximizer and the iteration procedure of Algorithm 3 converges to this maximizer.

Proof. To prove the existence of a unique maximizer, we need to show that two conditions are satisfied.

1. The objective function is continuous and differentiable in the interior.
2. The boundary point cannot be a global optimum.

Since our objective function is the sum of linear terms, quadratic functions and logarithmic functions, the first condition is trivially satisfied. To check the second condition, we need to verify that the derivative is positive at $\mu_i^Q = 0$ and is negative at $\mu_i^Q = 1$. The derivative is:

$$\begin{aligned} \frac{\partial}{\partial \mu_i^Q} \mathcal{A}(\mu^Q, \mathcal{X}, D, G) &= \theta_0 + \theta_1 d_i + X_i' \theta_2 + X_i' \theta_3 d_i + A_N \sum_{j=1}^N \theta_4 m_{ij} G_{ij} d_j \\ &+ A_N \sum_{j=1}^N m_{ij} G_{ij} (\theta_5 + \theta_6 d_i d_j) \mu_j^Q - \log(\mu_i^Q) + \log(1 - \mu_i^Q). \end{aligned}$$

When $\mu_i^Q = 0$,

$$-\log(\mu_i^Q) + \log(1 - \mu_i^Q) = -\log(0) + \log(1) = +\infty.$$

When $\mu_i^Q = 1$,

$$-\log(\mu_i^Q) + \log(1 - \mu_i^Q) = -\log(1) + \log(0) = -\infty.$$

Since all of the elements in $\theta_0 + \theta_1 d_i + X_i' \theta_2 + X_i' \theta_3 d_i + A_N \sum_{j=1}^N \theta_4 m_{ij} G_{ij} d_j + A_N \sum_{j=1}^N m_{ij} G_{ij} (\theta_5 + \theta_6 d_i d_j) \mu_j^Q$ are bounded,

$$\frac{\partial}{\partial \mu_i^Q} \mathcal{A}(\mu^Q, \mathcal{X}, D, G) \big|_{\mu_i^Q=0} = +\infty,$$

and

$$\frac{\partial}{\partial \mu_i^Q} \mathcal{A}(\mu^Q, \mathcal{X}, D, G) \big|_{\mu_i^Q=1} = -\infty.$$

Away from the boundary, the objective function increases. A global optimum, if it exists, then has to be in the interior. Since the objective function is continuous and differentiable, the first-order condition has to be satisfied at the global optimum. The coordinate ascent property that holds at every iterative step implies that the objective function increases at each step, which further implies that the point that the algorithm converges to cannot be a local maximum. \square

A.1.6 Proof of Contraction Mapping

Under Assumptions 4 to 6, and 9, the iteration procedure in Algorithm 3 is a *contraction mapping* for all $i \in \mathcal{N}$, for all $\{d_i\}_{i=1}^N \in \{0, 1\}^N$, for all $\mathcal{X} \in \mathbb{R}^{N \times k}$, and for all $G \in \{0, 1\}^{N \times N}$.

Proof. This proof is very similar to Lemma A.1.3. Recall the iteration in Algorithm 3:

$$\tilde{\mu}_i^{t+1} = \Lambda \left[\theta_0 + \theta_1 d_i + X'_i(\theta_2 + \theta_3 d_i) + A_N \theta_4 \sum_{j \neq i} m_{ij} G_{ij} d_j + A_N \sum_{j \neq i} m_{ij} G_{ij} (\theta_5 + \theta_6 d_i d_j) \tilde{\mu}_j^t \right].$$

We denote this iteration process as $\{\tilde{\mu}^t\}_{t=1}^T$ and show the above mapping $T : [0, 1]^N \rightarrow [0, 1]^N$ is a contraction mapping. To prove the above iteration is a contraction mapping, we use ℓ_1 -distance. For any $t \geq 1$,

$$d(T(\tilde{\mu}^t), T(\tilde{\mu}^s)) = \sum_{i=1}^N |\tilde{\mu}_i^{t+1} - \tilde{\mu}_i^{s+1}|,$$

$$d(\tilde{\mu}^t, \tilde{\mu}^s) = \sum_{i=1}^N |\tilde{\mu}_i^t - \tilde{\mu}_i^s|.$$

First, we know $\Lambda(\cdot)$ is a sigmoid function. Hence, its largest slope is 0.25.

$$\tilde{\mu}_i^{t+1} - \tilde{\mu}_i^{s+1} \leq 0.25 A_N \sum_{j \neq i} (\theta_5 + \theta_6 d_i d_j) m_{ij} G_{ij} (\tilde{\mu}_j^t - \tilde{\mu}_j^s)$$

(By Multivariate Mean Value Theorem)

$$\leq 0.25 A_N (|\theta_5| + |\theta_6|) \sum_{j \neq i} m_{ij} G_{ij} (\tilde{\mu}_j^t - \tilde{\mu}_j^s)$$

$$\leq 0.25 A_N (|\theta_5| + |\theta_6|) \sum_{j \neq i} m_{ij} G_{ij} |\tilde{\mu}_j^t - \tilde{\mu}_j^s|.$$

Therefore,

$$|\tilde{\mu}_i^{t+1} - \tilde{\mu}_i^{s+1}| \leq \frac{A_N (|\theta_5| + |\theta_6|)}{4} \sum_{j \neq i} m_{ij} G_{ij} |\tilde{\mu}_j^t - \tilde{\mu}_j^s|.$$

Hence,

$$\begin{aligned} \sum_{i=1}^N |\tilde{\mu}_i^{t+1} - \tilde{\mu}_i^{s+1}| &\leq \frac{A_N(|\theta_5| + |\theta_6|)}{4} \sum_{i=1}^N \sum_{j \neq i} m_{ij} G_{ij} |\tilde{\mu}_j^t - \tilde{\mu}_j^s| \\ &\leq \frac{A_N(|\theta_5| + |\theta_6|)}{4} \sum_{i=1}^N |\tilde{\mu}_i^t - \tilde{\mu}_i^s| \max_{i \in \mathcal{N}} \sum_{j \neq i} m_{ij} G_{ij}. \end{aligned}$$

Therefore, under Assumption 9, T is a contraction mapping. In addition, since $\tilde{\mu}^t \in [0, 1]$ for all $t \geq 1$, the metric space $(\tilde{\mu}, d)$ is a complete metric space. By Theorem A.3.1, $\{\tilde{\mu}^t, \tilde{\mu}^1, \dots\}$ will converge to a unique fixed point, which is $\tilde{\mu}$. \square

A.2 Theorem

A.2.1 Proof of Theorem 2.3.1

Let Q^* denote the independent Bernoulli distribution that solves Eq. 2.10. Under Assumptions 4 to 8, the Kullback–Leibler divergence between Q^* and P is bounded from above by:

$$\mathbb{KL}(Q^* || P) \leq C_1 A_N \bar{N} + C_2 N + \mathcal{O}\left(\sqrt{A_N^2 \bar{N}^2 N}\right) + \mathcal{O}\left(\sqrt{A_N^3 \bar{N}^2 N^2}\right) + \mathcal{O}\left(\sqrt{A_N^3 \bar{N} N^3}\right) + o(N),$$

where C_1, C_2 are constants that depend only upon θ , \bar{m} , $\max_{1 \leq i \leq N} |X_i' \theta_2|$, and $\max_{1 \leq i \leq N} |X_i' \theta_3|$.

Proof. We need to apply Theorem A.3.2. We define $f : [0, 1]^N \rightarrow \mathbb{R}$ as:

$$f(\tilde{\mu}) = \sum_i (\theta_0 + \theta_1 d_i + X_i'(\theta_2 + \theta_3 d_i) + A_N \theta_4 \sum_{j \neq i} m_{ij} G_{ij} d_j) \tilde{\mu}_i + \frac{A_N}{2} \sum_i \sum_j (\theta_5 + \theta_6 d_i d_j) m_{ij} G_{ij} \tilde{\mu}_i \tilde{\mu}_j.$$

Therefore,

$$\begin{aligned}
\|f\| &\leq \sum_i (|\theta_0| + |\theta_1| + |X'_i\theta_2| + |\theta_3 X_i|) + A_N \sum_i \sum_j |\theta_4 m_{ij}| G_{ij} + \frac{A_N}{2} \sum_i \sum_j G_{ij} (|\theta_5 m_{ij}| + |\theta_6 m_{ij}|) \\
&\leq N(|\theta_0| + |\theta_1| + \max_i |X'_i\theta_2| + \max_i |X'_i\theta_3|) + \overline{m} A_N (|\theta_4| + |\theta_5| + |\theta_6|) \sum_i \sum_j G_{ij} \\
&\leq N(|\theta_0| + |\theta_1| + \max_i |X'_i\theta_2| + \max_i |X'_i\theta_3|) + \overline{m} A_N N \overline{N} (|\theta_4| + |\theta_5| + |\theta_6|) \\
&=: \tilde{a}.
\end{aligned}$$

The partial derivative of $f(\tilde{\mu})$ with respect to $\tilde{\mu}_i$ is:

$$\frac{\partial f(\tilde{\mu})}{\partial \tilde{\mu}_i} = \theta_0 + \theta_1 + X'_i\theta_2 + X'_i\theta_3 d_i + A_N \sum_{j \neq i} \theta_4 m_{ij} G_{ij} d_j + A_N \sum_{j \neq i} \theta_5 m_{ij} G_{ij} \tilde{\mu}_j + A_N \sum_{j \neq i} \theta_6 m_{ij} G_{ij} d_i d_j \tilde{\mu}_j. \quad (\text{A.28})$$

Therefore,

$$\|\nabla_i f(\tilde{\mu})\| \leq |\theta_0| + |\theta_1| + \max_i |X'_i\theta_2| + \max_i |X'_i\theta_3| + \overline{m} A_N \overline{N} (|\theta_4| + |\theta_5| + |\theta_6|) =: \tilde{b}.$$

The second partial derivative with respect to $\tilde{\mu}_j$ is:

$$\frac{\partial^2 f(\tilde{\mu})}{\partial \tilde{\mu}_i \partial \tilde{\mu}_j} = A_N \theta_5 m_{ij} G_{ij} + A_N \theta_6 m_{ij} G_{ij} d_i d_j.$$

Therefore, for all $j \neq i$,

$$\|\nabla_i \nabla_j g(\tilde{\mu})\| \leq \overline{m} A_N (|\theta_5| + |\theta_6|) G_{ij} = \tilde{c} G_{ij} =: \tilde{c}_{ij},$$

with the second derivative zero if $i = j$. Next, we need to compute $|\mathcal{M}(\epsilon)|$, where $\mathcal{M}(\epsilon)$ is the finite subset of \mathbb{R}^N such that for any $\tilde{\mu} \in \{0, 1\}^N$, there exists $\eta = (\eta_1, \dots, \eta_N) \in \mathcal{M}(\epsilon)$ such that

$$\sum_i \left(\frac{\partial f(\tilde{\mu})}{\partial \tilde{\mu}_i} - \eta_i \right)^2 \leq N \epsilon^2.$$

We note that this is requirement on the derivative at the boundary; we only need the above condition to hold when $\tilde{\mu} \in \{0, 1\}^N$ to define $\mathcal{M}(\epsilon)$, and not more generally for $\tilde{\mu} \in (0, 1)^N$.

Recalling Eq. A.28 and defining T_1 and T_2 as:

$$T_1(\tilde{\mu}) := \frac{A_N \theta_5}{2} \sum_i \sum_j m_{ij} G_{ij} \tilde{\mu}_i \tilde{\mu}_j, \quad T_2(\tilde{\mu}) := \frac{A_N \theta_6}{2} \sum_i \sum_j m_{ij} G_{ij} d_i d_j \tilde{\mu}_i \tilde{\mu}_j,$$

we state the partial derivative of $f(\tilde{\mu})$ as

$$\frac{\partial f(\tilde{\mu})}{\partial \tilde{\mu}_i} = \theta_0 + \theta_1 + X'_i \theta_2 + X'_i \theta_3 d_i + \theta_4 A_N \sum_j m_{ij} G_{ij} d_j + \frac{\partial T_1(\tilde{\mu})}{\partial \tilde{\mu}_i} + \frac{\partial T_2(\tilde{\mu})}{\partial \tilde{\mu}_i}.$$

Let $\mathcal{M}_1(\epsilon)$ be the finite subset of \mathbb{R}^N such that for any $\tilde{\mu} \in \{0, 1\}^N$, there exists $\lambda = (\lambda_1, \dots, \lambda_N) \in \mathcal{M}_1(\epsilon)$ such that

$$\sum_i \left(\frac{\partial T_1(\tilde{\mu})}{\partial \tilde{\mu}_i} - \lambda_i \right)^2 \leq N \epsilon^2.$$

Let $\mathcal{M}_2(\epsilon)$ be the finite subset of \mathbb{R}^N such that for any $\tilde{\mu} \in \{0, 1\}^N$, there exists $\vartheta = (\vartheta_1, \dots, \vartheta_N) \in \mathcal{M}_2(\epsilon)$ such that

$$\sum_i \left(\frac{\partial T_2(\tilde{\mu})}{\partial \tilde{\mu}_i} - \vartheta_i \right)^2 \leq N \epsilon^2.$$

Defining λ_i and ϑ_i for all $i \in \mathcal{N}$ to be:

$$\lambda_i = A_N \theta_5 \sum_j m_{ij} G_{ij} \tilde{y}_j, \quad \vartheta_i = A_N \theta_6 \sum_j m_{ij} G_{ij} d_i d_j \tilde{v}_j,$$

where $\tilde{y}_j \in \{0, 1\}$ and $\tilde{v}_j \in \{0, 1\}$ by Lemma A.1.2, we can then define $\mathcal{M}(\epsilon)$ as:

$$\begin{aligned} \mathcal{M}(\epsilon) := & \left\{ \theta_0 + \theta_1 + X'_i \theta_2 + X'_i \theta_3 d_i + A_N \theta_4 \sum_j m_{ij} G_{ij} d_j + \ell_1 + \ell_2 \right. \\ & \left. : \ell_1 \in \mathcal{M}_1 \left(\frac{\epsilon}{\sqrt{2}} \right), \ell_2 \in \mathcal{M}_2 \left(\frac{\epsilon}{\sqrt{2}} \right), i \in \mathcal{N} \right\}. \end{aligned}$$

Then,

$$\begin{aligned}
\left(\frac{\partial T_1(\tilde{\mu})}{\partial \tilde{\mu}_i} - \lambda_i\right)^2 &= \left(A_N \theta_5 \sum_j m_{ij} G_{ij} (\tilde{\mu}_j - \tilde{y}_j)\right)^2 \\
&\leq \left(A_N \theta_5 \sum_j m_{ij}^2 G_{ij}\right) \left(A_N \theta_5 \sum_j (\tilde{\mu}_j - \tilde{y}_j)^2\right) \\
&\quad \text{(By the Cauchy–Schwarz inequality)} \\
&\leq A_N^2 \theta_5^2 N \sum_j (\tilde{\mu}_j - \tilde{y}_j)^2 \max_{i,j} m_{ij}^2 \\
&\leq A_N^2 \theta_5^2 N \tau_1^2 \max_{i,j} m_{ij}^2 \\
&\quad \text{(By Lemma A.1.2)} \\
&= \frac{\epsilon^2}{2} \quad \text{By choosing } \tau_1 = \sqrt{\frac{\epsilon^2}{2A_N^2 \theta_5^2 N \max_{i,j} m_{ij}^2}}.
\end{aligned}$$

$$\begin{aligned}
\left(\frac{\partial T_2(\tilde{\mu})}{\partial \tilde{\mu}_i} - \vartheta_i\right)^2 &= \left(A_N \theta_6 \sum_j m_{ij} G_{ij} d_i d_j (\tilde{\mu}_j - \tilde{v}_j)\right)^2 \\
&\leq \left(A_N \theta_6 \sum_j m_{ij}^2 G_{ij} d_i d_j\right) \left(A_N \theta_6 \sum_j (\tilde{\mu}_j - \tilde{v}_j)^2\right) \\
&\quad \text{(By the Cauchy–Schwarz inequality)} \\
&\leq A_N^2 \theta_6^2 N \sum_j (\tilde{\mu}_j - \tilde{v}_j)^2 \max_{i,j} m_{ij}^2 \\
&\leq A_N^2 \theta_6^2 N \tau_2^2 \max_{i,j} m_{ij}^2 \\
&\quad \text{(By Lemma A.1.2)} \\
&= \frac{\epsilon^2}{2} \quad \text{By choosing } \tau_2 = \sqrt{\frac{\epsilon^2}{2A_N^2 \theta_6^2 N \max_{i,j} m_{ij}^2}}.
\end{aligned}$$

Therefore,

$$|\mathcal{M}(\epsilon)| \leq N \cdot \frac{N(N+1)}{2} \cdot \left| \mathcal{M}_1 \left(\frac{\epsilon}{\sqrt{2}} \right) \right| \cdot \left| \mathcal{M}_2 \left(\frac{\epsilon}{\sqrt{2}} \right) \right| \leq 2^{2N-1} N^2 (N+1).$$

We now apply Theorem A.3.2, choosing $\epsilon = N^{-1}$:

$$\begin{aligned}
\mathbb{K}(Q^* \| P) &\leq \frac{1}{4} \left(N \sum_i b_i^2 \right)^{\frac{1}{2}} N^{-1} + 3 + \log(2^{2N-1}(N^3 + N)) \\
&\quad + 4 \left(\sum_i b_i^2 + \frac{1}{4} \sum_{i,j} (ac_{ij}^2 + b_i b_j c_{ij} + 4b_i c_{ij}) \right)^{\frac{1}{2}} + \log 2 \\
&\leq \frac{\tilde{b}}{4} + 3 + 2N \log 2 + \log(N^3 + N) + 4 \left(\tilde{b}^2 N + \frac{1}{4} \sum_{i,j} (\tilde{a} \tilde{c}^2 G_{ij} + \tilde{b}^2 \tilde{c} G_{ij} + 4\tilde{b} \tilde{c} G_{ij}) \right)^{\frac{1}{2}} \\
&\leq \frac{\tilde{b}}{4} + 3 + 2N \log 2 + \log(N^3 + N) + 4 \left(\tilde{b}^2 N + \frac{1}{4} \bar{N} N (\tilde{a} \tilde{c}^2 + \tilde{b}^2 \tilde{c} + 4\tilde{b} \tilde{c}) \right)^{\frac{1}{2}} \\
&= C_1 A_N \bar{N} + C_2 N \\
&\quad + \sqrt{(C_3 A_N \bar{N} + C_4 A_N^2 \bar{N}^2) N + (C_5 A_N + C_6 A_N^2 \bar{N} + C_7 A_N^3 \bar{N}^2 + C_8 A_N^2 + C_9 \bar{N} A_N^3) N^2 + o(N)} \\
&= C_1 A_N \bar{N} + C_2 N + \mathcal{O} \left(\sqrt{A_N^2 \bar{N}^2 N} \right) + \mathcal{O} \left(\sqrt{A_N^3 \bar{N}^2 N^2} \right) + \mathcal{O} \left(\sqrt{A_N^3 \bar{N} N^2} \right) + o(N),
\end{aligned}$$

where $o(N)$ collects those elements that are constant or that grow at a slower rate than N , and

$$C_1 = \frac{1}{4} \bar{m} (|\theta_4| + |\theta_5| + |\theta_6|),$$

$$C_2 = 2 \log 2,$$

$$C_3 = 32 \bar{m} (|\theta_0| + |\theta_1| + \max_i |X'_i \theta_2| + \max_i |X'_i \theta_3|) (|\theta_4| + |\theta_5| + |\theta_6|)$$

$$C_4 = 16 \bar{m}^2 (|\theta_4| + |\theta_5| + |\theta_6|)^2$$

$$C_5 = 4 \bar{m} (|\theta_0| + |\theta_1| + \max_i |X'_i \theta_2| + \max_i |X'_i \theta_3| + 4) (|\theta_0| + |\theta_1| + \max_i |X'_i \theta_2| + \max_i |X'_i \theta_3|) (|\theta_5| + |\theta_6|).$$

$$C_6 = 8 \bar{m}^2 [(|\theta_0| + |\theta_1| + \max_i |X'_i \theta_2| + \max_i |X'_i \theta_3|) + 2] (|\theta_4| + |\theta_5| + |\theta_6|) (|\theta_5| + |\theta_6|),$$

$$C_7 = 4 \bar{m}^3 (|\theta_4| + |\theta_5| + |\theta_6|)^2 (|\theta_5| + |\theta_6|),$$

$$C_8 = 4 \bar{m}^2 (|\theta_0| + |\theta_1| + \max_i |X'_i \theta_2| + \max_i |X'_i \theta_3|) (|\theta_5| + |\theta_6|)^2,$$

$$C_9 = 4 \bar{m}^3 (|\theta_4| + |\theta_5| + |\theta_6|) (|\theta_5| + |\theta_6|)^2.$$

□

A.2.2 Proof of Theorem 2.3.2

Under Assumptions 4 to 11, the curvature ξ and submodularity ratio γ of $\tilde{W}(\mathcal{D})$ are in $(0, 1)$. The greedy Algorithm enjoys an approximation guarantee for solving the problem in Eq. 2.13 of:

$$\tilde{W}(D_G) \geq \frac{1}{\xi}(1 - e^{-\xi\gamma})\tilde{W}(\tilde{D}),$$

where D_G is the treatment assignment rule obtained by Algorithm 4.

For completeness, we present the statement of Theorem 1 in [Bian et al. \(2017\)](#).

Theorem A.2.1. *Theorem 1 ([Bian et al., 2017](#)) Let $F(\cdot)$ be a non-negative nondecreasing set function with submodularity ratio $\gamma \in [0, 1]$ and curvature $\xi \in [0, 1]$. The greedy algorithm enjoys the following approximation guarantee for solving the maximization problem with cardinality constraint:*

$$F(D_G) \geq \frac{1}{\xi}(1 - e^{-\xi\gamma})F(D^*),$$

where D_G is the result of the greedy algorithm and D^* is the optimal solution.

Proof. The first statement in Theorem 2.3.2 directly follows Lemma A.1.4. Then we apply Theorem A.2.1 to get the second statement. □

A.3 Results from Previous Literature

Lemma A.3.1. ([Götze et al., 2019](#), §Theorem 4.2) *Let δ be a measure with full support on \mathcal{Y}^N . Define*

$$\chi = \min_{1 \leq i \leq N} \min_{Y \in \mathcal{Y}^N} \delta_i(Y_i | Y_{-i}).$$

Let $A = (a_{ij})_{i,j \leq N}$ satisfy $a_{ii} = 0$ for all i and for $i \neq j$

$$\|\delta_i(\cdot \mid Y_{-i}) - \delta_i(\cdot \mid Y'_{-i})\|_{\text{TV}} \leq a_{ij},$$

whenever $Y, Y' \in \mathcal{Y}^N$ differ only at the j -th coordinate. Assume moreover that

$$\|A\|_{2,2} < 1.$$

Then δ satisfies the approximate tensorization property $\text{AT}(C)$ with

$$C = \frac{1}{\chi (1 - \|A\|_{2,2})^2}.$$

Lemma A.3.2. ([Adamczak et al. \(2019, §Proposition 5.4\)](#)) If X is a $[-1, 1]^N$ - valued random vector with law μ , which satisfies the approximate tensorization $\text{AT}(C)$, then X satisfies the dimension-free convex concentration inequality with constant K depending only on C .

Lemma A.3.3. ([Adamczak et al., 2019, §Theorem 5.3](#)) Let X be a random vector in \mathbb{R}^N with distribution μ . The following conditions are equivalent:

- (i) There exists K such that X has the dimension-free convex concentration property with constant K .
- (ii) There exists c such that μ satisfies the inequality $\overline{\text{T}}_f$ with $f(x) = c\|x\|_1^2$.
- (iii) There exist $D, \lambda > 0$ such that for every convex Lipschitz function and every concave function whose Hessian is bounded below by $(-\lambda)\text{Id}$,

$$\text{Ent}(e^{f(X)}) \leq D \mathbb{E}[\|\nabla f(X)\|^2] \mathbb{E}[e^{f(X)}].$$

Moreover, for any two of the above assertions, the constants in one may be chosen to depend only on the constants in the other.

Theorem A.3.1. Banach Fixed-Point Theorem: Let (X, d) be a non-empty complete metric space with a contraction mapping $T : X \rightarrow X$. Then T admits a unique fixed-point x^* .

Theorem A.3.2. Chatterjee and Dembo (2016) Suppose that $f : [0, 1]^N \rightarrow \mathbb{R}$ is twice continuously differentiable in $(0, 1)^N$, so that f and all of its first- and second-order derivatives extend continuously to the boundary. Let $\|f\|$ denote the supremum norm of $f : [0, 1]^N \rightarrow \mathbb{R}$. For each i and j , denote

$$f_i := \frac{\partial f}{\partial x_i}, \quad f_{ij} := \frac{\partial^2 f}{\partial x_i \partial x_j},$$

and let

$$a := \|f\|, \quad b_i := \|f_i\|, \quad c_{ij} := \|f_{ij}\|.$$

Given $\epsilon > 0$, $\mathcal{M}(\epsilon)$ is a finite subset of \mathbb{R}^N such that for any $\tilde{\mu} \in \{0, 1\}^N$, there exists $\eta = (\eta_1, \dots, \eta_N) \in \mathcal{M}(\epsilon)$ such that

$$\sum_i \left(\frac{\partial f(\tilde{\mu})}{\partial \tilde{\mu}_i} - \eta_i \right)^2 \leq N\epsilon^2.$$

Let us define for any $\tilde{\mu} = (\tilde{\mu}_1, \dots, \tilde{\mu}_N) \in [0, 1]^N$,

$$I(\tilde{\mu}) = \sum_{i=1}^N [\tilde{\mu}_i \log \tilde{\mu}_i + (1 - \tilde{\mu}_i) \log(1 - \tilde{\mu}_i)].$$

Let us define

$$F = \sup_{\tilde{\mu} \in [0, 1]^N} (f(\tilde{\mu}) - I(\tilde{\mu})) + \text{lower order terms}.$$

Then for an $\epsilon > 0$,

$$F \leq \sup_{\tilde{\mu} \in [0, 1]^N} (f(\tilde{\mu}) - I(\tilde{\mu})) + \text{complexity term} + \text{smoothness term}.$$

where

$$\text{complexity term} = \frac{1}{4} \left(n \sum_{i=1}^N b_i^2 \right)^{1/2} \epsilon + 3N\epsilon + \log |\mathcal{M}(\epsilon)|,$$

and

$$\begin{aligned} \text{smoothness term} = & 4 \left(\sum_{i=1}^N (ac_{ii} + b_i^2) + \frac{1}{4} \sum_{i,j=1}^N (ac_{ij}^2 + b_i b_j c_{ij} + 4b_i c_{ij}) \right)^{1/2} \\ & + \frac{1}{4} \left(\sum_{i=1}^N b_i^2 \right)^{1/2} \left(\sum_{i=1}^N c_{ii}^2 \right)^{1/2} + 3 \sum_{i=1}^N c_{ii} + \log 2. \end{aligned}$$

Chapter 3

Robust Network Targeting with Multiple Nash Equilibria

3.1 Introduction

Many policy problems involve allocating treatment among a network of interacting agents. Examples include technology diffusion ([Parente and Prescott, 1994](#); [Alvarez et al., 2023](#)), teenage smoking ([Nakajima, 2007](#)), consumer adoption decisions ([Banerjee et al., 2013](#); [Keane and Wasi, 2013](#)), and education and migration ([Hsiao, 2022](#)). Research in these fields highlights the role of spillover effects, particularly those arising from strategic interactions.

Among other things, these strategic interactions lead to the presence of multiple Nash equilibria, which complicates the process of finding an optimal treatment allocation policy. To handle this multiplicity, counterfactual policy analysis “*has made simplifying assumptions which either change the outcome space or impose ad hoc selection mechanisms in regions of multiplicity*” ([Tamer, 2003](#)). Consequently, this approach “*potentially introduces misspecifications and nonrobustness in the analysis of substantive questions*” ([De Paula, 2013](#)). To address this problem, can we develop a treatment allocation rule that remains optimal even under the least favorable equilibrium?

Focusing on a class of network models where units participate in a simultaneous decision game with strategic complementarity (Brock and Durlauf, 2001; Ballester et al., 2006; Molinari and Rosen, 2008; Jia, 2008; Echenique and Komunjer, 2009; Lazzati, 2015; Graham and Pelican, 2023), this paper develops a method for constructing a maximin optimal treatment allocation rule that is robust to the presence of multiple Nash equilibria. A planner allocates a binary treatment among a target population of N units embedded within a network, where each unit’s covariates and the network structure are observable. Each unit then simultaneously chooses a binary action to maximize its own utility, which depends on its own characteristics and treatment, as well as the characteristics, treatments, and expected choice of its neighbors¹. Our goal is to learn a treatment allocation policy that maximizes social welfare for the target population.

To determine the optimal treatment allocation rule for the target population, we assume that there exists data for a social network of units who have been assigned treatment in the past. This sample may differ from the target population in terms of both the number of units and the network structure. Data is assumed to be available for each unit’s covariates, decisions, and assignments, as well as those of their neighbors. After assessing how individual outcomes vary in response to different treatment allocations among the training sample, we analyze the optimal treatment allocation strategy, taking into account the covariates and network structure of the target population. Consider, for example, targeted information provision in villages with the aim of increasing microfinance adoption, as discussed in Banerjee et al. (2013). By analyzing heterogeneous choices among the units in villages selected by policymakers, we then estimate whom to better target in external villages.

There are both theoretical and practical challenges to studying optimal treatment allocation in the presence of strategic interactions. The primary theoretical challenge is incompleteness (Jovanovic, 1989) of the model when there are multiple Nash equilibria². With-

¹This incomplete information setting (Brock and Durlauf, 2001; Bajari et al., 2010a; de Paula and Tang, 2012), is our primary focus. Section 3.7 extends our results to the complete information setting.

²See the detailed surveys by De Paula (2013); Molinari (2020); Kline and Tamer (2020).

out assuming an equilibrium selection mechanism, our model predicts a set of equilibrium outcomes under a counterfactual policy. From a theoretical perspective, one cannot judge which equilibrium outcome is more likely than the others. This paper allows for these multiple equilibria and imposes no assumptions on equilibrium selection. Instead, we provide set-identified equilibrium social welfare for any treatment allocation policy, along with a closed-form expression that characterizes the bounds of this set.

As the counterfactual equilibrium social welfare is only set-identified, we cannot directly target equilibrium welfare when designing a treatment allocation rule. To address this uncertainty, we refine the optimality of treatment allocation using the maximin welfare criterion. This criterion is employed in the robust decision theory literature (e.g., [Chamberlain, 2000a](#)), and the robust mechanism design literature (e.g., [Morris et al., 2024](#)). Under the maximin welfare criterion, our objective is to design a treatment allocation policy that maximizes social welfare evaluated under the least favourable equilibrium selection rule.

In terms of implementation, there are two challenges. We adopt a parametric utility function specification, and the first challenge is estimating the parameters of this utility function. We assume the existence of a one-period training data set that contains a finite number n of units, along with their covariates and the network structure³. An existing treatment allocation policy is assumed, and we observe each unit’s choice under this policy. We estimate parameters using the two-step maximum likelihood estimator proposed by [Leung \(2015\)](#). However, in the context of a network game setting, the asymptotic behavior of this estimator cannot be characterized without assuming how the network structure changes as the number of units increases (i.e., whether the network is dense or sparse). Although non-asymptotic results could elucidate how the sampling uncertainty of this estimator is influenced by network structure, the current literature lacks such analysis. Addressing this gap is one of the primary focuses of our paper.

The second challenge to implementation is finding the maximin optimal treatment alloca-

³We allow the training data to come from our target population, with caveats about the private information of each unit. A more detailed discussion is provided in Section [3.4.1](#).

tion, which requires optimizing an objective function dependent on a system of simultaneous equations. In the presence of strategic interactions, when a treatment is assigned to a unit, it not only influences their behavior but also that of their neighbors. This, in turn, affects the payoff of their neighbors' neighbors, propagating feedback effects throughout the network and presenting a complex combinatorial optimization problem. To tackle this complexity, we propose a greedy algorithm. This algorithm sequentially assigns treatment to the agent who yields the highest marginal welfare gain at each step. However, this class of algorithm generally lacks a performance guarantee. We address this by characterizing the performance guarantee through the features of our objective function.

We evaluate the performance of our proposed method based on its regret, which is defined as the difference between the largest achievable welfare and the welfare achieved by our proposed method, evaluated under the least favourable equilibrium selection rule. Regret arises from two sources of uncertainty: The first is due to the use of estimated structural parameters, and reflects sampling uncertainty. The second is due to the use of a greedy algorithm.

This paper makes three theoretical contributions: (i) It provides a closed-form expression for the identified set corresponding to the equilibrium outcomes under any arbitrary policy intervention. The heavy computation costs due to the large number of equilibria have limited the range of empirical applications in the literature to static models with a small number of players and choice alternatives. Our approach avoids computing the set of equilibria and hence allows for a feasible characterization of the identified region for the equilibrium social welfare; (ii) It presents the first non-asymptotic result on regret with strategic interactions. It shows that, under regularity conditions, the regret introduced by sampling uncertainty shrinks at the rate $\log(n)/\sqrt{n}$; (iii) It offers a theoretical performance guarantee for the regret associated with using a greedy algorithm to solve optimization problems involving systems of simultaneous equations, a topic previously unexplored in the existing literature.

To demonstrate how our method can be implemented and quantitatively evaluate its per-

formance, we apply it to the data of [Banerjee et al. \(2013\)](#). We design a policy to maximize the take-up rate of microfinance products among households across various villages. For each village in the sample, we estimate the utility function parameters. These estimates are then used to assess the presence of strategic complementarity in each village. We find that strategic complementarities are present in 16 out of the 43 villages. For these villages, we construct an individualized treatment allocation rule using our greedy algorithm. Empirical results revealed that the occurrence of multiple equilibria can vary depending on the allocation method used. We compare the welfare outcomes achieved by our algorithm with those obtained by the NGO Bharatha Swamukti Samsthe (BSS). Our results indicate that, for all 16 villages exhibiting strategic complementarity, our method achieves notably higher welfare levels, with improvements ranging from 20% to 270%, and an average improvement of 116%. Additionally, the lower bound of welfare under our method consistently exceeds the maximal welfare attained under the allocation rule used by BSS and a rule that assigns treatment at random. These substantial welfare gains highlight the benefits of individualized targeting in the presence of strategic interference, which demonstrates the efficacy of our approach in optimizing resource allocation and improving social welfare.

3.1.1 Literature Review

This paper is related to several literatures in economics and econometrics, including strategic interactions, statistical treatment rules, robust decision theory, robust mechanism design, and greedy algorithms.

Pioneering contributions to the econometric aspects of game-theoretic models include works by [Jovanovic \(1989\)](#) and [Bresnahan and Reiss \(1991\)](#), which explore the empirical challenges associated with models that feature multiple equilibria. The recent literature on the econometrics of strategic interactions includes simultaneous decision games with complete information, such as [Tamer \(2003\)](#), [Bajari et al. \(2010a,b\)](#), [De Paula et al. \(2018\)](#), [Sheng \(2020\)](#), and [Chesher and Rosen \(2020\)](#); simultaneous decision games with incom-

plete information, such as [de Paula and Tang \(2012, 2020\)](#), [Menzel \(2016\)](#), and [Ridder and Sheng \(2020\)](#); and sequential decision games: [Aguirregabiria and Mira \(2007, 2019\)](#), [Mele \(2017\)](#), [Leung \(2019\)](#), and [Christakis et al. \(2020\)](#).

Focusing on a game with complete information, [Tamer \(2003\)](#) obtains bounds for structural parameters while remaining fully agnostic about the equilibrium selection mechanism. Motivated by this, [Sheng \(2020\)](#) uses a sub-network approach to provide bounds for structural parameters in a network formation setting. [Chesher and Rosen \(2020\)](#) partially identifies structural parameters using the Generalized Instrumental Variable approach ([Chesher and Rosen, 2017](#)). [Bajari et al. \(2010a\)](#) point identifies the parameters for a game with incomplete information, along with providing a semi-parametric estimator. However, estimation methods typically require observing repeated samples of the game, which may not be feasible for social network games. Motivated by this, [Leung \(2015\)](#) studies a two-step maximum likelihood estimator in a large network setting, while [Ridder and Sheng \(2020\)](#) studies a two-step GMM estimator in a similar setting. The paper adopts an existing estimator for structural parameters and treats this as an intermediate step in estimating an optimal policy.

One important task in obtaining an optimal policy is predicting the equilibrium outcome under counterfactual policies. In the strategic interaction literature, counterfactual analysis has been studied among others by [Jia \(2008\)](#), [Aguirregabiria and Mira \(2010\)](#), and [Canen and Song \(2020\)](#) under various assumptions on the equilibrium selection mechanism. [Ciliberto and Tamer \(2009\)](#) does not restrict the equilibrium selection rule but considers only some candidate counterfactual policies. Additionally, [Lee and Pakes \(2009\)](#) investigates ATM network games by enumerating all Nash equilibria and analyzing how different learning algorithms select among them. None of these studies considers the aggregate equilibrium outcome, which aggregates the equilibrium outcomes of each unit, under a counterfactual policy. Here, we consider a social planner, and hence it is crucial to evaluate the aggregate social welfare. Remaining fully agnostic about the equilibrium selection mechanism, we provide a counterfactual analysis of the aggregate social welfare.

Strategic interactions are closely related to social interaction models. These were introduced by [Manski \(1993\)](#), which examines spillover effects through strategic interactions using a linear social interaction model with unique equilibrium. [Brock and Durlauf \(2001\)](#) extends this model to a nonlinear setting and considers multiple equilibria. [Goldsmith-Pinkham and Imbens \(2013\)](#) considers the endogeneity of the network formation process. [De Paula et al. \(2024\)](#) recovers unknown network structure using a linear social interaction model.

This paper contributes to the growing literature on statistical treatment rules, which were introduced into econometrics by [Manski \(2004\)](#) and [Dehejia \(2005\)](#). The recent literature includes [Stoye \(2009, 2012\)](#), [Hirano and Porter \(2009, 2020\)](#), [Chamberlain \(2011\)](#), [Kitagawa and Tetenov \(2018\)](#), [Ananth \(2020b\)](#), [Athey and Wager \(2021\)](#), [Mbakop and Tabord-Meehan \(2021\)](#), [Kitagawa et al. \(2021\)](#), [Sun \(2021\)](#), [Munro et al. \(2021\)](#), [Christensen et al. \(2022\)](#), [Adjaho and Christensen \(2022\)](#), [Kitagawa et al. \(2022b\)](#), [Kitagawa and Wang \(2023a,b\)](#), [Viviano \(2024\)](#), [Fernandez et al. \(2024\)](#), and [Munro \(2024\)](#). In contrast to the i.i.d. setting considered in most of these papers, we consider a setting where the spillover effects of treatment assignment are important. A small number of papers in the literature considers spillover effects. These include [Viviano \(2024\)](#), [Ananth \(2020b\)](#), [Munro et al. \(2021\)](#), and [Kitagawa and Wang \(2023a,b\)](#). Apart from [Kitagawa and Wang \(2023a\)](#), none of those papers consider spillover effects introduced by strategic interaction or the related complications, such as multiple equilibria.

[Viviano \(2024\)](#) and [Ananth \(2020b\)](#) focus on estimating direct and indirect treatment effects to derive optimal allocation policies based on data. We focus on the strategic interaction setting where each unit's behavior is influenced by the behaviors of nearby individuals within a network. These interactions are naturally modeled using game theory ([Jackson and Zenou, 2015](#)), which we adopt in our analysis. In addition, using a game theoretical approach enables us to evaluate social welfare directly through the individual's utility and account for the general equilibrium effects. [Munro et al. \(2021\)](#) focuses on a competitive equilibrium where spillover effects are mediated through the equilibrium price. Their approach uses the mean-

field limit to characterize the asymptotic behavior of treatment effects, focusing on settings where a unique mean-field equilibrium price exists. [Kitagawa and Wang \(2023a\)](#) focuses on a sequential decision game with a Markovian structure, which leads to a unique stationary joint distribution of units' decisions ([Mele, 2017](#)). This stationary distribution allows each state to be revisited instead of converging to multiple distinct equilibria. As a consequence, their stationary welfare differs from the equilibrium welfare that we consider here and we do not require the existence of a potential function. [Kitagawa and Wang \(2023b\)](#) considers the spillover effects from vaccination.

Although the source of uncertainty is different, our paper robustly addresses the incompleteness introduced by multiple equilibria in a manner inspired by the robust decision theory literature (see the recent survey by [Chamberlain, 2020](#)). [Chamberlain \(2000a,b\)](#) consider decision-making when there is uncertainty due to a partially specified subjective distribution. Their robust decision rule maximizes the risk function evaluated at the least-favourable distribution. [Hansen and Sargent \(2001, 2008\)](#) achieve robustness by working within a neighborhood of a reference model and maximizing the minimum of expected utility over that neighborhood. [Manski \(2003\)](#) faces a similar problem to us where some part of the model is missing from the data, and obtains a robust identification region by incorporating the maximum and minimum value of the unobserved component. [Giacomini and Kitagawa \(2021\)](#) applies the robust Bayes approach of [Berger \(1994\)](#) to a set-identified model, and shows asymptotic equivalence between the identified set and the set of posterior means obtained from using a multiple priors. See also [Giacomini et al. \(2021\)](#) and references therein. [Christensen and Connault \(2023\)](#) relaxes parametric assumptions about the distribution of latent variables in a structural model. Their robust counterfactual set is obtained by maximizing (minimizing) the counterfactual through the distribution of latent variables over a neighborhood of the prespecified parametric distribution.

Finally, this paper is closely related to network games and mechanism design, as exemplified by [Morris \(2000\)](#), [Ballester et al. \(2006\)](#), [Galeotti et al. \(2010\)](#), and [Galeotti et al. \(2020\)](#)

for network games, and [Mathevet \(2010\)](#), [Gonçalves and Furtado \(2020\)](#), [Fu et al. \(2021\)](#), [Morris et al. \(2024\)](#), and [Brooks and Du \(2024\)](#) for mechanism design. Network games explore how network characteristics influence behavior. [Jackson et al. \(2008\)](#) and [Jackson and Zenou \(2015\)](#) provide comprehensive summaries. [Galeotti et al. \(2020\)](#) employs a principal component approach to analyze how interventions that change characteristics impact outcomes and develops strategies for optimal interventions within network games. However, they assume a unique equilibrium, whereas this paper focuses on models with multiple equilibria. Following the same setting, [Sun et al. \(2023\)](#) examines optimal interventions that alter network structure, while [Kor and Zhou \(2022\)](#) considers interventions that affect both characteristics and network structure. Nonetheless, these studies differ from ours in terms of utility specification, objective function, and the definition of the action space.

This paper can be viewed as a specific instance of mechanism design, where treatments are allocated to incentivize units' equilibrium behavior towards achieving desired objectives. Closely related is [Morris et al. \(2024\)](#), which characterizes the set of outcomes achievable from the smallest equilibrium, referred to as the smallest implementable outcome, in a supermodular game. Moreover, within a convex potential game, they show that the optimal outcome—realized by implementing information to maximize the smallest equilibrium—results in all players selecting the same action. While our implementation approach differs, the lower bound of the set-identified social welfare in this paper is similar in concept to these smallest implementable outcomes. However, it is more complex to characterize this set as the number of players increases. Additionally, [Morris et al. \(2024\)](#) leaves open the question of which implementation strategies are needed to achieve these outcomes, a gap this paper addresses.

Outline. The rest of this paper proceeds as follows: Section [3.2](#) introduces the game setting and the solution concept. Section [3.3](#) discusses counterfactual analysis. Section [3.4](#) focuses on treatment allocation and implementation. Section [3.5](#) presents theoretical results related to the implementation of our proposed method. We apply our proposed method to the Indian micro- finance data, which is studied by [Banerjee et al. \(2013\)](#), and demonstrate

its performance in Section 3.6. Section 3.7 extends our analysis to the complete information setting. Section 3.8 concludes. All proofs and derivations are shown in Appendix A.2 to Appendix A.5.

3.2 Model

3.2.1 Setup

Let $\mathcal{N} = \{1, 2, \dots, N\}$ be the target population. Each unit i has a K -dimensional vector of characteristics X_i observable to the researcher. X_i is assumed to have bounded support, and we standardize the measurements of X_i to be nonnegative, such that $X_i \in \mathcal{X} \in \mathbb{R}_+^K$. Let $X = [X_1^\top, \dots, X_N^\top] \in \mathcal{X}^N$ be an $N \times K$ matrix whose i th row contains the characteristics of unit i , and let \mathcal{X}^N represent the set of all such possible matrices \mathcal{X} . Let $D = \{D_1, \dots, D_N\} \in \mathcal{D} = \{0, 1\}^N$ be a vector of binary treatment allocations. For $i \in \mathcal{N}$, $D_i = 1$ if unit i is treated and $D_i = 0$ if not.

The social network is represented by an $N \times N$ binary adjacency matrix, denoted by $G = \{G_{ij}\}_{i,j \in \mathcal{N}} \in \mathcal{G} = \{0, 1\}^{N \times N}$. G is assumed to be fixed and exogenous, irrelevant to treatment allocation. $G_{ij} = 1$ indicates that units i and j are connected, while $G_{ij} = 0$ indicates that they are not. Let $\mathcal{N}_i := \{j : G_{ij} \neq 0\}$ denote the set of neighbors of unit i . \bar{N} denotes the maximum number of edges connected to any unit in the network (i.e., $\bar{N} = \max_i |\mathcal{N}_i|$), while \underline{N} denotes the minimum (i.e., $\underline{N} = \min_i |\mathcal{N}_i|$). We adopt the convention of no self-links (i.e., $G_{ii} = 0$ for all $i \in \mathcal{N}$). This framework can accommodate both directed networks, where G_{ij} and G_{ji} can differ, and undirected networks, where $G_{ij} = G_{ji}$ for all $i, j \in \mathcal{N}$. Additionally, we allow the strength of spillover effects to depend not only on the adjacency matrix G_{ij} but also on the covariates and treatment statuses of units i and j .

We consider a counterfactual equilibrium social welfare in the context of a large simultaneous decision game. We use the following notation for our simultaneous decision game.

$Y_i \in \mathcal{Y} = \{0, 1\}$ denotes unit i 's decision. The decision vector for all units is denoted by $Y = (Y_1, \dots, Y_N) \in \mathcal{Y}^N$, with $y \in \{0, 1\}^N$ representing the realized decision outcomes. Additionally, we define a vector of idiosyncratic shocks $\varepsilon = \{\varepsilon_1, \dots, \varepsilon_N\}$, where ε_i is the shock for unit $i \in \mathcal{N}$.

The game, denoted by Γ , comprises:

Players: A set of individuals that we label \mathcal{N} , a social planner;

Payoffs: The preferences (utilities) of units are denoted by $\{U_i(y, X, D, G; \theta)\}_{i=1}^N$. Following [de Paula and Tang \(2012\)](#) and [Galeotti et al. \(2020\)](#), we endow units with a quadratic utility function

$$U_i(y, X, D, G; \theta) = (\alpha_i - \varepsilon_i)y_i + \sum_{j \neq i} \beta_{ij}y_iy_j.$$

where $\alpha_i := \alpha_i(X, D, G)$ and $\beta_{ij} := \beta_{ij}(X, D, G)$ are heterogeneous functions that capture unit i 's *individual utility* and *spillover utility*. The utility of $Y_i = 0$ is normalised to 0.

Given a network G , covariates $X = (X_1, \dots, X_N)$, and a treatment allocation $D = (D_1, \dots, D_N)$, the coefficient α_i on unit i 's choice depends upon their own covariates and treatment status as well as those of all of their neighbors; the coefficient β_{ij} multiplying the quadratic term y_iy_j depends upon their own covariates and treatment status as well as those of unit j . Since the choice variable is binary, if α_i and β_{ij} are unconstrained, then this specification of the utility function is without loss of generality. We endow these utilities with certain properties, which are specified in [Section 3.2.3](#) and [Section 3.4.1](#).

Information: The literature delineates two information environments: *complete information* and *incomplete information*. In a *complete information* setting, players can observe all characteristics of other units. This setting is studied in [Tamer \(2003\)](#), [Ciliberto and Tamer \(2009\)](#), [Bajari et al. \(2010b\)](#) and [Chesher and Rosen \(2020\)](#). Since we consider a large network setting, it may not be plausible for players to have perfect information about all the other units ([Ridder and Sheng, 2020](#)). Therefore, in our headline setting, we follow [Brock](#)

and Durlauf (2001), Aguirregabiria and Mira (2007), Bajari et al. (2010a), and de Paula and Tang (2012), and consider an *incomplete information* setting. All units and the social planner are assumed to observe characteristics X and the network structure G , but the vector of idiosyncratic shocks of units is assumed to be unobservable. The realization of ε_i is unit i 's private information. All players are assumed to have a common belief about the distribution of ε . Formally,

Assumption 12. *The set of idiosyncratic shocks ε must satisfy the following conditions:*

- (i) *The $\{\varepsilon_i\}_{i=1}^N$ is i.i.d. with a known distribution F_ε , which is common knowledge for all the players;*
- (ii) *The distribution of ε_i has a density f_ε , which is bounded above by a constant τ . In addition, f_ε is continuously differentiable;*
- (iii) *$\varepsilon_i \perp X, G, D$ for all $i \in \mathcal{N}$.*

These assumptions are standard in the literature (de Paula and Tang, 2012; Leung, 2015; Ridder and Sheng, 2020). The third assumption can be replaced by a conditional independence assumption if we assume that $F_{\varepsilon|X,D,G}(\cdot|X, D, G)$ is known.

Actions: At the beginning of the game, the social planner assigns treatment D_i to each unit $i \in \mathcal{N}$ to maximize the *planner's welfare*:

$$W_{X,G}(D) = \frac{1}{N} \sum_{i=1}^N \mathbb{E}_\varepsilon [g_i(Y, X, D, G)|X, G], \quad (3.1)$$

subject to the capacity constraint κ (i.e., $\sum_{i=1}^N D_i \leq \kappa$). The expectation in Eq.3.1 is taken with respect to choices Y given the observed covariates X , network structure G , and the treatment allocation rule D^4 . The function $g_i : \mathcal{Y}^N \times \mathcal{X}^N \times \mathcal{D} \times \mathcal{G} \rightarrow \mathbb{R}$ allows social welfare to

⁴With multiple equilibria and no assumption imposed on the equilibrium selection mechanism, the expectation becomes a set in which each element is conditional on a specific equilibrium selection mechanism. This concept will be further formalized in Section 3.2.2.

deviate from the utilitarian welfare function, which corresponds to $g_i(\cdot) = U_i(\cdot)$. We explore two common types of social welfare functions: Utilitarian welfare, and Engagement welfare. Section 3.3 discusses each in detail.

After receiving their allocated treatment, units choose action Y simultaneously to maximize their own payoff given the realization of ε . With complete information unit i 's decision rule would be:

$$Y_i = \mathbb{1}\left\{U_i(1, Y_{-i}, X, D, G) \geq 0\right\}, \quad \forall i \in \mathcal{N}.$$

However, since unit i only has partial information about other units, the realization of Y_{-i} is not observed. Therefore, units make decisions that are best responses given their belief about other units' decisions given the public information and their own type. Formally, in the incomplete information setting,

$$Y_i = \mathbb{1}\left\{\mathbb{E}_\varepsilon[U_i(1, Y_{-i}, X, D, G)|X, D, G, \varepsilon_i] \geq 0\right\}, \quad \forall i \in \mathcal{N}. \quad (3.2)$$

with

$$\mathbb{E}_\varepsilon[U_i(1, Y_{-i}, X, D, G)|X, D, G, \varepsilon_i] = \alpha_i + \sum_{j \neq i} \beta_{ij} \mathbb{E}_\varepsilon[Y_j|X, D, G, \varepsilon_i] - \varepsilon_i.$$

As ε is i.i.d. by Assumption 12, this can be simplified to:

$$\mathbb{E}_\varepsilon[U_i(1, Y_{-i}, X, D, G)|X, D, G, \varepsilon_i] = \alpha_i + \sum_{j \neq i} \beta_{ij} \mathbb{E}_\varepsilon[Y_j|X, D, G] - \varepsilon_i. \quad (3.3)$$

We have now established the game setting. To further elaborate, we introduce additional notation. Consider the action set \mathcal{Y} , defined as $\{0, 1\}$. This set is a totally ordered set, endowed with the usual ordering relation \leq , characterized by reflexivity, antisymmetry, and transitivity⁵. The action profile space \mathcal{Y}^N , formed as a direct product of \mathcal{Y} , also constitutes a partially ordered set (Topkis, 1998, §Example 2.2.1). It is equipped with the *product relation*

⁵**Reflexive:** \leq is reflexive if $y \leq y$, for all $y \in \mathcal{Y}^N$. **Antisymmetric:** \leq is antisymmetric if $y \leq y'$ and $y' \leq y$ implies $y = y'$. **Transitive:** if $y \leq y'$ and $y' \leq y''$ implies $y \leq y''$.

\leq , where for any $y, y' \in \mathcal{Y}$, we have $y \leq y'$ if and only if $y_i \leq y'_i$ for all $i \in \mathcal{N}$. Given that \mathcal{Y}^N is a partially ordered set, we can define a *greatest* and *least* element on it. A strategy profile y is a greatest (least) element on \mathcal{Y}^N if $y \geq y'$ ($y \leq y'$) for all $y' \in \mathcal{Y}^N$. In addition, The *join* of any two elements $y, y' \in \mathcal{Y}^N$, written as $y \vee y'$, is defined as $\inf\{x \in \mathcal{Y}^N : x \geq y, x \geq y'\}$. The *meet*, denoted as $y \wedge y'$, is symmetrically defined as: $\sup\{x \in \mathcal{Y}^N : x \leq y, x \leq y'\}$. In addition, a partially ordered set is called a *lattice* if the join and meet of any pair of elements exist. A lattice is a *complete lattice* if it contains the supremum and infimum of any subsets of it.

3.2.2 Equilibrium

As the game introduced in the previous section features incomplete information, it is a Bayesian game (Harsanyi, 1967), and its Nash equilibria are **Bayesian Nash equilibria** (BNE). We use the pure strategy BNE solution concept. This is defined as:

Definition 3.2.1. (Pure Strategy Bayesian Nash equilibrium) Let Y be the set of all possible decision rules $\{y_i(\varepsilon_i)\}_{i=1}^N$, where $y_i(\varepsilon_i) : \mathbb{R} \rightarrow \{0, 1\}$ specifies unit i 's choice for each realization of their private information ε_i . A pure strategy BNE of game Γ is a strategy profile (y_1^*, \dots, y_N^*) such that, for every $i \in \mathcal{N}$,

$$\mathbb{E}_\varepsilon[U_i(y_i^*, y_{-i}^*)|X, D, G, \varepsilon_i] \geq \mathbb{E}_\varepsilon[U_i(y'_i, y_{-i}^*)|X, D, G, \varepsilon_i],$$

for all $y'_i \in Y$, where $\mathbb{E}_\varepsilon[U_i(\cdot)|X, D, G, \varepsilon_i]$ is defined as in Eq.3.3.

Following Bajari et al. (2010a), we represent the Bayesian Nash equilibrium in the conditional choice probability space. Denote the conditional choice probability (CCP) profile as $\sigma(X, D, G) = \{\sigma_i(X, D, G)\}_{i=1}^N$. An element of the CCP profile:

$$\sigma_i(X, D, G) := \mathbb{E}_\varepsilon[Y_i|X, D, G]. \quad (3.4)$$

Combining the specification of Y_i (Eq.3.2 and Eq.3.3) with Eq.3.4, we have:

$$\sigma_i(X, D, G) = \int \mathbb{1}\left\{\alpha_i + \sum_{j \neq i} \beta_{ij} \sigma_j(X, D, G) \geq \varepsilon_i\right\} dF_\varepsilon. \quad (3.5)$$

Let Ω be a mapping from $[0, 1]^N$ to $[0, 1]^N$ that collects Eq.3.5 for all units. This is a non-linear simultaneous equation system. An equilibrium CCP profile $\sigma^*(X, D, G)$ is a fixed point of this simultaneous equation system:

$$\sigma^*(X, D, G) = \Omega(\sigma^*(X, D, G)). \quad (3.6)$$

This is one representation of the Bayesian Nash equilibrium. Alternatively, given an equilibrium CCP profile σ^* , a fixed X, D, G , and a realization of ε , we can define a Bayesian Nash equilibrium $\{y_i^*\}_{i=1}^N$ as:

$$y_i^* = \mathbb{1}\left\{\alpha_i + \sum_{j \neq i} \beta_{ij} \sigma_j^*(X, D, G) \geq \varepsilon_i\right\}, \quad \forall i \in \mathcal{N}. \quad (3.7)$$

As the right hand side of Eq.3.5 is equal to $F_\varepsilon(\alpha_i + \sum_{j \neq i} \beta_{ij} \sigma_j)$, the existence of a fixed point is guaranteed by the *Brouwer fixed-point theorem* (Brouwer, 1911). As noted in Echenique and Komunjer (2009), this type of simultaneous equation system can have multiple fixed points. In particular, games with strategic complementarity, as in our setting, tend to have a large number of equilibria (Takahashi, 2008). Let $\Sigma := \{\sigma : \sigma = \Omega(\sigma)\}$ denote the set of equilibria. Any equilibrium outcome in this set is a reasonable prediction. In other words, for a given X, D, G and θ , the model predicts a set of equilibrium outcomes σ^* . If we do not assume an equilibrium selection mechanism, this multiplicity introduces incompleteness (Jovanovic, 1989). Incompleteness dramatically increases the difficulty of counterfactual analysis since the model can only identify a set of equilibrium CCP profiles Σ with a newly implemented policy (i.e., a new treatment allocation rule D).

With a newly implemented policy, the realized equilibrium depends on an equilibrium

selection mechanism. Let $\xi : \Sigma \rightarrow [0, 1]$ denote the probability distribution over equilibria, and let $\Delta(\Sigma) := \{\xi : \sum_{\sigma^* \in \Sigma} \xi(\sigma^*) = 1\}$ denote the set of all the probability distributions. The equilibrium selection mechanism is a mapping from the public information (i.e., X, D, G) to one particular element of $\Delta(\Sigma)$. Formally:

Definition 3.2.2. (Equilibrium Selection Mechanism) The equilibrium selection mechanism is denoted by $\lambda(\cdot|X, D, G)$ and the equilibrium selection mechanism space is defined as:

$$\Lambda := \{\lambda : \mathcal{X}^N \times \mathcal{D} \times \mathcal{G} \rightarrow \Delta(\Sigma)\}.$$

If the equilibrium selection mechanism is observable, the conditional choice probability becomes complete by conditioning on λ . Since

$$\Pr[Y_i = 1|X, D, G, \lambda] = \sum_{\sigma^* \in \Sigma} \lambda(\sigma^*|X, D, G)\sigma_i^*, \quad \forall i \in \mathcal{N}. \quad (3.8)$$

There are two main difficulties in characterizing the equilibrium outcome under a newly implemented policy. First, the equilibrium selection mechanism is not directly observable. The identification of an equilibrium selection mechanism from data is studied in [Bajari et al. \(2010b\)](#) and [Aguirregabiria and Mira \(2019\)](#), among others. This is useful in the identification of parameters, since parameter values are independent of λ . For counterfactual analysis, however, there is no guarantee that the equilibrium selection mechanism remains fixed when X, D, G changes. The second difficulty is that the cardinality of Σ increases dramatically with the number of units in the network. Hence, it is not feasible to evaluate the summation in Eq.3.8. To improve the tractability of counterfactual analysis, we focus on a game with strategic complementarity.

3.2.3 Complementarity and Supermodular Games

Strategic complementarity in games implies that, given an ordering of strategies, a player's choice of a higher action incentivizes other players to similarly choose a higher action (Bulow et al., 1985). In economics, complementarity is an important and empirically relevant concept (Molinari and Rosen, 2008). It has many policy applications, such as price setting (Alvarez et al., 2022), house prices (Guren, 2018), technology adoption (Alvarez et al., 2023), as well as the additional examples given in Molinari and Rosen (2008), Lazzati (2015), and Graham and Pelican (2023). The theoretical literature has established that games with strategic complementarities have *robust dynamic stability properties* (Milgrom and Roberts, 1991; Milgrom and Shannon, 1994). This means they converge to the set of Nash equilibria even with simple learning dynamics (Fudenberg and Levine, 1998; Chen and Gazzale, 2004). Topkis (1998) shows that strategic complementarity and supermodularity are equivalent in finite strategy games. The mathematical property *supermodularity* simplifies analysis. It captures the idea of increasing returns between the choice variables. Therefore, to analyze the Bayesian Nash equilibrium of our game, we characterize it as a supermodular game. The definition of supermodular game is:

Definition 3.2.3. Supermodular Game (Milgrom and Roberts, 1990): A game Γ is a supermodular game if, for each $i \in \mathcal{N}$:

- (i) Strategy set \mathcal{Y} is a complete lattice;
- (ii) Payoff $U_i : \mathcal{Y}^N \rightarrow \mathbb{R}$ is order upper semi-continuous in y_i (for fixed y_{-i}) and order continuous in y_{-i} (for fixed y_i), and has a finite upper bound;
- (iii) Payoff U_i is supermodular in y_i (for a fixed y_{-i});
- (iv) Payoff U_i has increasing differences in y_i and y_{-i} .

The definitions of a *supermodular function* and *increasing differences* are:

Definition 3.2.4. Supermodular Function: A function $U: \mathcal{Y}^N \rightarrow \mathbb{R}$ is *supermodular* on \mathcal{Y}^N if for all $y, y' \in \mathcal{Y}^N$:

$$U(y) + U(y') \leq U(y \wedge y') + U(y \vee y').$$

Definition 3.2.5. Increasing Differences: A function $U: \mathcal{Y}^N \rightarrow \mathbb{R}$ has *increasing differences* if for all $y_{-i} < y'_{-i}$ and $y_i < y'_i$:

$$U(y_i, y'_{-i}) - U(y_i, y_{-i}) \leq U(y'_i, y'_{-i}) - U(y'_i, y_{-i}).$$

[Topkis \(1998, §Chapter 2.6.1\)](#) shows that, for a real valued utility function, increasing differences is equivalent to complementarity between units' decisions. Given the definition of a supermodular game above, U_i is a supermodular function on \mathcal{Y}^N if and only if U_i exhibits increasing differences on \mathcal{Y}^N ([Topkis, 1998, §Theorem 2.6.1; §Corollary 2.6.1](#)). Therefore, we have equivalence between complementarity and supermodularity in our game. *Topkis's characterization theorem* ([Topkis, 1978, §Section 3](#)) shows that

$$\frac{\partial^2 U_i(y)}{\partial y_i \partial y_j} \geq 0, \quad \forall j \neq i$$

is a necessary and sufficient condition to guarantee a utility function is a supermodular function on \mathcal{Y}^N . In our specification, this is equivalent to $\beta_{ij} \geq 0$ for $j \neq i$.

Assuming that $\beta_{ij} \geq 0$ for $j \neq i$, our game is a supermodular game since $\{0, 1\}^N$ is a complete lattice and our utility function is continuous. *Tarski's fixed point theorem* ([Tarski, 1955, §Theorem 1](#)) then guarantees the existence of pure strategy Bayesian Nash equilibrium y^* . In particular, there always exists a least BNE \underline{y}^* and a greatest BNE \bar{y}^* ([Milgrom and Roberts, 1990, §Theorem 5](#)). *Tarski's fixed point theorem* can be applied to the conditional choice probability space instead of the strategy profile space to obtain an equivalent result. $[0, 1]^N$ is also a complete lattice, and $\Omega: [0, 1]^N \rightarrow [0, 1]^N$ in [Eq. 3.6](#) is an increasing function given $\beta_{ij} \geq 0$. Therefore, we have a maximal equilibrium CCP profile $\bar{\sigma}^*$ and a minimal

equilibrium CCP profile $\underline{\sigma}^*$. In section 3.3, we show how strategic complementarity simplifies counterfactual analysis.

3.3 Counterfactual Analysis for the Target Population

The goal of this paper is to obtain a treatment allocation that maximizes the equilibrium social welfare of the target population. To achieve this, we first need to characterize the counterfactual equilibrium social welfare if we implement a policy in the target population, which may have a different network structure to the training sample. As the equilibrium selection mechanism is unobservable, the literature typically obtains a point-identified prediction for welfare by assuming how counterfactual policies affect the equilibrium selection mechanism. For example, [Jia \(2008\)](#) assumes that a specific equilibrium is always played, [Aguirregabiria and Mira \(2010\)](#) assumes the equilibrium remains the same after intervention, and [Canen and Song \(2020\)](#) assumes that the equilibrium selection mechanism is invariant to the intervention. However, it is impossible to test the appropriateness of these assumptions given the existing method. In contrast, following [Tamer \(2003\)](#), we are fully agnostic about how policy changes the equilibrium selection mechanism. In other words, the question we focus on is: *If we are agnostic about the equilibrium selection mechanism, what counterfactual outcome does the model predict?*

Recall that our social welfare function is:

$$W_{X,G}(D) = \frac{1}{N} \sum_{i=1}^N \mathbb{E}_{\varepsilon} [g_i(Y, X, D, G) | X, G] . \quad (3.9)$$

With multiple equilibria and no assumption imposed on the equilibrium selection mechanism, our model provides a set-valued equilibrium probability distribution for Y conditional on X, D, G . Therefore, the expectation in Eq.3.9 is also a set, with each element an expectation

conditional on a particular λ . Formally,

$$W_{X,G}(D) = \{W_{X,G,\lambda}(D) : \lambda \in \Lambda\},$$

where

$$W_{X,G,\lambda}(D) = \frac{1}{N} \sum_{i=1}^N \mathbb{E}_{\varepsilon} [g_i(Y, X, D, G) | X, G, \lambda] .$$

This paper considers counterfactual analysis for two standard social welfare functions.

- **Engagement Welfare:** In certain scenarios, a social planner may prioritize goals other than maximizing utilitarian welfare. For instance, in tax auditing, the planner might individualize the assignment of tax audits. Generally, units prefer not to pay taxes, so if maximizing utilitarian welfare were the sole objective, no one would be audited. In this case, a more appropriate target might be the average rate of tax compliance⁶. Engagement welfare is defined as

$$W_{X,G,\lambda}(D) = \frac{1}{N} \sum_{i=1}^N \Pr(Y_i = 1 | X, D, G, \lambda).$$

- **Utilitarian Welfare at Equilibrium:** Utilitarian welfare at equilibrium is the average of the expected utilities of individuals when the system is in equilibrium. This measure is often targeted in policy interventions as it comprehensively reflects overall societal benefit (e.g., [Brock and Durlauf, 2001](#); [Galeotti et al., 2020](#)). An example where the utilitarian welfare target is used is job training programs. Here policymakers allocate limited training resources to unemployed workers to assist them in finding new jobs ([Bloom et al., 1997](#)). In such scenarios, social welfare is defined as:

$$W_{X,G,\lambda}(D) = \frac{1}{N} \sum_{i=1}^N \mathbb{E}_{\varepsilon} [U_i(Y, X, D, G) + \varepsilon_i Y_i | X, D, G, \lambda] , \quad (3.10)$$

⁶This concept of welfare can be broadened to include situations where the policymaker aims to influence outcomes indirectly affected by individual decisions, such as total tax revenue, which depends on individuals' decisions to pay taxes.

which only depends on the expectation of the deterministic component in the utility function⁷. Plugging in our utility function specification, we have:

$$W_{X,G,\lambda}(D) = \frac{1}{N} \sum_{i=1}^N \alpha_i \Pr(Y_i = 1|X, G, \lambda) + \frac{1}{N} \sum_{i=1}^N \sum_{j \neq i} \beta_{ij} \Pr(Y_i Y_j = 1|X, D, G, \lambda).$$

Consider the engagement welfare function. We define bounds for equilibrium welfare, given covariates X , network G and an arbitrary treatment allocation rule D , as:

$$W_{X,G,\lambda}(D) \in \left[\inf_{\lambda \in \Lambda} W_{X,G,\lambda}(D), \sup_{\lambda \in \Lambda} W_{X,G,\lambda}(D) \right].$$

Accordingly, let $\underline{\lambda}$ be the least-favorable equilibrium selection mechanism and $\bar{\lambda}$ the most-favorable equilibrium selection mechanism :

$$\underline{\lambda} := \arg \inf_{\lambda \in \Lambda} W_{X,G,\lambda}(D), \quad \bar{\lambda} := \arg \sup_{\lambda \in \Lambda} W_{X,G,\lambda}(D).$$

In general it is not possible to solve for these two extreme points. There are two obstacles. First, the number of equilibria increases rapidly with the number of units in the network. Evaluating the expectation with respect to the joint distribution of Y thus becomes infeasible. Second, the space of the equilibrium selection mechanisms Λ may be infinite. This complicates any search for the infimum and supremum λ across Λ .

In the existing literature, counterfactual analysis (Ciliberto and Tamer, 2009) often focuses instead on the conditional choice probability (CCP). With no assumptions on the equilibrium selection mechanism, the bounds of the counterfactual CCP are:

$$\Pr(Y_i = 1|X, D, G, \lambda) \in \left[\inf_{\lambda \in \Lambda} \Pr(Y_i = 1|X, D, G, \lambda), \sup_{\lambda \in \Lambda} \Pr(Y_i = 1|X, D, G, \lambda) \right].$$

These bounds can be computed using off-the-shelf methods (e.g., Sheng (2020) for complete

⁷Brock and Durlauf (2001, Section 4) show that introducing a shock term in Eq.3.10 would render the model analytically intractable.

information settings). However, exact bounds of social welfare cannot be directly obtained from the bounds of the CCP. This is because:

$$\inf_{\lambda \in \Lambda} \sum_{i=1}^N \Pr(Y_i = 1 | X, D, G, \lambda) \leq \sum_{i=1}^N \inf_{\lambda \in \Lambda} \Pr(Y_i = 1 | X, D, G, \lambda), \quad (3.11)$$

and

$$\sup_{\lambda \in \Lambda} \sum_{i=1}^N \Pr(Y_i = 1 | X, D, G, \lambda) \geq \sum_{i=1}^N \sup_{\lambda \in \Lambda} \Pr(Y_i = 1 | X, D, G, \lambda). \quad (3.12)$$

That is, the lower (and upper) bound of unit i 's conditional choice probability may be obtained under a different equilibrium selection mechanism to the bound for some unit $j \neq i$. Therefore, bounds for social welfare obtained by summing the bounds on the CCP will generally be loose. However, we show that Eq.3.11 and Eq.3.12 hold with equality in a supermodular game. Formally,

Theorem 3.3.1. (*Engagement Welfare*) *For a supermodular game, the least favorable equilibrium selection rule $\underline{\lambda}$ and the most favorable equilibrium selection rule $\bar{\lambda}$ are:*

$$\underline{\lambda} := \delta_{\underline{\sigma}^*}, \quad \bar{\lambda} := \delta_{\bar{\sigma}^*},$$

where δ_{σ} is the Dirac measure on the set of equilibria Σ . In addition, the following conditions are satisfied:

$$\begin{aligned} \inf_{\lambda \in \Lambda} \sum_{i=1}^N \Pr(Y_i = 1 | X, D, G, \lambda) &= \sum_{i=1}^N \inf_{\lambda \in \Lambda} \Pr(Y_i = 1 | X, D, G, \lambda), \\ \sup_{\lambda \in \Lambda} \sum_{i=1}^N \Pr(Y_i = 1 | X, D, G, \lambda) &= \sum_{i=1}^N \sup_{\lambda \in \Lambda} \Pr(Y_i = 1 | X, D, G, \lambda). \end{aligned}$$

A proof of Theorem 3.3.1 is provided in Appendix A.3.1. This new result characterizes the most and least favorable equilibrium selection rules for aggregate social welfare. This approach enables us to leverage Tarski's fixed point theorem, which significantly reduces

the computational burden by obviating the need to calculate all possible Nash equilibria. Furthermore, it establishes equivalence between the identified set of aggregate social welfare and the aggregation of identified sets of conditional choice probabilities, concepts studied in [Sheng \(2020\)](#) and [Gu et al. \(2022\)](#). This equivalence is not guaranteed to hold in the absence of complementarity. When there are values of ε with unordered multiple equilibria, such as $(Y_1 = 1, Y_2 = 0)$ and $(Y_1 = 0, Y_2 = 1)$ in the two-unit case, the process of identifying the least and most favorable λ is significantly more complicated. Intuitively, the bounds coincide because strategic complementarity guarantees the existence of a least BNE and a greatest BNE for all the values of ε . Since the social welfare function is a monotonically increasing function of σ , it achieves its lower bound at the least equilibrium $\underline{\sigma}^*$ and its upper bound at the greatest equilibria $\bar{\sigma}^*$. By definition, the conditional choice probability $\Pr(Y_i = 1|X, D, G, \lambda)$ also achieves its lower bound under $\underline{\sigma}^*$ and its upper bound under $\bar{\sigma}^*$. The same argument can be applied to utilitarian social welfare to obtain the following corollary.

Corollary 3.3.1. (Utilitarian Welfare at Equilibrium) *Under Assumption 12, given the specification of our utility function, the predicted set of the expected utilitarian welfare under a counterfactual policy D is given as:*

$$W_{X,G,\lambda}(D) \in \left[\frac{1}{N} \sum_{i=1}^N \alpha_i f(\alpha_i) + \frac{1}{N} \sum_{i=1}^N \sum_{j \neq i} \beta_{ij} \underline{\sigma}_i^* \underline{\sigma}_j^*, \frac{1}{N} \sum_{i=1}^N \alpha_i f(\alpha_i) + \frac{1}{N} \sum_{i=1}^N \sum_{j \neq i} \beta_{ij} \bar{\sigma}_i^* \bar{\sigma}_j^* \right],$$

where

$$f(\alpha_i) = \begin{cases} \Pr(Y_i = 1|X, D, G, \underline{\lambda}) & \text{if } \alpha_i > 0 \\ \Pr(Y_i = 1|X, D, G, \bar{\lambda}) & \text{if } \alpha_i \leq 0. \end{cases}$$

A proof of Corollary 3.3.1 is provided in Appendix A.2.1. This result implies that it is sufficient to compute the minimal and maximal equilibrium CCP profile for the utilitarian welfare in the incomplete information setting. This result does not hold in the complete information setting, where we provide an alternative approach to compute the bounds of the identified set.

Remark 3.3.1. The infimum and supremum of the planner’s welfare, calculated over the equilibrium selection mechanism, are equivalent to the *Choquet integral* (see [Denneberg, 1994](#) and [Gilboa, 2009](#)) of the planner’s welfare with respect to the capacity and its conjugate. The capacity v and its conjugate v^* are non-additive probability measures defined on the set of equilibrium Σ . In our case, $\Pr(Y_i = 1|X, D, G, \underline{\lambda})$ is equivalent to Choquet integration with respect to the capacity $v(A)$ where $A = \{y^* : y_i^* = 1\}$. Analogously, $\Pr(Y_i = 1|X, D, G, \bar{\lambda})$ is equivalent to Choquet integration with respect to the conjugate $v^*(A)$ where $A = \{y^* : y_i^* = 1\}$. We refer to [Kaido and Zhang \(2023\)](#) for the definition of capacity and a more detailed discussion on this topic. Despite the complexity typically associated with Choquet integration, which often requires approximate solution by simulation methods, our method provides a closed-form expression for the identified set which can be solved without numerical error. Moreover, the applications of Choquet integration extend to robust Bayesian analysis to manage multiple priors (see [Chamberlain \(2000a\)](#) and [Giacomini et al. \(2021\)](#)), which is analogous to a setting with multiple equilibria.

The details of the computation of the maximal and minimal equilibrium conditional choice probabilities (CCPs) are discussed in Section 3.4.1. The arguments above apply not only to the counterfactual analysis of treatment allocation policies but also to policy interventions that alter covariates or the network structure.

3.4 Treatment Allocation

Our model allows for multiple equilibria, but can only predict a set of possible equilibrium outcomes, denoted as $W_{X,G}(D)$. Consequently, the expected value calculation that determines social welfare is not well-defined without specifying the equilibrium selection mechanism. Drawing on game theory ([Morris et al., 2024](#)) and robust decision theory ([Chamberlain, 2000a](#)), we apply the maximin welfare criterion to select a treatment allocation rule. This approach involves a social planner opting for choices that lead to higher welfare while

preparing for the worst-case scenario of the least favorable equilibrium. Essentially, the planner anticipates the minimal equilibrium will be realized. For example, [Segal \(2003, Section 4.1.3\)](#) discusses scenarios in contracting where the worst-case equilibrium corresponds to the Pareto-efficient outcome for the parties involved. Moreover, in settings where action 0 is the default, games exhibiting strategic complementarity tend to converge toward their minimal equilibrium.

The planner chooses D to maximise welfare under the assumption that the minimal equilibrium, conditional on the chosen D , will be realized. We denote the set of feasible allocations by $\mathcal{D}_\kappa := \{D \in \mathcal{D} : \sum_{i=1}^N D_i \leq \kappa\}$. Formally:

$$D^* = \arg \max_{D \in \mathcal{D}_\kappa} \min_{\lambda \in \Lambda} W_{X,G,\lambda}(D). \quad (3.13)$$

Recall that, by [Theorem 3.3.1](#), the lower bound of equilibrium social welfare equals the summation of individual welfares. Thus, the maximin welfare optimisation problem simplifies to:

$$\max_{D \in \mathcal{D}_\kappa} W_{X,G,\underline{\lambda}}(D),$$

where $W_{X,G,\underline{\lambda}}(D)$ is the social welfare function evaluated at the minimal equilibrium. This formulation converts the maximin welfare problem into a straightforward maximization problem, providing a clear framework for solving the optimal treatment allocation problem.

3.4.1 Implementation

Identification

The preceding discussion has assumed that true parameter values are observed. To implement our proposed method, we first describe the identification of structural parameters using a training sample. Details on the estimation procedure are provided in [Section 3.4.1](#).

The discussion of counterfactual analysis in Section 3.3 makes no assumptions about the functional form of parameters $\{\alpha_i\}_{i \in \mathcal{N}}$ and $\{\beta_{ij}\}_{i,j \in \mathcal{N}}$. However, observable data is limited to units' choices, covariates X , the network structure G , and a predetermined treatment allocation D . In practice, we are restricted by what it is possible to identify given this data.

For identification, we follow [Bajari et al. \(2010a\)](#) and adopt the inverse-CDF procedure⁸. Let ε^n be the private information in the training data, which is distinct from ε in the target population. Recall from Eq. 3.7 that, given an equilibrium conditional choice probability profile in the training data σ^{data} , unit i chooses their actions according to the decision rule

$$Y_i^n = \mathbb{1}\left\{\alpha_i + \sum_{j \neq i} \beta_{ij} \sigma_j^{data} \geq \varepsilon_i^n\right\}, \quad \forall i \in \mathcal{N}.$$

The equilibrium CCP profile is thus:

$$\sigma_i^{data} = \int \mathbb{1}\left\{\varepsilon_i^n \leq \alpha_i + \sum_{j \neq i} \beta_{ij} \sigma_j^{data}\right\} dF_{\varepsilon^n} = F_{\varepsilon^n}\left[\alpha_i + \sum_{j \neq i} \beta_{ij} \sigma_j^{data}\right]. \quad (3.14)$$

Taking the inverse of the CDF of ε on both sides in Eq. 3.14 yields:

$$F_{\varepsilon^n}^{-1}(\sigma_i^{data}) = \alpha_i + \sum_{j \neq i} \beta_{ij} \sigma_j^{data}. \quad (3.15)$$

Even assuming that the equilibrium CCP profile in the training data is observable, identifying all the parameters in Eq. 3.15 remains challenging. Determining all utility parameters involves solving for $N \times N$ unknown parameters on the right-hand side of the above equation. However, the left-hand side of Eq. 3.15, only provides information about N scalars. Given these limitations, we define our utility function as follows to ensure identifiability and allow

⁸This approach builds on [Hotz and Miller \(1993\)](#) and [Aguirregabiria and Mira \(2007\)](#).

for the analysis of general treatment effects:

$$\begin{aligned}
U_i(y, X, D, G) = & \overbrace{y_i(\theta_0 + \theta_1 D_i + X_i^\top \theta_2 + X_i^\top \theta_3 D_i)}^{\alpha_i} + \frac{1}{|\mathcal{N}_i|} \sum_{j \neq i} \theta_4 m_{ij} G_{ij} D_j - \varepsilon_i \\
& + \sum_{j \neq i} \underbrace{\frac{1}{|\mathcal{N}_i|} (\theta_5 + \theta_6 D_i D_j)}_{\beta_{ij}} m_{ij} G_{ij} y_i y_j,
\end{aligned} \tag{3.16}$$

where $m_{ij} = m(X_i, X_j)$ is a (bounded) real-valued function of personal characteristics. m_{ij} measures the distance between unit i 's characteristics and unit j 's characteristics; the spillover effect is weighted by how similar two units appear. The utility that unit i derives from an action is the sum of the net benefits that they accrue from their own actions and from those of their neighbors. We assume that a unit's utility is only affected by the actions of their direct neighbors, not one-link-away contacts. The payoff of action $Y_i = 1$ has six components. When unit i chooses action $Y_i = 1$, they receive utility θ_0 irrespective of their allocated treatment. They also receive additional utility $\theta_1 D_i$ depending upon their own treatment status. Their utility also includes a heterogeneous component $X_i^\top (\theta_2 + \theta_3 D_i)$, which depends upon their characteristics X_i . Next, there is a spillover effect from the action of unit j . If unit j is a neighbor of unit i that receives treatment, then this provides $\theta_4 m_{ij}$ additional utility to unit i . The fifth and sixth components represent strategic complementarity. If unit j is a neighbor of unit i and selects $Y_j = 1$, then unit i 's payoff is increased by $\theta_5 m_{ij}$. The final component corresponds to choice spillovers between neighbors who receive treatment. If both unit i and unit j receive treatment and both choose action 1, unit i receives additional utility $\theta_6 m_{ij}$.

Accordingly, the structural parameters θ are uniquely determined by the conditional choice probabilities in the training sample, thus identifying the payoff function. This discussion provides only an informal overview of the identification process; a formal proof is available in [Bajari et al. \(2010a\)](#).

Estimation

For estimation, we employ the two-step maximum likelihood estimation procedure of [Leung \(2015\)](#). The first step involves estimating the equilibrium conditional choice probability from the training data. In the second step, structural parameters are estimated by maximizing the likelihood function given the estimated CCP profiles. To distinguish the training data from the target population, we denote covariates as $X = \{X_i\}_{i=1}^n$, the treatment allocation as $D = \{D_i\}_{i=1}^n$, decisions as $Y = \{Y_i\}_{i=1}^n$, and the network structure as $G = \{G_{ij}\}_{i,j=1}^n$. In addition, let $S = (X, D, G)$.

Let $\{\sigma_i^{data}\}_{i=1}^N$ be the CCP in the training data. Given that the training dataset contains only a single large network, two necessary conditions on the training data are required to estimate the conditional choice probability: *symmetric equilibrium*⁹ ([Leung, 2015](#)) and *network decaying dependence condition* ([Xu, 2018](#)). In general, each unit's choice depends on all public information across the network G (i.e., X and D of all units), although direct payoffs may depend only on immediate spillovers. Under the *network decaying dependence condition*, it is sufficient to consider only interactions within a relatively small distance.

Several estimation approaches have been proposed for CCP. These including the empirical frequency estimator ([Hotz and Miller, 1993](#)), sieve estimation ([Bajari et al., 2010a](#)), flexible logit estimation ([Arcidiacono and Miller, 2011](#)), and logit Lasso estimation ([Chernozhukov et al., 2022](#)). Here we leave aside the question of the most suitable procedure. Instead, we assume the existence of an estimator that satisfies the following statistical property:

Assumption 13. (Sub-Gaussian CCP Estimator) *There exists a positive constant C_σ such that for every $t \geq 0$, we have*

$$\Pr \left(|\hat{\sigma}_i^{data} - \sigma_i^{data}| \geq t \mid S, \sigma^{data} \right) \leq 2 \exp \left(- nt^2 / C_\sigma^2 \right), \quad \forall i = 1, \dots, n.$$

⁹A symmetric equilibrium implies that two units will exhibit identical conditional choice probabilities if they receive the same treatment, share identical covariates, and have comparable neighbors, specifically in terms of the neighbors' treatments and covariates.

This assumption is satisfied by the empirical frequency estimator (Leung, 2015; Ridder and Sheng, 2020) and logit/probit estimation¹⁰. The ultimate goal is to choose $\hat{\theta}$ that maximizes the likelihood function. As σ^{data} is unobserved, we replace σ^{data} in the likelihood function with $\hat{\sigma}^{data}$ and estimate θ by maximizing the quasi-likelihood function $\hat{Q}_n(\hat{\sigma}^{data}, \theta)$.

$$\hat{Q}_n(\hat{\sigma}^{data}, \theta) = \frac{1}{n} \sum_{i=1}^n Y_i \log (F_{\varepsilon}(\hat{Z}_i^T \theta)) + (1 - Y_i) \log (1 - F_{\varepsilon}(\hat{Z}_i^T \theta)),$$

where

$$\hat{Z}_i = \left(1, D_i, X_i^T, X_i^T D_i, \frac{1}{|\mathcal{N}_i|} \sum_{j \neq i} m_{ij} G_{ij} D_j, \frac{1}{|\mathcal{N}_i|} \sum_{j \neq i} m_{ij} G_{ij} \hat{\sigma}_j^{data}, \frac{1}{|\mathcal{N}_i|} \sum_{j \neq i} m_{ij} G_{ij} \hat{\sigma}_j D_i D_j \right)^T. \quad (3.17)$$

In addition, Z_i denotes the vector of regressors that would be obtained if we replaced $\hat{\sigma}_i^{data}$ in \hat{Z}_i with the true conditional choice probability σ_i^{data} .

Computation of Equilibria

After obtaining estimated parameters, we compute the set of equilibrium social welfare for given covariates X , network structure G , and a treatment allocation D in the target population. The lower bound and upper bound of this set are: $\Pr(Y_i = 1 | X, D, G, \underline{\lambda}; \hat{\theta})$ and $\Pr(Y_i = 1 | X, D, G, \bar{\lambda}; \hat{\theta})$ for all unit $i \in \mathcal{N}$. We first rewrite these two conditional probabilities as:

$$\begin{aligned} \Pr(Y_i = 1 | X, D, G, \bar{\lambda}; \hat{\theta}) &= \int \mathbb{1} \left\{ \hat{\alpha}_i + \sum_{j \neq i} \hat{\beta}_{ij} \Pr(Y_j = 1 | X, D, G, \bar{\lambda}; \hat{\theta}) \geq \varepsilon_i \right\} dF_{\varepsilon}, \\ \Pr(Y_i = 1 | X, D, G, \underline{\lambda}; \hat{\theta}) &= \int \mathbb{1} \left\{ \hat{\alpha}_i + \sum_{j \neq i} \hat{\beta}_{ij} \Pr(Y_j = 1 | X, D, G, \underline{\lambda}; \hat{\theta}) \geq \varepsilon_i \right\} dF_{\varepsilon}. \end{aligned} \quad (3.18)$$

¹⁰By Hoeffding's inequality, the frequency estimator easily satisfies the Assumption 13. The probit/logit estimator satisfies the Assumption 13 by the same argument as the second-stage MLE estimator in Section 3.5.

From Theorem 3.3.1, $\Pr(Y_i = 1|X, D, G, \lambda)$ achieves its upper (lower) bound when the equilibrium is $\bar{\sigma}^*$ ($\underline{\sigma}^*$) for all $i \in \mathcal{N}$. Therefore,

$$\Pr(Y_i = 1|X, D, G, \bar{\lambda}; \hat{\theta}) = \hat{\sigma}_i^*, \quad \forall i \in \mathcal{N},$$

$$\Pr(Y_i = 1|X, D, G, \underline{\lambda}; \hat{\theta}) = \hat{\underline{\sigma}}_i^*, \quad \forall i \in \mathcal{N},$$

where $\hat{\sigma}_i^*$ and $\hat{\underline{\sigma}}_i^*$ represent the estimators for the maximal and minimal equilibria, respectively. Hence, we need only compute the least and greatest equilibrium CCP profile. Topkis (1979) provides an easily implemented algorithm that is guaranteed to converge to the least and greatest equilibrium point of a supermodular game. Hold X, D, G fixed. To obtain the greatest fixed point $\bar{\sigma}^*$, begin with $\bar{\sigma}^0 = \{1, \dots, 1\}$. Define a sequence $\{\bar{\sigma}^t\}_{t=0}^T : \bar{\sigma}^{t+1} = \Omega(\bar{\sigma}^t)$. By construction, $\bar{\sigma}^0 \geq \Omega(\bar{\sigma}^0)$. Since $\Omega(\cdot)$ is an increasing function, $\Omega(\bar{\sigma}^0) \geq \Omega(\bar{\sigma}^1)$. Therefore, $\bar{\sigma}^0 \geq \bar{\sigma}^1 \geq \dots \geq \bar{\sigma}^T$. Suppose the iteration converges on the M -th step. Then $\bar{\sigma}^M$ is the greatest equilibrium since, for all the other σ^* , $\bar{\sigma}^M = \Omega^M(\sigma^0) \geq \Omega^M(\sigma^*) = \sigma^*$.

With a symmetric argument, we can obtain the least equilibrium CCP profile. Here we begin with $\underline{\sigma}^0 = \{0, \dots, 0\}$. Define a sequence $\{\underline{\sigma}^t\}_{t=0}^T : \underline{\sigma}^{t+1} = \Omega(\underline{\sigma}^t)$. By construction, we have $\underline{\sigma}^0 \leq \Omega(\underline{\sigma}^0)$. Again, $\Omega(\underline{\sigma}^0) \leq \Omega(\underline{\sigma}^1)$. Therefore, $\underline{\sigma}^0 \leq \underline{\sigma}^1 \leq \dots \leq \underline{\sigma}^T$. Suppose the iteration converges on the M -th step. Then $\underline{\sigma}^M$ is the least equilibrium since, for all the other σ^* , $\underline{\sigma}^M = \Omega^M(\sigma^0) \leq \Omega^M(\sigma^*) = \sigma^*$.

Greedy Algorithm

The previous sections describe the estimation of parameters and computation of the least equilibrium CCP profile. In this section, we propose an algorithm to allocate treatment in a manner that maximizes the worst-case social welfare given the estimated parameters. Define the empirical welfare function to be the welfare function with estimated structural parameters:

$$W_{X,G,\lambda}^n(D) = W_{X,G,\lambda}(D; \hat{\theta}).$$

We seek to maximize the empirical welfare evaluated at the minimal equilibrium:

$$\tilde{D} = \arg \max_{D \in \mathcal{D}_\kappa} W_{X,G,\underline{\sigma}}^n(D) = \arg \max_{D \in \mathcal{D}_\kappa} W_{X,G,\underline{\sigma}^*}^n(D). \quad (3.19)$$

As shown in Eq.3.18, $\underline{\sigma}^*$ is a solution to a non-linear simultaneous equation system. The conditional choice probability $\underline{\sigma}_i^*$ of unit i depends non-linearly on the conditional choice probability $\underline{\sigma}_j^*$ and treatment assignment of their neighbors $\{D_j : j \in \mathcal{N}_i\}$. Therefore, when a treatment is assigned to one unit, it not only influences their behavior but also leads to spillover effects through the network. Hence, Eq.3.19 is a complicated combinatorial optimization problem. We propose a greedy algorithm¹¹ (Algorithm 5) to solve this problem heuristically.

Intuitively, our greedy algorithm assigns treatment to the unit that contributes most to the welfare objective, and repeats this until a capacity constraint binds. Specifically, in each round, Algorithm 5 computes the marginal gain of receiving treatment for each untreated unit, evaluated at the least equilibrium CCP profile. We refer to the unit whose treatment induces the largest increase in the worst-case welfare as the most influential unit for that round.

¹¹A greedy algorithm is a heuristic approach used in optimization problems; it makes a series of choices that appear to offer the most immediate benefit, building a solution step by step to achieve locally optimal results.

Algorithm 5: Maximizing Over Treatment Allocation Rules

Input: adjacency matrix G , covariates X , parameters $\hat{\theta}$, capacity constraint κ

Output: Treatment allocation regime \hat{D}_G

Initialization: $D \leftarrow 0_{N \times 1}$

if $\sum_{i=1}^N D_i < \kappa$ **then**

for i with $D_i = 0$ **do**

$D_i \leftarrow 1$, denote new treatment vector as D'

$\underline{\sigma}^*(D') \leftarrow$ Computing the minimal equilibrium CCP profile given D'

$\Delta_i \leftarrow W_{X,G,\underline{\sigma}^*}^n(D') - W_{X,G,\underline{\sigma}^*}^n(D)$

end

$i^* \leftarrow \arg \max_i \Delta_i$

$D_{i^*} \leftarrow 1$

else

$\hat{D}_G \leftarrow D$

end

3.5 Theoretical Analysis

In this section, we analyze the theoretical properties of our proposed treatment allocation method. To simplify notation, denote the welfare of the targeted population $W_{X,G,\underline{\sigma}^*}(D)$ as $W(D)$, and empirical welfare $W_{X,G,\underline{\sigma}^*}^n(D)$ as $W_n(D)$. In addition, let $W(D^*)$ denote the welfare of the target population at its global optimizer D^* , and $W(\hat{D}_G)$ denote the welfare of the target population welfare under the treatment allocation rule obtained by our proposed method. Let the *regret* of the proposed treatment allocation policy be:

$$R(\hat{D}_G) := \max_{D \in \mathcal{D}} W(D) - W(\hat{D}_G).$$

We evaluate the performance of our proposed treatment allocation method using *expected regret*, which is defined as:

$$\mathbb{E}_{\varepsilon^n} \left[R(\hat{D}_G) | S, \sigma^{data} \right] := \max_{D \in \mathcal{D}} W(D) - \mathbb{E}_{\varepsilon^n} \left[W(\hat{D}_G) | S, \sigma^{data} \right],$$

where the expectation $\mathbb{E}_{\varepsilon^n}[\cdot]$ is taken with respect to the uncertainty in the training data¹² conditional on the observed covariates X , treatment allocation D , network G , and equilibrium σ^{data} . This is because the randomness in our proposed method primarily arises from utilizing the estimated parameters, which involve only the training data. This criterion captures the average welfare loss when implementing estimated policy \hat{D}_G relative to the maximum feasible population welfare. Recall that $W_n(\tilde{D})$ is the maximum for empirical welfare. We decompose regret into (eight terms):

$$\begin{aligned} W(D^*) - W(\hat{D}_G) &= W(D^*) - W_n(D^*) + W_n(D^*) - W_n(\tilde{D}) \\ &\quad + W_n(\tilde{D}) - W_n(\hat{D}_G) + W_n(\hat{D}_G) - W(\hat{D}_G). \end{aligned}$$

The first term measures the deviation arising from the use of the empirical social welfare function. This term is bounded by:

$$W(D^*) - W_n(D^*) \leq \sup_{D \in \mathcal{D}} |W_n(D) - W(D)|.$$

The second term measures the performance of the population welfare maximizer in the empirical social welfare function. This term is bounded by:

$$W_n(D^*) - W_n(\tilde{D}) \leq 0.$$

The third term measures the loss caused by using a greedy algorithm to solve the optimization problem. This is discussed in Section 3.5.3. The final term also measures regret introduced by using the empirical social welfare function. This term is bounded by:

$$W_n(\hat{D}_G) - W(\hat{D}_G) \leq \sup_{D \in \mathcal{D}} |W_n(D) - W(D)|.$$

¹²In the incomplete information setting, the unobserved variables represent units' private information. If the units in the training data are the target population, this may coincide for the training data and the target population. In this case, the expectation in the regret is taken with respect to the uncertainty in the target population. All discussions in this section are otherwise unchanged.

Combining all the above results, we conclude that expected regret is bounded by:

$$\mathbb{E}_{\varepsilon^n} \left[R(\hat{D}_G) | S, \sigma^{data} \right] \leq 2\mathbb{E}_{\varepsilon^n} \left[\max_{D \in \mathcal{D}} |W_n(D) - W(D)| | S, \sigma^{data} \right] + \mathbb{E}_{\varepsilon^n} \left[W_n(\tilde{D}) - W_n(\hat{D}_G) | S, \sigma^{data} \right]. \quad (3.20)$$

In the remainder of this section, we provide a non-asymptotic upper bound for expected regret.

3.5.1 Sampling Uncertainty

For illustrative purposes, this section focuses on engagement welfare. We begin by addressing the regret resulting from the use of estimates in place of true parameters in the payoff function. This represents the sampling uncertainty of the proposed method. We impose the following assumption on the parameter space:

Assumption 14. (Compactness) *The parameter θ lies in a compact set $\Theta \subseteq \mathbb{R}^{d_\theta}$.*

Assumption 14 is standard. We now proceed to characterize the sampling uncertainty associated with using the empirical welfare function.

Lemma 3.5.1. *Under Assumptions 12 and 14,*

$$\mathbb{E}_{\varepsilon^n} \left[\max_{D \in \mathcal{D}} |W_n(D) - W(D)| | S, \sigma^{data} \right] \leq C_1 \mathbb{E}_{\varepsilon^n} \left[\|\hat{\theta} - \theta_0\|_1 | S, \sigma^{data} \right],$$

where C_1 is a constant that depends on the distribution F_{ε^n} , and the supports of the parameter space, the covariates space \mathcal{X} , the network space \mathcal{G} and the treatment allocation space \mathcal{D} .

A proof of Lemma 3.5.1 is provided in Appendix A.2.2. Lemma 3.5.1 enables us to characterize the regret of maximizing the empirical welfare through the sampling uncertainty of the structural parameter estimators (i.e., $\mathbb{E}_{\varepsilon^n} [\|\hat{\theta} - \theta\|_1 | S, \sigma^{data}]$). As there is no closed-form expression for MLE $\hat{\theta}$ in our case, we study the sampling uncertainty of $\hat{\theta}$ through the sampling

uncertainty of its associated empirical process

$$\mathbb{E}_{\varepsilon^n} [\mathbb{G}_n(\hat{\theta}) - \mathbb{G}_n(\theta_0) | S, \sigma^{data}], \quad (3.21)$$

where the empirical process $\mathbb{G}_n(\theta)$ is defined as

$$\mathbb{G}_n(\theta) := \hat{\mathbb{M}}(\theta) - M(\theta),$$

$$\hat{\mathbb{M}}(\theta) = \frac{1}{n} \sum_{i=1}^n Y_i \log(F_\varepsilon(\hat{Z}_i^\top \theta)) + (1 - Y_i) \log(1 - F_\varepsilon(\hat{Z}_i^\top \theta)),$$

and

$$M(\theta) = \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{\varepsilon^n} \left[Y_i \log(F_\varepsilon(Z_i^\top \theta)) + (1 - Y_i) \log(1 - F_\varepsilon(Z_i^\top \theta)) \middle| S, \sigma^{data} \right].$$

Recall that \hat{Z}_i , as defined in Eq.3.17, serves as the regressor in our likelihood function. Since we are using a quasi-likelihood ML estimator, the criterion function $\hat{\mathbb{M}}(\theta)$ is evaluated at the estimated equilibrium in the data $\hat{\sigma}^{data}$. As a result, Eq.3.21 contains two sources of sampling uncertainty: uncertainty from $\hat{\theta}$, and uncertainty from $\hat{\sigma}^{data}$. The difference between $\mathbb{G}_n(\hat{\theta})$ and $\mathbb{G}_n(\theta_0)$, is given by:

$$\mathbb{G}_n(\hat{\theta}) - \mathbb{G}_n(\theta_0) = \hat{\mathbb{M}}(\hat{\theta}) - \hat{\mathbb{M}}(\theta_0) + M(\theta_0) - M(\hat{\theta}). \quad (3.22)$$

To study the relationship between the estimator and its associated empirical process, we start with a second-order Taylor expansion with Lagrange remainder for both terms in Eq.3.22:

$$\hat{\mathbb{M}}(\theta_0) - \hat{\mathbb{M}}(\hat{\theta}) = \frac{1}{2}(\hat{\theta} - \theta_0)^\top \nabla_{\hat{\theta}}^2 \hat{\mathbb{M}}(\hat{\theta})(\hat{\theta} - \theta_0), \quad (3.23)$$

$$M(\hat{\theta}) - M(\theta_0) = \frac{1}{2}(\hat{\theta} - \theta_0)^\top \nabla_{\hat{\theta}}^2 M(\hat{\theta})(\hat{\theta} - \theta_0), \quad (3.24)$$

for some $\hat{\theta} \in \mathbb{R}^{d_\theta}$ and $\hat{\theta} \in \mathbb{R}^{d_\theta}$ on the segment from θ_0 to $\hat{\theta}$. Let η_{max}^0 denote the largest

eigenvalue (in magnitude) of $\nabla_{\hat{\theta}}^2 \hat{\mathbb{M}}(\hat{\theta})$, and η_{max}^1 denote the largest eigenvalue of $\nabla_{\hat{\theta}}^2 M(\hat{\theta})$. By Assumption 12 (ii), the Hessian matrix is symmetric. Therefore, by the *Courant-Fischer Theorem*¹³, we can characterize the relationship between the parameter sampling uncertainty and the deviation of the criterion function through:

$$\hat{\mathbb{M}}(\theta_0) - \hat{\mathbb{M}}(\hat{\theta}) \leq \frac{1}{2} \eta_{max}^0 \|\hat{\theta} - \theta_0\|_2^2. \quad (3.25)$$

$$M(\hat{\theta}) - M(\theta_0) \leq \frac{1}{2} \eta_{max}^1 \|\hat{\theta} - \theta_0\|_2^2. \quad (3.26)$$

Combining Eq.3.22 with Eq.3.25 and Eq.3.26 yields

$$-(\mathbb{G}_n(\hat{\theta}) - \mathbb{G}_n(\theta_0)) \leq \frac{1}{2} (\eta_{max}^0 + \eta_{max}^1) \|\hat{\theta} - \theta_0\|_2^2. \quad (3.27)$$

Applying the mean value Theorem to the left-hand side of Eq.3.27, we have

$$\mathbb{G}_n(\hat{\theta}) - \mathbb{G}_n(\theta_0) = (\hat{\theta} - \theta_0)^\top \nabla_{\tilde{\theta}} \mathbb{G}_n(\tilde{\theta}), \quad (3.28)$$

for some $\tilde{\theta} \in \mathbb{R}^{d_\theta}$ on the segment from θ_0 to $\hat{\theta}$. Since $\hat{\theta}$ is the maximizer of $\hat{\mathbb{M}}(\cdot)$ and θ_0 is the maximizer of $M(\cdot)$, Eq.3.28 must be positive given the definition of \mathbb{G}_n . By the Cauchy–Schwarz inequality,

$$\mathbb{G}_n(\hat{\theta}) - \mathbb{G}_n(\theta_0) = (\hat{\theta} - \theta_0)^\top \nabla_{\tilde{\theta}} \mathbb{G}_n(\tilde{\theta}) = |(\hat{\theta} - \theta_0)^\top \nabla_{\tilde{\theta}} \mathbb{G}_n(\tilde{\theta})| \leq \|\hat{\theta} - \theta_0\|_2 \|\nabla_{\tilde{\theta}} \mathbb{G}_n(\tilde{\theta})\|_2. \quad (3.29)$$

Combining Eq.3.27 and Eq.3.29,

$$-\|\hat{\theta} - \theta_0\|_2 \|\nabla_{\tilde{\theta}} \mathbb{G}_n(\tilde{\theta})\|_2 \leq \frac{1}{2} (\eta_{max}^0 + \eta_{max}^1) \|\hat{\theta} - \theta_0\|_2^2.$$

Assuming $\hat{\theta}$, an MLE estimator, lies in the interior of the parameter space, the Hessian matrix

¹³The largest eigenvalue η_{max} of a $C \times C$ symmetric matrix M is given by the maximum Rayleigh quotient (i.e., $\eta_{max} = \max_{A \in \mathbb{R}^C \setminus \{0\}} \frac{A^\top M A}{A^\top A}$).

$\nabla_{\theta}^2 \hat{\mathbb{M}}(\hat{\theta})$ is negative definite. Therefore, η_{max}^0 is negative. In addition, as θ_0 is the maximizer of $M(\theta)$, η_{max}^1 also must be negative. Hence,

$$\begin{aligned} \|\hat{\theta} - \theta_0\|_1 &\leq d_{\theta} \|\hat{\theta} - \theta_0\|_2 \leq -\frac{2d_{\theta}}{\eta_{max}^0 + \eta_{max}^1} \|\nabla_{\theta} \mathbb{G}_n(\tilde{\theta})\|_2 \\ &\leq -\frac{2d_{\theta}}{\eta_{max}^0 + \eta_{max}^1} \|\nabla_{\theta} \mathbb{G}_n(\tilde{\theta})\|_1. \end{aligned} \quad (3.30)$$

If η_{max}^0 and η_{max}^1 did not depend on the sample size n (i.e., if they were constant), we could study the finite sample properties of $\|\hat{\theta} - \theta_0\|_1$ through the finite sample properties of the empirical process $\nabla_{\theta} \mathbb{G}_n(\tilde{\theta})$. However, the Hessian matrix is a function that depends on the sample, so η_{max}^0 and η_{max}^1 also depend on n . This prevents us from characterizing the finite sample properties of our estimator.

To overcome this difficulty, we establish a uniform constant upper bound for the largest eigenvalue of the Hessian matrices, which guarantees their strict negativity. We denote the uniform constant upper bound for the largest eigenvalue as maximal largest eigenvalue and as smallest To obtain a over all the possible samples, we impose the following three assumptions:

Assumption 15. (Shape) $F_{\varepsilon}(x)$ satisfies the condition: $\frac{F'_{\varepsilon}(x)^2}{F_{\varepsilon}(x)-1} < F''_{\varepsilon}(x) < \frac{F'_{\varepsilon}(x)^2}{F_{\varepsilon}(x)}$ for all $x \in \mathbb{R}$.

Assumption 16. (Full Rank) Let Z denote $[Z_1, \dots, Z_n]$ and \hat{Z} denote $[\hat{Z}_1, \dots, \hat{Z}_n]$. We assume that the matrices Z and \hat{Z} each have full row rank.

Assumption 17. (Non-Zero Treatment) There exists a constant $C_d > 0$ such that $\frac{1}{n} \sum_{i=1}^n D_i \geq C_d, \forall n \in \mathbb{Z}_+$.

Assumption 15 imposes a regularity condition on the shape of the CDF function. The following are two examples of common distributions that satisfy this assumption:

- **Logistic Distribution:** The cumulative distribution function of the Logistic distribution is: $F_{\varepsilon}(x) = \frac{1}{1+\exp(-x)}$. The corresponding probability density function is: $F'_{\varepsilon}(x) =$

$\frac{\exp(-x)}{(1+\exp(-x))^2}$. Finally, the second derivative of the CDF is: $F''_\varepsilon(x) = \frac{\exp(-3x) - \exp(-x)}{(1+\exp(-x))^4}$. Therefore:

$$\frac{F'_\varepsilon(x)^2}{F_\varepsilon(x) - 1} = \frac{-\exp(-x) - \exp(-2x)}{(1 + \exp(-x))^4} < F''_\varepsilon(x) < \frac{\exp(-3x) + \exp(-2x)}{(1 + \exp(-x))^4} = \frac{F'_\varepsilon(x)^2}{F_\varepsilon(x)}.$$

- **Gaussian Distribution:** Denote the Gaussian cumulative distribution function by $F_\varepsilon(x) = \Phi(x)$, its first derivative (the probability density function) by $F'_\varepsilon(x) = \phi(x)$, and its second derivative by $F''_\varepsilon(x) = -x\phi(x)$. We aim to demonstrate that: $F''_\varepsilon(x) < \frac{F'_\varepsilon(x)^2}{F_\varepsilon(x)}$. Substituting the known forms of $F'_\varepsilon(x)$ and $F_\varepsilon(x)$, this inequality simplifies to: $\frac{\phi(x)}{\Phi(x)} > -x$. For $x \geq 0$, this inequality is always satisfied. When $x < 0$, we require that: $\frac{\phi(-x)}{\Phi(-x)} > x$ for all $x > 0$. As $\frac{\phi(x)}{\Phi(x)}$ is the inverse Mills' ratio,

$$\frac{\phi(-x)}{\Phi(-x)} = \frac{\phi(x)}{1 - \Phi(x)} = \mathbb{E}[X|X > x] > x, \forall x > 0.$$

By employing a symmetric argument, we find that: $\frac{F'_\varepsilon(x)^2}{F_\varepsilon(x) - 1} < F''_\varepsilon(x)$.

Assumption 17 guarantees that the average number of treated units in the training data is non-zero for any network size. Building on Eq.3.30, the following Lemma uniformly characterizes the relationship between the sampling uncertainty of $\hat{\theta}$ and the sampling uncertainty inherent in the empirical process $\mathbb{G}_n(\cdot)$. Formally:

Lemma 3.5.2. (Sampling Uncertainty) Under Assumption 12, 15, 16, and 17, we have

$$\mathbb{E}_{\varepsilon^n} [\|\hat{\theta} - \theta_0\|_1 | S, \sigma^{data}] \leq C_2 \mathbb{E}_{\varepsilon^n} [\|\nabla_{\theta} \mathbb{G}_n(\tilde{\theta})\|_1 | S, \sigma^{data}],$$

where C_2 is a constant that depends on the distribution F_{ε^n} , and the dimension and supports of the parameter space, the covariates space \mathcal{X} , the network space \mathcal{G} and the treatment allocation space \mathcal{D} .

Proof of Lemma 3.5.2 is provided in Appendix A.2.3, where we establish a uniform upper

bound for the largest eigenvalues (i.e., η_{max}^0 and η_{max}^1), termed the maximal largest eigenvalue. We show that this value is strictly negative and is encapsulated within the constant C_2 . As C_2 in Lemma 3.5.2 is a constant, characterising the concentration of the empirical process $\nabla_{\theta}\mathbb{G}_n(\tilde{\theta})$ is sufficient. The next section does so.

3.5.2 Finite Sample Result

Recall we are using a two-step ML estimation procedure, so the first step of estimation introduces additional sampling uncertainty through $\hat{\sigma}^{data}$. We incorporate these two layers of sampling uncertainty in the following lemma:

Lemma 3.5.3. *Under Assumption 12 to 16, we have*

$$\mathbb{E}_{\varepsilon^n} \left[\sup_{\theta \in \Theta} \|\nabla_{\theta}\mathbb{G}_n(\theta)\|_1 \middle| S, \sigma^{data} \right] \leq \frac{C_3 + C_4 \sqrt{\log(n)}}{\sqrt{n}},$$

where C_3 and C_4 are constants that depend only on the support of covariates, the distribution of ε , C_{σ} , the covariates space \mathcal{X} , the network space \mathcal{G} and the treatment allocation space \mathcal{D} .

A proof of Lemma 3.5.3 is provided in Appendix A.2.4. Lemma 3.5.3 analyzes the finite sample property of the empirical process (i.e., $\mathbb{E}_{\varepsilon^n} [\|\nabla_{\theta}\mathbb{G}_n(\theta)\|_1 | S, \sigma^{data}]$). By combining the results of Lemma 3.5.1, Lemma 3.5.2 and Lemma 3.5.3, we have our main theorem. This theorem characterizes the sampling uncertainty of using the empirical welfare:

Theorem 3.5.1. (Sampling Uncertainty of Regret) *Under Assumption 12 to 17, the sampling uncertainty of the two-step MLE estimator is bounded by:*

$$\mathbb{E}_{\varepsilon^n} [\|\hat{\theta} - \theta_0\|_1 | S, \sigma^{data}] \leq C_2 \frac{C_3 + C_4 \log(n)}{\sqrt{n}}$$

In addition, the sampling uncertainty of the empirical welfare is bounded by:

$$\mathbb{E}_{\varepsilon^n} \left[\max_{D \in \mathcal{D}} |W_n(D) - W(D)| \middle| S, \sigma^{data} \right] \leq C_1 C_2 \frac{C_3 + C_4 \log(n)}{\sqrt{n}}.$$

A proof of Theorem 3.5.1 is provided in Appendix A.3.2. This new result characterizes the finite sample properties of the sampling uncertainty that emerges when utilizing empirical welfare in settings of strategic interaction. It shows that the regret associated with empirical welfare converges at a rate influenced by the size of the network in the training data, as well as by the covariates and the chosen distribution for private information. Furthermore, this result characterizes the performance of the two-step maximum likelihood estimation from a finite sample perspective. This analysis can be extended to general M-estimators, including the Generalized Method of Moments and broader MLE frameworks.

3.5.3 Regret due to our Greedy Algorithm

Now, we evaluate the second term of Eq.3.20, which is the regret introduced by our greedy algorithm,

$$\mathbb{E}_{\varepsilon^n} \left[W_n(\tilde{D}) - W_n(\hat{D}_G) \middle| S, \sigma^{data} \right].$$

In general, the gap between a greedy optimizer and the global optimizer in terms of the value of the objective function is unknown. For monotone non-decreasing submodular set functions, Nemhauser et al. (1978) shows that a greedy algorithm achieves results within $(1 - 1/e)$ of the global maximum value. Although our optimization problem does not involve a submodular function, our empirical findings in Section 3.6 indicate that our greedy algorithm performs well, a result echoed in other applications such as experimental design (Lawrence et al., 2002). Building on these findings, Bian et al. (2017) provides a theoretical performance guarantee for using a greedy algorithm on non-submodular functions by leveraging the submodularity ratio and curvature of the objective function.

Submodularity, the submodularity ratio, and the curvature of a set function f are defined as follows.

Definition 3.5.1. (Submodularity) A set function is a submodular function if:

$$\sum_{k \in R \setminus S} [f(S \cup \{k\}) - f(S)] \geq f(S \cup R) - f(S), \quad \forall S, R \subseteq \mathcal{N}.$$

Definition 3.5.2. (Submodularity Ratio) The submodularity ratio of a non-negative set function $f(\cdot)$ is the largest γ such that

$$\sum_{k \in R \setminus S} [f(S \cup \{k\}) - f(S)] \geq \gamma [f(S \cup R) - f(S)], \quad \forall S, R \subseteq \mathcal{N}.$$

Definition 3.5.3. (Curvature) The curvature of a non-negative set function $f(\cdot)$ is the smallest value of ξ such that

$$f(R \cup \{k\}) - f(R) \geq (1 - \xi)[f(S \cup \{k\}) - f(S)], \quad \forall S \subseteq R \subseteq \mathcal{N}, \forall k \in \mathcal{N} \setminus R.$$

Submodularity is similar to diminishing returns. It states that adding an element to a smaller set yields a greater benefit than adding it to a larger set. [Lovász \(1983\)](#) highlights that, in discrete optimization, submodularity plays a role analogous to convexity in continuous optimization. The submodularity ratio measures how close a set function is to being submodular ([Das and Kempe, 2011](#)). Curvature quantifies the extent to which a set function deviates from being additive.

We evaluate the theoretical performance of our greedy algorithm in scenarios where the treatment exerts both direct and indirect positive effects on equilibrium welfare, as indicated by positive values for $(\hat{\theta}_1 + X_i^\top \hat{\theta}_3)$ and $\hat{\theta}_5$. Additionally, our empirical analysis explores a variety of other scenarios, including those with a negative direct effect but a positive indirect effect, among others. The results indicate that the algorithm performs well across a range of conditions.

To characterize the submodularity ratio and curvature of the objective function, we first represent it as a set function, which is a real-valued mapping defined over treatment alloca-

tions sets, $\mathcal{D} \subset \mathcal{N}$ (i.e., $\mathcal{D} = \{i \in \mathcal{N} : D_i = 1\}$):

$$W_n(\mathcal{D}) = \sum_{i \in \mathcal{D}} F_\varepsilon \left[\hat{\theta}_0 + \hat{\theta}_1 + X_i^\top (\hat{\theta}_2 + \hat{\theta}_3) + \frac{1}{|\mathcal{N}_i|} \sum_{j \in \mathcal{D} \setminus \{i\}} (\hat{\theta}_4 + \hat{\theta}_6 \underline{\sigma}_j) m_{ij} G_{ij} + \frac{\hat{\theta}_5}{|\mathcal{N}_i|} \sum_{j \in \mathcal{N} \setminus \{i\}} m_{ij} G_{ij} \underline{\sigma}_j \right] \\ + \sum_{k \in \mathcal{N} \setminus \mathcal{D}} F_\varepsilon \left[\hat{\theta}_0 + X_k^\top \hat{\theta}_2 + \frac{\hat{\theta}_4}{|\mathcal{N}_k|} \sum_{\ell \in \mathcal{D}} m_{k\ell} G_{k\ell} + \frac{\hat{\theta}_5}{|\mathcal{N}_k|} \sum_{\ell \in \mathcal{N} \setminus \{k\}} m_{k\ell} G_{k\ell} \underline{\sigma}_\ell \right].$$

Let γ denote the submodularity ratio. For a nondecreasing function, γ ranges between $[0, 1]$ and is 1 if and only if the function is submodular. Similarly, the curvature, denoted by ξ , of a nondecreasing function ranges between $[0, 1]$, and is 0 if and only if the function is supermodular. As our objective function involves a system of simultaneous equations, evaluating its curvature and submodularity ratio directly is challenging. Instead, we focus on the upper bound for curvature and submodularity, and ensure that their values remain within $(0, 1)$. Combining this result with [Bian et al. \(2017, Theorem 1\)](#) of leads to:

Proposition 3.5.1. *Under Assumption 12 and Assumption 14, the curvature ξ of $W_n(\mathcal{D})$ and the submodularity ratio γ of $W_n(\mathcal{D})$ are in $(0, 1)$. The greedy algorithm enjoys the following approximation guarantee for the problem in Eq.3.19:*

$$W_n(D_G) \geq \frac{1}{\xi} (1 - e^{-\xi\gamma}) W_n(\tilde{D}),$$

where D_G is the treatment assignment rule that is obtained by Algorithm 5.

A proof is provided in Appendix A.3.3. This proof is similar to [Kitagawa and Wang \(2023a\)](#). The first part of Proposition 3.5.1 implies that the performance guarantee is a non-trivial bound. Although the curvature and submodularity ratio of our objective function are unknown, for a particular application, it is possible to evaluate them empirically. As a consequence,

$$\mathbb{E}_{\varepsilon^n} \left[W_n(\tilde{D}) - W_n(\hat{D}_G) \middle| S, \sigma^{data} \right] \leq \mathcal{O}(1) \left(1 - \frac{1}{\xi} (1 - e^{-\xi\gamma}) \right), \quad (3.31)$$

where $\mathcal{O}(1)$ captures the $\mathbb{E}_{\varepsilon^n}[W_n(\tilde{D})|S, \sigma^{data}]$. Combining Eq.3.31 with Theorem 3.5.1, we obtain our main theorem:

Theorem 3.5.2. (Regret Bound) *Let D^* denote the maximizer of $W(D)$ and D_G be the assignment vector obtained by Algorithm 5. Under Assumptions 12 to 16, given curvature ξ and submodularity ratio γ , the regret is bounded from above by:*

$$\mathbb{E}_{\varepsilon^n} \left[R(\hat{D}_G) | S, \sigma^{data} \right] \leq C_1 C_2 \frac{C_3 + C_4 \log(n)}{\sqrt{n}} + \mathcal{O}(1) \left(1 - \frac{1}{\xi} (1 - e^{-\xi\gamma}) \right). \quad (3.32)$$

Theorem 3.5.2 is our key result. The first term in Eq.3.32 characterizes the sampling uncertainty, whose convergence rate depends on the network size. The dependence upon the parameters in the utility function, network structure, and private information distribution are shown implicitly via the terms C_1 in Lemma 3.5.1, C_2 in Lemma 3.5.2, and C_3 and C_4 in Lemma 3.5.3. The second term comes from the use of a greedy algorithm, and converges to a constant.

3.6 Empirical Application

We illustrate our proposed method using data from Banerjee et al. (2013), which explores the impact of information provision on microfinance adoption. Banerjee et al. (2013) studies a microfinance loan program. This program was introduced by Bharatha Swamukti Samsthe (BSS), a non-governmental microfinance institution in India, and implemented across 43 villages in Karnataka. BSS invited influential units, such as teachers, leaders of self-help groups, and shopkeepers, to an informational meeting about the availability of microfinance (the treatment). In total, 1262 units were assigned treatment, an average of 25.75 per village. After the intervention, researchers collected data on the network structure and household characteristics—including access to electricity, latrine quality, and per capita counts of beds and rooms in all participating villages. The number of households in each village varied from

107 to 341, with 10 to 51 households per village receiving information about the program. The program commenced in 2007, and the survey of microfinance adoption was completed by early 2011. We treat each household's decision to purchase microfinance as an equilibrium outcome within a simultaneous decision network game.

We consider each village as a distinct target population. Structural parameters in the payoff function for each village are estimated separately using the two-step maximum likelihood estimation (MLE) method of [Leung \(2015\)](#). This setup assumes that the training data, which includes several villages, acts as a representative sample, with each village in the training dataset mirroring a corresponding village in the target population in terms of covariates and network structure.

In the first stage of our analysis, following [Arcidiacono and Miller \(2011\)](#), we estimate the conditional choice probability using a flexible Logit approach (Chi-square goodness of fit test result is provided in Appendix [A.1](#)). This includes each unit's covariates and their second powers, as well as the covariates of directly linked neighbors and interactions among these covariates, which is under the network decaying dependence assumption ([Xu, 2018](#)). Although our estimator allows for incorporating covariates from neighbors at higher levels of linkage, we focus on directly linked neighbors' covariates in this estimation. We treat these estimates as the true parameters and assess the presence of strategic complementarity in each village. We find strategic complementarities in 16 of the 43 villages in the dataset¹⁴, which are the focus of this exercise. We assume the policymaker utilizes all available covariates to determine the treatment allocation mechanism. We assume that the private information follows a logistic distribution, and we define the measure of closeness between units i and j to be $m(X_i, X_j) = \frac{1}{1+|X_i - X_j|}$.

¹⁴The indices of those 16 villages in the original data set are: 1, 4, 6, 7, 12, 14, 17, 18, 20, 24, 25, 29, 31, 39, 40, and 41. To enhance clarity, we discard their original indices and re-label them as villages 1 to 16.

3.6.1 Policy evaluation

In this application, the objective is to maximize engagement welfare, measured as the microfinance participation rate, evaluated at the minimal equilibrium under a treatment capacity constraint (as in Eq.3.13) within our target population. To ensure comparability with the original study, we set the capacity constraint equal to the number of treatments used by Bharatha Swamukti Samsthe (BSS). We compare our method ('Robust') with two different treatment allocation regimes: the allocation rule adopted by BSS in the original study ('Original'), and a random allocation rule ('Random').

Table 3.1 presents predicted village-level microfinance take-up probabilities under three different treatment allocations. For each allocation rule, we report both the upper and lower bounds of the prediction set. The first column lists the 16 villages that exhibit strategic complementarities. The second column (*Sample Avg.*) contains the empirical average take-up rate for these villages. The third column (*Welfare under Original*) shows the average adoption rates for the original treatment allocation used by BSS. **Four villages—Villages 2, 3, 6, and 12—exhibit multiple equilibria under this rule.** To further assess our proposed method's performance, we generate 500 random treatment allocations within the capacity constraint for each village. The average purchasing probability across these 500 allocations is reported in the fourth column (*Welfare under Random*). Under random allocation, **multiple equilibria arise in Villages 2, and 9, highlighting that the occurrence of multiple equilibria can vary with the allocation method used.** The share of households adopting microfinance according to the robust optimal treatment allocation is shown as *Welfare under Robust*, **where the multiple equilibria only presents in the Village 2.**

Note first that the equilibrium average share of households adopting microfinance under the original allocation closely tracks the observed data for all villages except for Village 7 ¹⁵. Second, we find that the equilibrium average share of households purchasing microfinance

¹⁵The goodness of fit test indicates that the estimation for Village 7, referred to as Village 17 in Table A.2, may not adequately fit the data, potentially due to inaccuracies in the first-stage Conditional Choice Probability (CCP) estimation.

<i>Village</i>	<i>Sample Avg.</i>	<i>Original</i>	<i>Welfare under Random</i>	<i>Robust</i>	<i>Welfare Gain* Level</i>	<i>Percentage</i>
1	0.24	[0.25, 0.25]	[0.20, 0.20]	[0.41, 0.41]	0.16	66%
2	0.08	[0.04, 0.07]	[0.03, 0.04]	[0.13, 0.16]	0.10	270%
3	0.18	[0.17, 0.23]	[0.25, 0.25]	[0.37, 0.37]	0.20	122%
4	0.30	[0.26, 0.26]	[0.29, 0.29]	[0.44, 0.44]	0.18	68%
5	0.15	[0.15, 0.15]	[0.16, 0.16]	[0.37, 0.37]	0.22	146%
6	0.17	[0.15, 0.19]	[0.17, 0.17]	[0.39, 0.39]	0.24	158%
7	0.19	[0.48, 0.48]	[0.40, 0.40]	[0.66, 0.66]	0.18	37%
8	0.19	[0.17, 0.17]	[0.18, 0.18]	[0.28, 0.28]	0.11	64%
9	0.19	[0.19, 0.19]	[0.20, 0.23]	[0.33, 0.33]	0.14	72%
10	0.24	[0.24, 0.24]	[0.23, 0.23]	[0.28, 0.28]	0.05	20%
11	0.23	[0.22, 0.22]	[0.22, 0.22]	[0.34, 0.34]	0.12	57%
12	0.10	[0.09, 0.10]	[0.11, 0.11]	[0.30, 0.30]	0.20	223%
13	0.15	[0.13, 0.13]	[0.15, 0.15]	[0.45, 0.45]	0.31	238%
14	0.21	[0.23, 0.23]	[0.19, 0.19]	[0.46, 0.46]	0.23	101%
15	0.16	[0.18, 0.18]	[0.18, 0.18]	[0.41, 0.41]	0.24	132%
16	0.16	[0.16, 0.16]	[0.19, 0.19]	[0.29, 0.29]	0.13	85%

Table 3.1: Comparison using 16 Indian villages microfinance data from [Banerjee et al. \(2013\)](#)

* Minimal Welfare Gain compares the minimal simulated equilibrium welfare under the Robust allocation method and the simulated equilibrium welfare under the Original allocation implemented by BSS.

under random allocation is similar to the original BSS allocation method. When comparing our robust optimal treatment allocation regime with the original allocation rule, our method consistently outperforms the original rule in terms of both minimal and maximal equilibrium welfare. As depicted in Figure 3.1, improvements in welfare with minimal equilibrium vary from 20% to 270%. Notice that the welfare at the minimal equilibrium of our approach surpasses the maximal welfare under the other two approaches. This suggests that the information diffusion facilitated by the original treatment may not have significantly impacted adoption rates. Additionally, [Wang et al. \(2024\)](#) finds that households with higher centrality, such as the leaders selected by BSS, tend to have a lower borrowing probability compared to less central households. It is possible that more central households have greater access to alternative borrowing sources within their networks, thus diminishing their need for microfinance, and reducing the spillover effects through strategic interactions.




Figure 3.1: Comparisons between three approaches

3.7 Extension

3.7.1 Complete Information Game

In a complete information setting, units observe all the characteristics of other units participating in the game. This means that units are informed of others' choices before making their own decisions, allowing them to play the best response to the observed actions rather than basing their actions on beliefs, as is common in a private information setting. As a consequence, unit i 's decision rule is:

$$Y_i = \mathbb{1}\left\{U_i(1, Y_{-i}, X, D, G) \geq 0\right\}, \quad \forall i \in \mathcal{N}.$$

One main distinction from incomplete information settings is that the solution concept transitions to a pure-strategy Nash equilibrium. A pure-strategy Nash equilibrium is defined by a

set of actions $y^* = \{y_1^*, \dots, y_N^*\}$ such that

$$U_i(y_i^*, y_{-i}^*, X, D, G) \geq U_i(y_i', y_{-i}^*, X, D, G)$$

for any $y_i' \in \mathcal{Y}$ and for all $i \in \mathcal{N}$. We denote the set of all such equilibria as $\Sigma(X, D, G, \varepsilon) := \{y^*\}$, given covariates X , treatment allocation D , network structure G , and the idiosyncratic shock ε . To simplify the notation, we subsequently refer to it as $\Sigma(\varepsilon)$. Let $\xi : \Sigma \rightarrow [0, 1]$ denote the probability distribution over equilibria, and let $\Delta(\Sigma) := \{\xi : \sum_{y^* \in \Sigma} \xi(y^*) = 1\}$ denote the set of all the probability distributions.

In scenarios with strategic complementarity, there exists a maximal and a minimal Nash equilibrium, denoted by \bar{y}^* and \underline{y}^* . For our counterfactual analysis, which is analogous to the framework established in Theorem 3.3.1 under an incomplete information setting, we propose the following:

Proposition 3.7.1. *For a supermodular game, the least favorable equilibrium selection rule $\underline{\lambda}$ and the most favorable equilibrium selection rule $\bar{\lambda}$ are:*

$$\underline{\lambda} := \delta_{\underline{y}^*}, \quad \bar{\lambda} := \delta_{\bar{y}^*},$$

where δ_y is the Dirac measure on Σ . In addition, the following conditions are satisfied:

$$\inf_{\lambda \in \Lambda} \sum_{i=1}^N \Pr(Y_i = 1 | X, D, G, \lambda) = \sum_{i=1}^N \inf_{\lambda \in \Lambda} \Pr(Y_i = 1 | X, D, G, \lambda),$$

$$\sup_{\lambda \in \Lambda} \sum_{i=1}^N \Pr(Y_i = 1 | X, D, G, \lambda) = \sum_{i=1}^N \sup_{\lambda \in \Lambda} \Pr(Y_i = 1 | X, D, G, \lambda).$$

The proof of Proposition 3.7.1 mirrors that of Theorem 3.3.1, with the primary modification being the substitution of Bayesian Nash equilibrium with Nash equilibrium. Computing the conditional choice probability differs from the previous analysis since it is no longer a simultaneous equation system. With complete information, the conditional choice probability

is given by:

$$\Pr(Y_i = 1|X, D, G, \underline{\lambda}) = \int 1 \{ \exists y_{-i} : (1, y_{-i}) \in \Sigma(\varepsilon) \text{ and } \forall y_{-i}, (0, y_{-i}) \notin \Sigma(\varepsilon) \} dF_\varepsilon$$

$$\Pr(Y_i = 1|X, D, G, \bar{\lambda}) = \int 1 \{ \exists y_{-i} : (1, y_{-i}) \in \Sigma(\varepsilon) \} dF_\varepsilon$$

The above expression is hard to compute. To further simplify the computation, let us define event $A := \{ \exists y_{-i} : (1, y_{-i}) \in \Sigma(\varepsilon) \}$ and event $B := \{ \exists y_{-i}, (0, y_{-i}) \in \Sigma(\varepsilon) \}$. Given $\Pr(A \cap B^c|X, D, G) = \Pr(A \cup B|X, D, G) - \Pr(B|X, D, G)$, we hence have:

$$\Pr(Y_i = 1|X, D, G, \underline{\lambda}) = 1 - \Pr(Y_i = 0|X, D, G, \bar{\lambda}).$$

As a consequence, it is enough to compute $\Pr(Y_i = 1|X, D, G, \bar{\lambda})$ and $\Pr(Y_i = 0|X, D, G, \bar{\lambda})$, which are given by:

$$\Pr(Y_i = 1|X, D, G, \bar{\lambda}) = \int 1 \left\{ \max_{y_{-i} : (1, y_{-i}) \in \Sigma(\varepsilon)} U_i(1, y_{-i}, X, D, G) \geq 0 \right\} dF_\varepsilon, \quad (3.33)$$

$$\Pr(Y_i = 0|X, D, G, \bar{\lambda}) = \int 1 \left\{ \min_{y_{-i} : (1, y_{-i}) \in \Sigma(\varepsilon)} U_i(1, y_{-i}, X, D, G) < 0 \right\} dF_\varepsilon. \quad (3.34)$$

Let us define i^C as the complement of unit i and their neighbor set, i.e., $i^C := \mathcal{N} \setminus (\mathcal{N}_i \cup \{i\})$. Define $y_{\mathcal{N}_i}$ as the collection of neighbors' choices of unit i . Consequently, y_{-i} can be expressed as $(y_{\mathcal{N}_i}, y_{i^C})$, where y_{i^C} represents the choices of units in i^C . For the optimization problems defined in Eq.3.33 (maximization) and Eq.3.34 (minimization), it is necessary to explore all possible equilibria for each value of ε within the network game. Given our utility function specification, the choice of unit i depends only on $k \in i^C$ through the choices of units directly connected with i . Thus, we can simplify the maximization problem in Eq.3.33 to:

$$\max_{y_{\mathcal{N}_i}, y_{i^C}} U_i(1, y_{\mathcal{N}_i}, X, D, G)$$

with constraints:

$$y_j = 1\{U_j(1, y_{\mathcal{N}_j \setminus \{i\}}, X, D, G) \geq 0\}, \forall j \in \mathcal{N}_i, \quad (3.35)$$

$$y_k = 1\{U_k(1, y_{\mathcal{N}_k}, X, D, G) \geq 0\}, \forall k \in i^C. \quad (3.36)$$

These constraints (Eq.3.35 and Eq.3.36) ensure that $(1, y_{\mathcal{N}_i}, y_{i^C})$ forms a Nash equilibrium. In the optimization, $U_i(1, y_{\mathcal{N}_i}, X, D, G)$ does not depend directly on y_{i^C} but needs to confirm that $(1, y_{\mathcal{N}_i}, y_{i^C})$ is a Nash equilibrium for any given $y_{\mathcal{N}_i}$. If for some $y_{\mathcal{N}_i}$, multiple y_{i^C} ensure y_{-i} as a Nash equilibrium, the existence of any y_{i^C} that satisfies this condition is sufficient for our purposes. We search for $y_{\mathcal{N}_i} \in \mathcal{Y}^{|\mathcal{N}_i|}$ that maximizes $U_i(y_{\mathcal{N}_i}, X, D, G)$, denoted as $y_{\mathcal{N}_i}^*$. Given the supermodular nature of our game, where neighbors' choices are strategic complements to unit j 's choice, we select y_{i^C} such that $(y_{\mathcal{N}_i}, y_{i^C})$ constitutes the largest Nash equilibrium for the given $y_{\mathcal{N}_i}$, leveraging the increasing monotonicity between $y_{\mathcal{N}_i}$ and y_{i^C} . We then search the $y_{\mathcal{N}_i}$ that maximizes the objective function.

3.8 Conclusion

This paper proposes a method for constructing individualized treatment allocations to maximize equilibrium welfare robust to the presence of multiple equilibria in large simultaneous decision games with complementarity. Our approach, takes into account the inherent complexity introduced by the presence of multiple Nash equilibria, and the resulting incompleteness. We refrain from making assumptions about the equilibrium selection mechanism, which leads to both analytical and numerical challenges in evaluating counterfactual equilibrium welfare. Due to the inherent uncertainties in our model, we use the maximin welfare criterion to evaluate treatment allocation rules. This leads to treatment allocation rules that are optimized to maximize the worst-case equilibrium social welfare, ensuring their robustness. The use of a greedy optimization algorithm further enhances the applicability of our approach.

We acknowledge that several questions remain open, and there are multiple ways in which our work can be extended. First, we have not explored counterfactual analysis within the broader framework of general simultaneous decision games. Second, although we parametrize the utility function and the distribution of idiosyncratic shock in this work, adopting a non-parametric utility function and a non-parametric distribution of idiosyncratic shock could significantly enhance the robustness and applicability of our approach. Third, while we have assumed independence among idiosyncratic shocks, recent literature, such as [Grieco \(2014\)](#) and [de Paula and Tang \(2020\)](#), have begun to relax this assumption, suggesting another avenue for refining our model.

Supplement to “Robust Network Targeting with Multiple Nash Equilibria”

A.1 Chi-Square Goodness of Fit Test

Number of Rooms					Number of Rooms				
Village	1-2	3-4	5-6	≥ 7	Village	1-2	3-4	5-6	≥ 7
1	0.08	0.04	0.61	0.38	23	0.06	0.22	3.27	0.00
2	92.90	17.03	0.13	–	24	0.31	0.19	0.60	–
3	0.75	0.14	0.53	–	25	0.06	0.86	1.52	–
4	4.81	0.50	0.62	0.00	26	0.04	0.14	0.15	–
5	0.02	0.02	2.08	0.15	27	0.20	0.13	0.80	0.01
6	0.35	0.15	0.02	0.04	28	345.59	115.88	6.39	1.76
7	0.82	0.45	0.74	0.00	29	0.02	0.01	1.03	0.08
8	0.07	0.10	0.00	0.45	30	2.42	1.02	0.00	0.00
9	1.76	0.81	0.05	–	31	0.18	0.97	0.19	–
10	0.01	0.02	1.38	0.01	32	0.76	0.36	0.00	0.87
11	0.09	0.03	0.34	0.10	33	1.61	0.12	0.72	1.55
12	0.00	1.16	1.47	0.35	34	0.12	0.31	0.11	–
13	1.67	0.06	0.26	3.79	35	0.21	0.45	0.45	0.09
14	0.01	0.10	0.92	0.00	36	0.11	0.80	0.68	0.93
15	0.05	0.00	0.58	1.11	37	0.04	0.03	0.17	0.03
16	0.03	0.04	0.24	0.01	38	0.00	0.09	0.32	–
17	80.36	7.58	0.37	0.34	39	0.35	0.03	0.09	–
18	2.20	1.76	0.73	5.54	40	0.43	0.02	0.03	–
19	0.00	2.02	0.60	0.24	41	0.01	0.00	0.07	0.00
20	0.04	0.23	0.01	0.89	42	49.43	10.70	1.26	0.00
21	0.04	0.00	0.25	0.14	43	0.82	0.00	0.58	0.03
22	0.42	0.93	2.04	0.05					

Table A.2: Chi-square values based on number of rooms for 43 Indian villages microfinance data from [Banerjee et al. \(2013\)](#)

At 0.05 significance level, the critical value is given by 7.815.

‘–’ symbol indicates that there are no households with more than 7 rooms.

A.2 Lemmas

We first introduce notation. We define $\bar{m} := \max_{ij} |m_{ij}|$, and $\underline{F}_\varepsilon := \min_{\substack{\theta \in \Theta \\ z \in \mathcal{Z}}} F_\varepsilon(z^\top \theta)$. In addition, we define $v := \min\{\underline{F}_\varepsilon, 1 - \bar{F}_\varepsilon\}$. We measure the distance between parameters θ with the

L_1 metric, which we denote by $\|\theta - \theta'\|_1 := \sum_{k=1}^{d_\theta} |\theta_k - \theta'_k|$. For a $K \times L$ matrix A , $\|A\|_\infty$ denotes the operator norm of A induced by the L_∞ norm, which is given as: $\|A\|_\infty = \max_{k=1, \dots, K} \sum_{l=1}^L |A_{kl}|$.

A.2.1 Proof of Corollary 3.3.1

Corollary A.2.1. (Utilitarian Welfare at Equilibrium) Under Assumption 12, given the specification of our utility function, the predicted set of the expected utilitarian welfare under a counterfactual policy D is given as:

$$W_{X,G,\lambda}(D) \in \left[\frac{1}{N} \sum_{i=1}^N \alpha_i f(\alpha_i) + \frac{1}{N} \sum_{i=1}^N \sum_{j \neq i} \beta_{ij} \underline{\sigma}_i^* \underline{\sigma}_j^*, \frac{1}{N} \sum_{i=1}^N \alpha_i f(\alpha_i) + \frac{1}{N} \sum_{i=1}^N \sum_{j \neq i} \beta_{ij} \bar{\sigma}_i^* \bar{\sigma}_j^* \right],$$

where

$$f(\alpha_i) = \begin{cases} \Pr(Y_i = 1 | X, D, G, \underline{\lambda}) & \text{if } \alpha_i > 0 \\ \Pr(Y_i = 1 | X, D, G, \bar{\lambda}) & \text{if } \alpha_i \leq 0. \end{cases}$$

Proof. Given

$$W_{X,G,\lambda}(D) = \frac{1}{N} \sum_{i=1}^N \mathbb{E} [U_i(Y, X, D, G) - \varepsilon_i Y_i | X, D, G, \lambda],$$

we have,

$$W_{X,G,\lambda}(D) = \frac{1}{N} \sum_{i=1}^N \alpha_i \Pr(Y_i = 1 | X, D, G, \lambda) + \frac{1}{N} \sum_{i=1}^N \sum_{j \neq i} \beta_{ij} \Pr(Y_i Y_j = 1 | X, D, G, \lambda). \quad (\text{A.37})$$

Therefore, if $\alpha_i > 0$, $\alpha_i \Pr(Y_i = 1 | X, D, G, \lambda)$ achieves its lower bound by choosing the equilibrium selection rule $\underline{\lambda}$. When $\alpha_i \leq 0$, $\alpha_i \Pr(Y_i = 1 | X, D, G, \lambda)$ achieves its lower bound by choosing the equilibrium selection rule $\bar{\lambda}$. For the second term, since $\beta_{ij} \geq 0$ for all $i, j \in \mathcal{N}$, it achieves its upper bound by choosing the equilibrium selection rule as $\bar{\lambda}$ and it achieves its lower bound by choosing the equilibrium selection rule as $\underline{\lambda}$. This is because

$$\begin{aligned} & \Pr(Y_i Y_j = 1 | X, D, G, \lambda) \\ &= \sum_{\sigma^* \in \Sigma} \lambda(\sigma^* | X, D, G) \int \mathbb{1} \left\{ \alpha_i + \sum_{k \neq i} \beta_{ik} \sigma_k^*(X, D, G) \geq \varepsilon_i \right\} \mathbb{1} \left\{ \alpha_j + \sum_{k \neq j} \beta_{jk} \sigma_k^*(X, D, G) \geq \varepsilon_j \right\} dF_{\varepsilon_i} dF_{\varepsilon_j} \\ &= \sum_{\sigma^* \in \Sigma} \lambda(\sigma^* | X, D, G) \int \mathbb{1} \left\{ \alpha_i + \sum_{k \neq i} \beta_{ik} \sigma_k^*(X, D, G) \geq \varepsilon_i \right\} dF_{\varepsilon} \int \mathbb{1} \left\{ \alpha_j + \sum_{k \neq j} \beta_{jk} \sigma_k^*(X, D, G) \geq \varepsilon_j \right\} dF_{\varepsilon}, \end{aligned}$$

where the second equality holds by Assumption 12. Therefore,

$$\Pr(Y_i Y_j = 1 | X, D, G, \lambda) = \sum_{\sigma^* \in \Sigma} \lambda(\sigma^* | X, D, G) \Pr(Y_i = 1 | X, D, G, \sigma^*) \Pr(Y_j = 1 | X, D, G, \sigma^*).$$

From Theorem 3.3.1, $\{\Pr(Y_j = 1|X, D, G, \sigma^*)\}_{j=1}^N$ achieves their upper bound under the most favorable equilibrium selection rule $\bar{\lambda}$, where $\bar{\sigma}^*$ happens with probability 1. Therefore,

$$\begin{aligned}\Pr(Y_i Y_j = 1|X, D, G, \bar{\lambda}) &= \Pr(Y_i = 1|X, D, G, \bar{\sigma}^*) \Pr(Y_j = 1|X, D, G, \bar{\sigma}^*) \\ &\geq \Pr(Y_i Y_j = 1|X, D, G, \lambda), \quad \forall \lambda \in \Lambda.\end{aligned}$$

By the symmetric argument, we have

$$\begin{aligned}\Pr(Y_i Y_j = 1|X, D, G, \underline{\lambda}) &= \Pr(Y_i = 1|X, D, G, \underline{\sigma}^*) \Pr(Y_j = 1|X, D, G, \underline{\sigma}^*) \\ &\leq \Pr(Y_i Y_j = 1|X, D, G, \lambda), \quad \forall \lambda \in \Lambda.\end{aligned}$$

Therefore, for all $i, j \in \mathcal{N}$,

$$\beta_{ij} \Pr(Y_i Y_j = 1|X, D, G, \lambda) \in [\beta_{ij} \underline{\sigma}_i^* \underline{\sigma}_j^*, \beta_{ij} \bar{\sigma}_i^* \bar{\sigma}_j^*]. \quad (\text{A.38})$$

Plugging Eq.A.38 into Eq.A.37 completes the proof. \square

A.2.2 Proof of Lemma 3.5.1

Lemma A.2.1. Under Assumptions 12 and 14,

$$\mathbb{E}_{\varepsilon^n} \left[\max_{D \in \mathcal{D}} |W_n(D) - W(D)| \middle| S, \sigma^{data} \right] \leq C_1 \mathbb{E}_{\varepsilon^n} \left[\|\hat{\theta} - \theta\|_1 \middle| S, \sigma^{data} \right]$$

where C_1 is a constant that only depends on the distribution F_{ε^n} , the maximum link in the network \bar{N} , and the support of true parameter θ_0 .

Proof. By the Triangle inequality, an upper bound for our objective function is:

$$|W_n(D) - W(D)| = \left| \frac{1}{N} \sum_{i=1}^N (\underline{\sigma}_i(\hat{\theta}) - \underline{\sigma}_i(\theta)) \right| \leq \frac{1}{N} \sum_{i=1}^N |\underline{\sigma}_i(\hat{\theta}) - \underline{\sigma}_i(\theta)|.$$

We know:

$$\begin{aligned}\underline{\sigma}_i(\theta) &= \int \mathbb{1} \left\{ \alpha_i + \sum_{j \neq i} \beta_{ij} \underline{\sigma}_j(\theta) - \varepsilon_i \geq 0 \right\} dF(\varepsilon) \\ &= F_{\varepsilon} \left(\theta_0 + \theta_1 D_i + \theta_2^T X_i + \theta_3^T X_i D_i + \frac{\theta_4}{N_i} \sum_{j \neq i} G_{ij} m_{ij} D_j + \frac{1}{N_i} \sum_{j \neq i} (\theta_5 + \theta_6 D_i D_j) G_{ij} m_{ij} \underline{\sigma}_j(\theta) \right).\end{aligned}$$

Let $r(i, \theta) := \theta_0 + \theta_1 D_i + \theta_2^T X_i + \theta_3^T X_i D_i + \frac{\theta_4}{N_i} \sum_{j \neq i} G_{ij} m_{ij} D_j + \frac{1}{N_i} \sum_{j \neq i} (\theta_5 + \theta_6 D_i D_j) G_{ij} m_{ij} \underline{\sigma}_j(\theta)$. Therefore,

$$|\underline{\sigma}_i(\hat{\theta}) - \underline{\sigma}_i(\theta)| = |F_{\varepsilon}(r(i, \hat{\theta})) - F_{\varepsilon}(r(i, \theta))|.$$

By the Mean Value Theorem,

$$\begin{aligned}
|\underline{\sigma}_i(\hat{\theta}) - \underline{\sigma}_i(\theta)| &= |\nabla_{\theta} F_{\varepsilon}(r_i(\tilde{\theta}))(\hat{\theta} - \theta)| \\
&= |F'_{\varepsilon}(r(i, \tilde{\theta})) \nabla_{\theta} r(i, \tilde{\theta})(\hat{\theta} - \theta)| \\
&= F'_{\varepsilon}(r(i, \tilde{\theta})) |\nabla_{\theta} r(i, \tilde{\theta})(\hat{\theta} - \theta)| \\
&\leq \tau |\nabla_{\theta} r(i, \tilde{\theta})(\hat{\theta} - \theta)|.
\end{aligned}$$

For some $\tilde{\theta} \in \mathbb{R}^{d_{\theta}}$ on the segment from θ to $\hat{\theta}$. By the Cauchy–Schwarz inequality, we have:

$$|\underline{\sigma}_i(\hat{\theta}) - \underline{\sigma}_i(\theta)| \leq \tau \|\nabla_{\theta} r(i, \tilde{\theta})\|_2 \|\hat{\theta} - \theta\|_2 \leq \tau \|\nabla_{\theta} r(i, \tilde{\theta})\|_1 \|\hat{\theta} - \theta\|_1. \quad (\text{A.39})$$

To deal with the simultaneity within $\nabla_{\theta} r(i, \tilde{\theta})$, define

$$\nabla_{\theta} r(\theta) := \begin{bmatrix} \nabla_{\theta_0} r(\theta) & \cdots & \nabla_{\theta_{d_{\theta}}} r(\theta) \end{bmatrix},$$

where

$$\nabla_{\theta_k} r(\theta) = \begin{bmatrix} \nabla_{\theta_k} r(1, \theta) \\ \vdots \\ \nabla_{\theta_k} r(N, \theta) \end{bmatrix},$$

for all $k = 1, \dots, d_{\theta}$. Then,

$$\|\nabla_{\theta} r(i, \theta)\|_1 \leq \|\nabla_{\theta} r(\theta)\|_{\infty}, \quad \forall i \in \mathcal{N}, \quad (\text{A.40})$$

where recall $\|\nabla_{\theta} r(\theta)\|_{\infty} = \max_{i=1, \dots, N} \|\nabla_{\theta} r(i, \theta)\|_1$ is the operator norm induced by the L_{∞} norm. To bound $\|\nabla_{\theta} r(\theta)\|_{\infty}$, we define an implicit function $I : \mathbb{R}^N \times \mathbb{R}^{d_{\theta}} \rightarrow \mathbb{R}^N$ such that

$$I(r, \theta) = \begin{bmatrix} r(1, \theta) - a_1 - \frac{\theta_4}{|\mathcal{N}_1|} \sum_{j \neq 1} G_{1j} m_{1j} D_j - \frac{1}{|\mathcal{N}_1|} \sum_{j \neq 1} (\theta_5 + \theta_6 D_1 D_j) G_{1j} m_{1j} F_{\varepsilon}(r(j, \theta)) \\ \vdots \\ r(N, \theta) - a_N - \frac{\theta_4}{|\mathcal{N}_N|} \sum_{j \neq N} G_{Nj} m_{Nj} D_j - \frac{1}{|\mathcal{N}_N|} \sum_{j \neq N} (\theta_5 + \theta_6 D_N D_j) G_{Nj} m_{Nj} F_{\varepsilon}(r(j, \theta)) \end{bmatrix},$$

where $a_i = \theta_0 + \theta_1 D_i - \theta_2^{\top} X_i - \theta_3^{\top} X_i D_i$. By the Implicit Function Theorem, $\nabla_{\theta} r(\theta)$ is given by:

$$\nabla_{\theta} r(\theta) = -(\nabla_r I(r, \theta))^{-1} \nabla_{\theta} I(r, \theta), \quad (\text{A.41})$$

where $\nabla_r I(r, \theta)$ is:

$$\nabla_r I(r, \theta) = \begin{bmatrix} 1 & -\frac{1}{|\mathcal{N}_1|} F'_\varepsilon(r(1, \theta))(\theta_5 + \theta_6 D_1 D_2) G_{12} m_{12} & \cdots \\ -\frac{1}{|\mathcal{N}_2|} F'_\varepsilon(r(2, \theta))(\theta_5 + \theta_6 D_2 D_1) G_{21} m_{21} & 1 & \cdots \\ \vdots & \vdots & \ddots \\ -\frac{1}{|\mathcal{N}_N|} F'_\varepsilon(r(N, \theta))(\theta_5 + \theta_6 D_N D_1) G_{N1} m_{N1} & -\frac{1}{|\mathcal{N}_N|} F'_\varepsilon(r(N, \theta))(\theta_5 + \theta_6 D_N D_2) G_{N2} m_{N2} & \cdots \end{bmatrix}$$

and $\nabla_\theta I(r, \theta)$ is:

$$\nabla_\theta I(r, \theta) = \begin{bmatrix} -1 & -D_1 & -X_1^\top & -D_1 X_1^\top & -\frac{\sum_{j \neq 1} G_{1j} m_{1j} D_j}{|\mathcal{N}_1|} & -\frac{\sum_{j \neq 1} F'_\varepsilon(r(j, \theta)) G_{1j} m_{1j}}{|\mathcal{N}_1|} & -\frac{\sum_{j \neq 1} F'_\varepsilon(r(j, \theta)) D_1 D_j G_{1j} m_{1j}}{|\mathcal{N}_1|} \\ -1 & -D_2 & -X_2^\top & -D_2 X_2^\top & -\frac{\sum_{j \neq 2} G_{2j} m_{2j} D_j}{|\mathcal{N}_2|} & -\frac{\sum_{j \neq 2} F'_\varepsilon(r(j, \theta)) G_{2j} m_{2j}}{|\mathcal{N}_2|} & -\frac{\sum_{j \neq 2} F'_\varepsilon(r(j, \theta)) D_2 D_j G_{2j} m_{2j}}{|\mathcal{N}_2|} \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ -1 & -D_N & -X_N^\top & -D_N X_N^\top & -\frac{\sum_{j \neq N} G_{Nj} m_{Nj} D_j}{|\mathcal{N}_N|} & -\frac{\sum_{j \neq N} F'_\varepsilon(r(j, \theta)) G_{Nj} m_{Nj}}{|\mathcal{N}_N|} & -\frac{\sum_{j \neq N} F'_\varepsilon(r(j, \theta)) D_N D_j G_{Nj} m_{Nj}}{|\mathcal{N}_N|} \end{bmatrix}.$$

Therefore, supremum norm of Eq.A.41 is bounded by

$$\|\nabla_\theta r(\theta)\|_\infty = \|(\nabla_r I(r, \theta))^{-1} \nabla_\theta I(r, \theta)\|_\infty \leq \|(\nabla_r I(r, \theta))^{-1}\|_\infty \|\nabla_\theta I(r, \theta)\|_\infty. \quad (\text{A.42})$$

The last inequality holds because norm $\|A\|_\infty$ on matrix A is the operator norm induced by L_∞ norm, which is a matrix norm. This ensures that it satisfies the submultiplicativity property. Therefore, we have the inequality. Assuming $\nabla_r I(r, \tilde{\theta})$ is non-singular, and by invoking Lemma A.4.1 — a corollary of Berge's Maximum Theorem — the norm $\|(\nabla_r I(r, \tilde{\theta}))^{-1}\|_\infty$ is a continuous function with respect to the entries of $\nabla_r I(r, \tilde{\theta})$. Given, for all $i, j \in \mathcal{N}$, D_i and G_{ij} are binary, X_i has bounded support, θ is in a compact parameter space (Assumption 14), and $F'_\varepsilon(\cdot) \in [0, \tau]$, the Extreme Value Theorem (Lemma A.5.1) guarantees the existence of a uniform maximum of $\|(\nabla_r I(r, \tilde{\theta}))^{-1}\|_\infty$ among all the values of X, D, G , which only depends on the support of each variable. We denote this uniform maximum as ζ . For $\|\nabla_\theta I(r, \tilde{\theta})\|_\infty$ in Eq.A.42, we know:

$$\begin{aligned} \|\nabla_\theta I(r, \tilde{\theta})\|_\infty &= \max_{i \in \mathcal{N}} 1 + D_i + \|X_i\|_1 + D_i \|X_i\|_1 + \frac{1}{|\mathcal{N}_i|} \sum_{j \neq i} D_j G_{ij} |m_{ij}| \\ &\quad + \frac{1}{|\mathcal{N}_i|} \sum_{j \neq i} F'_\varepsilon(r(j, \tilde{\theta})) (1 + D_i D_j) G_{ij} |m_{ij}|. \end{aligned} \quad (\text{A.43})$$

Furthermore, Eq.A.43 is upper bounded by $2 + 2 \max_{i \in \mathcal{N}} \|X_i\|_1 + 2\bar{m}\tau$, where $\bar{m} := \max_{i, j \in \mathcal{N}} |m_{ij}|$.

Therefore, we have

$$\|\nabla_{\theta} I(r, \tilde{\theta})\|_{\infty} \leq 2 + 2\|X\|_{\infty} + \bar{m} + 2\bar{m}\tau. \quad (\text{A.44})$$

Combining Eq.A.39, Eq.A.40 and Eq.A.44, we have:

$$\frac{1}{N} \sum_{i=1}^N |\underline{\sigma}_i(\hat{\theta}) - \underline{\sigma}_i(\theta)| \leq \zeta\tau(2 + 2\|X\|_{\infty} + \bar{m} + 2\bar{m}\tau)\|\hat{\theta} - \theta\|_1.$$

To complete the proof, let $C_1 = \zeta\tau(2 + 2\|X\|_{\infty} + \bar{m} + 2\bar{m}\tau)$.

□

A.2.3 Proof of Lemma 3.5.2

Lemma A.2.2. (Sampling Uncertainty) Under Assumption 12, 15, 16, and 17, we have

$$\mathbb{E}_{\varepsilon^n} [\|\hat{\theta} - \theta_0\|_1 | S, \sigma^{data}] \leq C_2 \mathbb{E}_{\varepsilon^n} [\|\nabla_{\theta} \mathbb{G}_n(\tilde{\theta})\|_1 | S, \sigma^{data}],$$

where C_2 is a constant that depends on the distribution F_{ε^n} , and the dimension and supports of the parameter space, the covariates space \mathcal{X} , the network space \mathcal{G} and the treatment allocation space \mathcal{D} .

Proof. The Hessian matrices in the Taylor expansions of $\hat{\mathbb{M}}(\theta_0) - \hat{\mathbb{M}}(\hat{\theta})$ and $M(\hat{\theta}) - M(\theta_0)$ are:

$$\nabla_{\theta}^2 \hat{\mathbb{M}}(\hat{\theta}) = \frac{1}{n} \sum_{i=1}^n [Y_i \omega_0(\hat{Z}_i^{\top} \hat{\theta}) - (1 - Y_i) \omega_1(\hat{Z}_i^{\top} \hat{\theta})] \hat{Z}_i \hat{Z}_i^{\top},$$

$$\nabla_{\theta}^2 M(\hat{\theta}) = \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{\varepsilon^n} \left[[Y_i \omega_0(Z_i^{\top} \hat{\theta}) - (1 - Y_i) \omega_1(Z_i^{\top} \hat{\theta})] Z_i Z_i^{\top} \middle| S, \sigma^{data} \right],$$

where

$$\omega_0(a) = \frac{F_{\varepsilon}''(a)F_{\varepsilon}(a) - F_{\varepsilon}'(a)^2}{F_{\varepsilon}(a)^2}, \quad \omega_1(a) = \frac{F_{\varepsilon}''(a)[1 - F_{\varepsilon}(a)] + F_{\varepsilon}'(a)^2}{[1 - F_{\varepsilon}(a)]^2},$$

and $a \in \mathbb{R}$. Combining the above equations with Eq.3.23 and Eq.3.24, we have:

$$\hat{\mathbb{M}}(\theta_0) - \hat{\mathbb{M}}(\hat{\theta}) = \frac{1}{2n} \sum_{i=1}^n \left[Y_i \omega_0(\hat{Z}_i^{\top} \hat{\theta}) - (1 - Y_i) \omega_1(\hat{Z}_i^{\top} \hat{\theta}) \right] \left[(\hat{\theta} - \theta_0)^{\top} \hat{Z}_i \right]^2 \leq 0. \quad (\text{A.45})$$

$$M(\hat{\theta}) - M(\theta_0) = \frac{1}{2n} \sum_{i=1}^n \left[\sigma_i \omega_0(Z_i^{\top} \hat{\theta}) - (1 - \sigma_i) \omega_1(Z_i^{\top} \hat{\theta}) \right] \left[(\hat{\theta} - \theta_0)^{\top} Z_i \right]^2 \leq 0, \quad (\text{A.46})$$

where $\sigma_i = \mathbb{E}_{\varepsilon^n} [Y_i | S, \sigma^{data}]$. Although Eq.A.45 and Eq.A.46 are non-positive, they do not pin down the sign of the coefficient on the quadratic term (i.e., $Y_i \omega_0(\hat{Z}_i^{\top} \hat{\theta}) - (1 - Y_i) \omega_1(\hat{Z}_i^{\top} \hat{\theta})$ and $\sigma_i \omega_0(Z_i^{\top} \hat{\theta}) - (1 - \sigma_i) \omega_1(Z_i^{\top} \hat{\theta})$). Under Assumption 15, we have $\omega_0(a) < 0$ and $\omega_1(a) > 0$ for all $a \in \mathbb{R}$. Furthermore, Assumption 14 ensures that Θ is a compact parameter space, which

guarantees that $\hat{\theta}$ resides within this compact set. Consequently, $\hat{\theta}$ and $\hat{\theta}$ are also confined within the same compact set. Given this setup, all elements in Z_i and \hat{Z}_i for each $i = 1, \dots, n$ also exist within a compact set. Therefore, the products $\hat{Z}_i^\top \hat{\theta}$ and $Z_i^\top \hat{\theta}$ are confined within a compact set, denoted as Ξ . We define the following bounds for ω_0 and ω_1 across Ξ :

$$\bar{\omega}_0 := \max_{x \in \Xi} \omega_0(x), \quad \bar{\omega}_1 := \max_{x \in \Xi} -\omega_1(x), \quad \underline{\omega}_0 := \min_{x \in \Xi} \omega_0(x), \quad \underline{\omega}_1 := \min_{x \in \Xi} -\omega_1(x).$$

By Assumption 12 and 15, the probability density function of ε^n and its derivative are bounded. In addition, under Assumption 12 they are continuous functions. Therefore, the Extreme Value Theorem guarantees the existence of $\bar{\omega}_0$ and $\bar{\omega}_1$. Define $\bar{\omega}$ as $\max\{\bar{\omega}_0, \bar{\omega}_1\}$, and $\underline{\omega}$ as $\min\{\underline{\omega}_0, \underline{\omega}_1\}$. We have $\omega_0(\hat{Z}_i^\top \hat{\theta}) \leq \bar{\omega}$ and $\omega_0(Z_i^\top \hat{\theta}) \leq \bar{\omega}$ for any $i = 1, \dots, n$. In addition, $-\omega_1(\hat{Z}_i^\top \hat{\theta}) \leq \bar{\omega}$ and $-\omega_1(Z_i^\top \hat{\theta}) \leq \bar{\omega}$ for any $i = 1, \dots, n$. Combining the above arguments with Eq.A.45 and Eq.A.46:

$$\mathbb{G}_n(\hat{\theta}) - \mathbb{G}_n(\theta_0) \geq -\bar{\omega}(\hat{\theta} - \theta_0)^\top \frac{1}{2n} \sum_{i=1}^n (Z_i Z_i^\top + \hat{Z}_i \hat{Z}_i^\top)(\hat{\theta} - \theta_0). \quad (\text{A.47})$$

Under Assumption 16, $\frac{1}{n} \sum_{i=1}^n \hat{Z}_i \hat{Z}_i^\top$ and $\frac{1}{n} \sum_{i=1}^n Z_i Z_i^\top$ are positive definite matrices for all $X \in \mathcal{X}^n$ and $G \in \mathcal{G}$. This guarantees that the smallest eigenvalues of $\frac{1}{n} \sum_{i=1}^n \hat{Z}_i \hat{Z}_i^\top$ and $\frac{1}{n} \sum_{i=1}^n Z_i Z_i^\top$ are positive. However, the smallest eigenvalue still depends on the size of the training sample. Lemma A.4.2 addresses this by guaranteeing the existence of a uniform lower bound on the smallest eigenvalue. Lemma A.4.2 is a corollary of *Berge's Maximum Theorem* (Lemma A.5.2). A proof is provided in Appendix A.4.2. Given that any element in Z_i and \hat{Z}_i is a linear combination of products between X_i , G_{ij} , σ_i^{data} , D_i , and $D \in \{0, 1\}^n$, $\sigma^{data} \in [0, 1]^n$, $G \in \{0, 1\}^{n \times n}$, and $X \in \mathcal{X}^n$ are both compact, the *Extreme Value Theorem* (Lemma A.5.1) guarantees the existence of a minimum smallest eigenvalue, which only depends on the bound of each element¹⁶ and is independent of the training data. Assumption 17 guarantees that the average number of treated units in the training data is non-zero for any network size. Combined with the above arguments, this ensures that the matrices formed by $\frac{1}{n} \sum_{i=1}^n \hat{Z}_i \hat{Z}_i^\top$ and $\frac{1}{n} \sum_{i=1}^n Z_i Z_i^\top$ have strictly positive minimum smallest eigenvalues. We denote these minimum smallest eigenvalues as $\underline{\varsigma}_{min}^0$ and $\underline{\varsigma}_{min}^1$, respectively. In addition, let $\underline{\varsigma}_{min} = \min\{\underline{\varsigma}_{min}^0, \underline{\varsigma}_{min}^1\}$. Therefore,

$$(\hat{\theta} - \theta_0)^\top \frac{1}{2n} \sum_{i=1}^n (Z_i Z_i^\top + \hat{Z}_i \hat{Z}_i^\top)(\hat{\theta} - \theta_0) \geq \underline{\varsigma}_{min} \|\hat{\theta} - \theta_0\|_2^2 > 0. \quad (\text{A.48})$$

Combining Eq.3.30 with Eq.A.47 and Eq.A.48, we conclude that the sampling uncertainty of $\hat{\theta}$ is characterized by the empirical process $\mathbb{G}_n(\cdot)$. Formally:

$$\mathbb{E}_{\varepsilon^n} [\|\hat{\theta} - \theta_0\|_1 | S, \sigma^{data}] \leq \frac{d_\theta}{-\bar{\omega} \underline{\varsigma}_{min}} \mathbb{E}_{\varepsilon^n} [\|\nabla_{\theta} \mathbb{G}_n(\tilde{\theta})\|_1 | S, \sigma^{data}].$$

¹⁶Although the equilibrium σ and $\hat{\sigma}^{data}$ may not be continuous functions of X , D , and G , the inherent compactness of Z_i and \hat{Z}_i for all $i = 1, \dots, n$ suffices to apply the Extreme Value Theorem.

To complete the proof, let $C_2 = \frac{d_\theta}{-\underline{\omega}_{\min}}$. □

A.2.4 Proof of Lemma 3.5.3

Lemma A.2.3. *Under Assumption 12 to 16, we have*

$$\mathbb{E}_{\varepsilon^n} \left[\sup_{\theta \in \Theta} \|\nabla_\theta \mathbb{G}_n(\theta)\|_1 \middle| S, \sigma^{data} \right] \leq \frac{C_3 + C_4 \sqrt{\log(n)}}{\sqrt{n}},$$

where C_3 and C_4 are constants that depend only on the support of covariates, the distribution of ε , C_σ , the covariates space \mathcal{X} , the network space \mathcal{G} and the treatment allocation space \mathcal{D} .

Proof. Recall we are using a two-step ML estimation procedure, so the first step of estimation introduces additional sampling uncertainty through $\hat{\sigma}^{data}$. To separate the sampling uncertainty of the first and the second steps, we introduce $\mathbb{M}(\theta)$, which is the likelihood function evaluated at the true equilibrium in the training data σ^{data} :

$$\mathbb{M}(\theta) = \frac{1}{n} \sum_{i=1}^n Y_i \log(F_\varepsilon(Z_i^\top \theta)) + (1 - Y_i) \log(1 - F_\varepsilon(Z_i^\top \theta)).$$

We then decompose the initial empirical process $\mathbb{G}_n(\theta)$ into two parts:

$$\mathbb{G}_n(\theta) = \hat{\mathbb{M}}(\theta) - \mathbb{M}(\theta) + \mathbb{M}(\theta) - M(\theta).$$

The first term measures the uncertainty of using the estimated equilibrium $\hat{\sigma}^{data}$, and the second term measures the uncertainty of using the estimated parameter $\hat{\theta}$. Rewrite the gradient of the empirical process as

$$\nabla_\theta \mathbb{G}_n(\tilde{\theta}) = \nabla_\theta \hat{\mathbb{M}}(\tilde{\theta}) - \nabla_\theta \mathbb{M}(\tilde{\theta}) + \nabla_\theta \mathbb{M}(\tilde{\theta}) - \nabla_\theta M(\tilde{\theta}).$$

By the triangle inequality, we have:

$$\|\nabla_\theta \mathbb{G}_n(\tilde{\theta})\|_1 \leq \|\nabla_\theta \hat{\mathbb{M}}(\tilde{\theta}) - \nabla_\theta \mathbb{M}(\tilde{\theta})\|_1 + \|\nabla_\theta \mathbb{M}(\tilde{\theta}) - \nabla_\theta M(\tilde{\theta})\|_1. \quad (\text{A.49})$$

The gradient of $\mathbb{G}_n(\cdot)$, $\hat{\mathbb{M}}(\cdot)$, $\mathbb{M}(\cdot)$ and $M(\cdot)$ are $d_\theta \times 1$ vectors. Let $\nabla_k \mathbb{G}_n(\cdot)$, $\nabla_k \hat{\mathbb{M}}(\cdot)$, $\nabla_k \mathbb{M}(\cdot)$ and $\nabla_k M(\cdot)$ denote their k -th elements. In addition, we define a sequence of empirical processes $\{\mathbb{B}_k(\theta)\}_{k=1}^{d_\theta}$ where $\mathbb{B}_k(\theta) := \nabla_k \mathbb{M}(\theta) - \nabla_k M(\theta)$, and a sequence of stochastic processes $\{\mathbb{A}_k(\theta)\}_{k=1}^{d_\theta}$ where $\mathbb{A}_k(\theta) := \nabla_k \hat{\mathbb{M}}(\theta) - \nabla_k \mathbb{M}(\theta)$. The first term in Eq. A.49 is then

$$\|\nabla_\theta \hat{\mathbb{M}}(\tilde{\theta}) - \nabla_\theta \mathbb{M}(\tilde{\theta})\|_1 = \sum_{k=1}^{d_\theta} |\mathbb{A}_k(\tilde{\theta})|, \quad (\text{A.50})$$

and the second term in Eq.A.49 is

$$\|\nabla_{\theta}\mathbb{M}(\tilde{\theta}) - \nabla_{\theta}M(\tilde{\theta})\|_1 = \sum_{k=1}^{d_{\theta}} |\mathbb{B}_k(\tilde{\theta})|. \quad (\text{A.51})$$

Combining Eq.A.49 with Eq.A.50 and Eq.A.51, we conclude

$$\mathbb{E}_{\varepsilon^n} [\|\nabla_{\theta}\mathbb{G}_n(\tilde{\theta})\|_1 | S, \sigma^{data}] \leq \sum_{k=1}^{d_{\theta}} \mathbb{E}_{\varepsilon^n} [|\mathbb{A}_k(\tilde{\theta})| | S, \sigma^{data}] + \sum_{k=1}^{d_{\theta}} \mathbb{E}_{\varepsilon^n} [|\mathbb{B}_k(\tilde{\theta})| | S, \sigma^{data}]. \quad (\text{A.52})$$

An upper bound for the first term in Eq.A.52 is provided by Lemma A.4.3. An upper bound for the second term in Eq.A.52 is provided by Lemma A.4.4. Combining them,

$$\mathbb{E}_{\varepsilon^n} \left[\sup_{\theta \in \Theta} \|\nabla_{\theta}\mathbb{G}_n(\theta)\|_1 | S, \sigma^{data} \right] \leq d_{\theta} \frac{C_A \sqrt{1 + \ln(2)} + 1}{\sqrt{n}} + \frac{d_{\theta}^{3/2} C_{B1}}{\sqrt{n}} \sqrt{\log(1 + C_{B2} \sqrt{n})}.$$

In addition, given $n \geq 2$,

$$\log(1 + C_{B2} \sqrt{n}) \leq 2 \log(C_{B2} \sqrt{n}) = 2 \log(C_{B2}) + \log(n).$$

As a consequence,

$$\mathbb{E}_{\varepsilon^n} \left[\sup_{\theta \in \Theta} \|\nabla_{\theta}\mathbb{G}_n(\theta)\|_1 | S, \sigma^{data} \right] \leq d_{\theta} \frac{C_A \sqrt{1 + \ln(2)} + 1}{\sqrt{n}} + \frac{d_{\theta}^{3/2} C_{B1}}{\sqrt{n}} \sqrt{\log(C_{B2})} + \frac{d_{\theta}^{3/2} C_{B1}}{\sqrt{n}} \sqrt{\log(n)}.$$

To complete the proof, let $C_3 = (C_A \sqrt{1 + \ln(2)} + 1) d_{\theta} + d_{\theta}^{3/2} C_{B1}$, and $C_4 = d_{\theta}^{3/2} C_{B1}$. \square

A.3 Theorems and Propositions

A.3.1 Proof of Theorem 3.3.1

Theorem A.3.1. *For a supermodular game, the least favorable equilibrium selection rule $\underline{\lambda}$ and the most favorable equilibrium selection rule $\bar{\lambda}$ are given as:*

$$\underline{\lambda} := \delta_{\underline{\sigma}^*}, \quad \bar{\lambda} := \delta_{\bar{\sigma}^*},$$

where δ_{σ} is the Dirac measure on Σ . In addition, the following conditions are satisfied:

$$\inf_{\lambda \in \Lambda} \sum_{i=1}^N \Pr(Y_i = 1 | X, D, G, \lambda) = \sum_{i=1}^N \inf_{\lambda \in \Lambda} \Pr(Y_i = 1 | X, D, G, \lambda),$$

$$\sup_{\lambda \in \Lambda} \sum_{i=1}^N \Pr(Y_i = 1 | X, D, G, \lambda) = \sum_{i=1}^N \sup_{\lambda \in \Lambda} \Pr(Y_i = 1 | X, D, G, \lambda).$$

Proof. Recall that the conditional choice probability is:

$$\begin{aligned}
\Pr(Y_i = 1|X, D, G, \lambda) &= \sum_{y_{-i} \in \mathcal{Y}^{N-1}} \Pr(Y_i = 1, Y_{-i} = y_{-i}|X, D, G, \lambda) \\
&= \sum_{y_{-i} \in \mathcal{Y}^{N-1}} \int \lambda(1, y_{-i}|X, D, G, \varepsilon) dF_\varepsilon \\
&= \int \sum_{y_{-i} \in \mathcal{Y}^{N-1}} \lambda(1, y_{-i}|X, D, G, \varepsilon) dF_\varepsilon.
\end{aligned}$$

Summing over all the units yields

$$\begin{aligned}
\sum_{i=1}^N \Pr(Y_i = 1|X, D, G, \lambda) &= \sum_{i=1}^N \sum_{y_{-i} \in \mathcal{Y}^{N-1}} \Pr(Y_i = 1, Y_{-i} = y_{-i}|X, D, G, \lambda) \\
&= \sum_{i=1}^N \sum_{y_{-i} \in \mathcal{Y}^{N-1}} \int \lambda(1, y_{-i}|X, D, G, \varepsilon) dF_\varepsilon \\
&= \sum_{i=1}^N \int \sum_{y_{-i} \in \mathcal{Y}^{N-1}} \lambda(1, y_{-i}|X, D, G, \varepsilon) dF_\varepsilon
\end{aligned}$$

Given the properties of a supermodular game, there always exists a maximal pure strategy Bayesian Nash equilibrium \bar{y}_ε and a minimal pure strategy Bayesian Nash equilibrium $\underline{y}_\varepsilon$ for all ε . Recall Eq.3.7, for a given $\varepsilon \in \mathbb{R}^N$, these two extreme equilibria can be represented by:

$$\begin{aligned}
\underline{y}_\varepsilon^i &= \mathbb{1} \left\{ \alpha_i + \sum_{j \neq i} \beta_{ij} \underline{\sigma}_j^*(X, D, G) \geq \varepsilon_i \right\}, \quad \forall i \in \mathcal{N}. \\
\bar{y}_\varepsilon^i &= \mathbb{1} \left\{ \alpha_i + \sum_{j \neq i} \beta_{ij} \bar{\sigma}_j^*(X, D, G) \geq \varepsilon_i \right\}, \quad \forall i \in \mathcal{N}.
\end{aligned}$$

Therefore, $\underline{y}_\varepsilon$ happens with probability 1 under our defined least favorable equilibrium selection rule $\underline{\lambda}$, and \bar{y}_ε happens with probability 1 under our defined most favorable equilibrium selection rule $\bar{\lambda}$. We know that $\bar{y}_\varepsilon \geq \underline{y}_\varepsilon$ for any ε where the order in here is *product order*. For any Bayesian Nash equilibrium $y_\varepsilon \in \bar{\Sigma}(X, D, G, \varepsilon)$, we must have $\bar{y}_\varepsilon^i \geq y_\varepsilon^i \geq \underline{y}_\varepsilon^i$ for any $i \in \mathcal{N}$, $\varepsilon \in \mathbb{R}^N$. Therefore, there are only three possible scenarios for each unit i :

- $\bar{y}_\varepsilon^i = y_\varepsilon^i = \underline{y}_\varepsilon^i = 1$.
- $\bar{y}_\varepsilon^i = y_\varepsilon^i = \underline{y}_\varepsilon^i = 0$.
- $\bar{y}_\varepsilon^i = 1, \underline{y}_\varepsilon^i = 0$ and $y_\varepsilon^i \in \{0, 1\}$.

Recall $\Pr(Y_i = 1|X, D, G, \lambda, \varepsilon) = \sum_{y_{-i} \in \mathcal{Y}^{N-1}} \lambda(1, y_{-i}|X, D, G, \varepsilon)$, for any $\varepsilon \in \mathbb{R}^N$, we must have:

- when $\bar{y}_\varepsilon^i = \underline{y}_\varepsilon^i = 1$, $\Pr(Y_i = 1|X, D, G, \lambda, \varepsilon) = 1$ for all $\lambda \in \Lambda$ and for all $i \in \mathcal{N}$;
- when $\bar{y}_\varepsilon^i = \underline{y}_\varepsilon^i = 0$, $\Pr(Y_i = 1|X, D, G, \lambda, \varepsilon) = 0$ for all $\lambda \in \Lambda$ and for all $i \in \mathcal{N}$;
- when $\bar{y}_\varepsilon^i = 1$ and $\underline{y}_\varepsilon^i = 0$, $\Pr(Y_i = 1|X, D, G, \underline{\lambda}, \varepsilon) = 0$ for all $i \in \mathcal{N}$;
- when $\bar{y}_\varepsilon^i = 1$ and $\underline{y}_\varepsilon^i = 0$, $\Pr(Y_i = 1|X, D, G, \bar{\lambda}, \varepsilon) = 1$ for all $i \in \mathcal{N}$.

Therefore, for all $\varepsilon \in \mathbb{R}^N$ and $i \in \mathcal{N}$,

$$\Pr(Y_i = 1|X, D, G, \underline{\lambda}, \varepsilon) \leq \Pr(Y_i = 1|X, D, G, \lambda, \varepsilon), \quad \forall \lambda \in \Lambda,$$

$$\Pr(Y_i = 1|X, D, G, \bar{\lambda}, \varepsilon) \geq \Pr(Y_i = 1|X, D, G, \lambda, \varepsilon), \quad \forall \lambda \in \Lambda.$$

As a consequence, the following two conditions are also satisfied:

$$\Pr(Y_i = 1|X, D, G, \underline{\lambda}) \leq \Pr(Y_i = 1|X, D, G, \lambda), \quad \forall \lambda \in \Lambda, \quad (\text{A.53})$$

$$\Pr(Y_i = 1|X, D, G, \bar{\lambda}) \geq \Pr(Y_i = 1|X, D, G, \lambda), \quad \forall \lambda \in \Lambda, \quad (\text{A.54})$$

Therefore, we must have:

$$\begin{aligned} \sum_{i=1}^N \Pr(Y_i = 1|X, D, G, \underline{\lambda}) &\leq \sum_{i=1}^N \Pr(Y_i = 1|X, D, G, \lambda), \quad \forall \lambda \in \Lambda, \\ \sum_{i=1}^N \Pr(Y_i = 1|X, D, G, \bar{\lambda}) &\geq \sum_{i=1}^N \Pr(Y_i = 1|X, D, G, \lambda), \quad \forall \lambda \in \Lambda. \end{aligned}$$

As a consequence,

$$\begin{aligned} \underline{\lambda} &= \arg \min_{\lambda \in \Lambda} \sum_{i=1}^N \Pr(Y_i = 1|X, D, G, \lambda). \\ \bar{\lambda} &= \arg \max_{\lambda \in \Lambda} \sum_{i=1}^N \Pr(Y_i = 1|X, D, G, \lambda). \end{aligned}$$

In addition, given Eq.A.53 and Eq.A.54, we have:

$$\underline{\lambda} = \arg \min_{\lambda \in \Lambda} \Pr(Y_i = 1|X, D, G, \lambda), \quad \forall i \in \mathcal{N}.$$

$$\bar{\lambda} = \arg \max_{\lambda \in \Lambda} \Pr(Y_i = 1|X, D, G, \lambda), \quad \forall i \in \mathcal{N}.$$

Therefore,

$$\inf_{\lambda \in \Lambda} \sum_{i=1}^N \Pr(Y_i = 1|X, D, G, \lambda) = \sum_{i=1}^N \inf_{\lambda \in \Lambda} \Pr(Y_i = 1|X, D, G, \lambda),$$

$$\sup_{\lambda \in \Lambda} \sum_{i=1}^N \Pr(Y_i = 1 | X, D, G, \lambda) = \sum_{i=1}^N \sup_{\lambda \in \Lambda} \Pr(Y_i = 1 | X, D, G, \lambda).$$

□

A.3.2 Proof of Theorem 3.5.1

Theorem A.3.2. (Sampling Uncertainty of Regret) Under Assumption 12 to 17, the sampling uncertainty of the two-step MLE estimator is bounded by:

$$\mathbb{E}_{\varepsilon^n} [\|\hat{\theta} - \theta_0\|_1 | S, \sigma^{data}] \leq C_2 \frac{C_3 + C_4 \log(n)}{\sqrt{n}}$$

In addition, the sampling uncertainty of the empirical welfare is bounded by:

$$\mathbb{E}_{\varepsilon^n} \left[\max_{D \in \mathcal{D}} |W_n(D) - W(D)| \middle| S, \sigma^{data} \right] \leq C_1 C_2 \frac{C_3 + C_4 \log(n)}{\sqrt{n}}.$$

Proof. Recall

$$\mathbb{E}_{\varepsilon^n} [\|\hat{\theta} - \theta_0\|_1 | S, \sigma^{data}] \leq C_2 \mathbb{E}_{\varepsilon^n} [\|\nabla_{\theta} \mathbb{G}_n(\tilde{\theta})\|_1 | S, \sigma^{data}]. \quad (\text{A.55})$$

Plugging Lemma 3.5.3 into Eq.A.55 leads to

$$\mathbb{E}_{\varepsilon^n} [\|\hat{\theta} - \theta_0\|_1 | S, \sigma^{data}] \leq C_2 \frac{C_3 + C_4 \log(n)}{\sqrt{n}}.$$

Combining Lemma 3.5.1 with Eq.A.55, we conclude:

$$\mathbb{E}_{\varepsilon^n} \left[\max_{D \in \mathcal{D}} |W_n(D) - W(D)| \middle| S, \sigma^{data} \right] \leq C_1 C_2 \frac{C_3 + C_4 \log(n)}{\sqrt{n}}.$$

□

A.3.3 Proof of Proposition 3.5.1

Proposition A.3.1. Under Assumptions 12 and Assumptions 14, the curvature ξ of $W_n(\mathcal{D})$ and the submodularity ratio γ of $W_n(\mathcal{D})$ are in $(0, 1)$. The greedy algorithm enjoys the following approximation guarantee for the problem in Eq.3.19:

$$W_n(D_G) \geq \frac{1}{\xi} (1 - e^{-\xi\gamma}) W_n(\tilde{D}),$$

where D_G is the treatment assignment rule that is obtained by Algorithm 5.

Proof. The curvature is defined as the smallest value of ξ such that

$$W_n(R \cup \{k\}) - W_n(R) \geq (1 - \xi) [W_n(S \cup \{k\}) - W_n(S)] \quad \forall S \subseteq R \subseteq \mathcal{N}, \forall k \in \mathcal{N} \setminus R.$$

As a consequence,

$$\xi = \max_{S \subseteq R \subset \mathcal{N}, k \in \mathcal{N} \setminus R} 1 - \frac{W_n(R \cup \{k\}) - W_n(R)}{W_n(S \cup \{k\}) - W_n(S)}.$$

The submodularity ratio of a non-negative set function is the largest γ such that

$$\sum_{k \in R \setminus S} W_n(S \cup \{k\}) - W_n(S) \geq \gamma [W_n(S \cup R) - W_n(S)], \quad \forall S, R \subseteq \mathcal{N}.$$

As a consequence,

$$\gamma = \min_{S \neq R} \frac{\sum_{k \in R \setminus S} [W_n(S \cup \{k\}) - W_n(S)]}{W_n(S \cup R) - W_n(S)}$$

Recall the utility specification in Eq.3.16, we denote $\hat{\theta}_0 + \hat{\theta}_1 D_i + X_i^\top \hat{\theta}_2 + X_i^\top \hat{\theta}_3 D_i$ as $\hat{\alpha}_{1i}$ and $\hat{\theta}_0 + X_i^\top \hat{\theta}_2$ as $\hat{\alpha}_{0i}$. To connect to the set function notation, we further denote $D_{\mathcal{R}} = \{D_i = 1 : i \in \mathcal{R}\}$. Therefore, for $i \in \mathcal{R}$, we have:

$$\begin{aligned} W_n^i(\mathcal{R} \cup \{k\}) - W_n^i(\mathcal{R}) &= F_\varepsilon \left(\hat{\alpha}_{1i} + \frac{1}{|\mathcal{N}_i|} \sum_{j \neq i} \hat{\theta}_4 m_{ij} G_{ij} D_j + \frac{1}{|\mathcal{N}_i|} \sum_{j \neq i} \hat{\theta}_5 m_{ij} G_{ij} \underline{\sigma}_j + \frac{1}{|\mathcal{N}_i|} \sum_{j \neq i} \hat{\theta}_6 D_j m_{ij} G_{ij} \underline{\sigma}_j \right) \\ &\quad - F_\varepsilon \left(\hat{\alpha}_{1i} + \frac{1}{|\mathcal{N}_i|} \sum_{j \neq i} \hat{\theta}_4 m_{ij} G_{ij} D'_j + \frac{1}{|\mathcal{N}_i|} \sum_{j \neq i} \hat{\theta}_5 m_{ij} G_{ij} \underline{\sigma}'_j + \frac{1}{|\mathcal{N}_i|} \sum_{j \neq i} \hat{\theta}_6 D'_j m_{ij} G_{ij} \underline{\sigma}'_j \right), \end{aligned}$$

where $\underline{\sigma}_b = \Pr(Y_b = 1 | X, G, D_{\mathcal{R} \cup \{k\}}, \underline{\lambda}; \hat{\theta})$ and $\underline{\sigma}'_b = \Pr(Y_b = 1 | X, G, D_{\mathcal{R}}, \underline{\lambda}; \hat{\theta})$ for all $b = 1, \dots, N$. For $m \in \mathcal{N} \setminus \mathcal{R} \cup \{k\}$, their empirical welfare is given as:

$$\begin{aligned} W_n^m(\mathcal{R} \cup \{k\}) - W_n^m(\mathcal{R}) &= F_\varepsilon \left(\hat{\alpha}_{0i} + \frac{1}{|\mathcal{N}_i|} \sum_{j \neq i} \hat{\theta}_4 m_{ij} G_{ij} D_j + \frac{1}{|\mathcal{N}_m|} \sum_{j \neq m} \hat{\theta}_5 m_{mj} G_{mj} \underline{\sigma}_j \right) \\ &\quad - F_\varepsilon \left(\hat{\alpha}_{0i} + \frac{1}{|\mathcal{N}_i|} \sum_{j \neq i} \hat{\theta}_4 m_{ij} G_{ij} D'_j + \frac{1}{|\mathcal{N}_m|} \sum_{j \neq m} \hat{\theta}_5 m_{mj} G_{mj} \underline{\sigma}'_j \right). \end{aligned}$$

For the unit k , her empirical welfare is given as:

$$\begin{aligned} &W_n^k(\mathcal{R} \cup \{k\}) - W_n^k(\mathcal{R}) \\ &= F_\varepsilon \left(\hat{\alpha}_{1k} + \frac{1}{|\mathcal{N}_k|} \sum_{j \neq k} \hat{\theta}_4 m_{kj} G_{kj} D_j + \frac{1}{|\mathcal{N}_k|} \sum_{j \neq k} \hat{\theta}_5 m_{kj} G_{kj} \underline{\sigma}_j + \frac{1}{|\mathcal{N}_k|} \sum_{j \neq k} \hat{\theta}_6 D_j m_{kj} G_{kj} \underline{\sigma}_j \right) \\ &\quad - F_\varepsilon \left(\hat{\alpha}_{0k} + \frac{1}{|\mathcal{N}_k|} \sum_{j \neq k} \hat{\theta}_4 m_{kj} G_{kj} D'_j + \frac{1}{|\mathcal{N}_k|} \sum_{j \neq k} \hat{\theta}_5 m_{kj} G_{kj} \underline{\sigma}'_j \right). \end{aligned}$$

In addition, the empirical welfare increments from assigning unit k treatment is given as:

$$W_n(\mathcal{R} \cup \{k\}) - W_n(\mathcal{R}) = \sum_{i \in \mathcal{R}} W_n^i(\mathcal{R} \cup \{k\}) - W_n^i(\mathcal{R}) + \sum_{m \in \mathcal{N} \setminus \mathcal{R} \cup \{k\}} W_n^m(\mathcal{R} \cup \{k\}) - W_n^m(\mathcal{R}) \\ + W_n^k(\mathcal{R} \cup \{k\}) - W_n^k(\mathcal{R}).$$

Applying the Mean Value Theorem, and Assumption 12, $W_n(\mathcal{R} \cup \{k\}) - W_n(\mathcal{R})$ is upper bounded by:

$$W_n(\mathcal{R} \cup \{k\}) - W_n(\mathcal{R}) \\ \leq \frac{\tau}{N} \sum_{i \in \mathcal{R}} \left(\frac{1}{|\mathcal{N}_i|} \sum_{j \neq i} \hat{\theta}_5 m_{ij} G_{ij}(\underline{\sigma}_j - \underline{\sigma}'_j) + \frac{1}{|\mathcal{N}_i|} \sum_{j \neq i, k} \hat{\theta}_6 D_j m_{ij} G_{ij}(\underline{\sigma}_j - \underline{\sigma}'_j) + \frac{1}{|\mathcal{N}_i|} (\hat{\theta}_4 + \hat{\theta}_6 \underline{\sigma}_k) m_{ik} G_{ik} \right) \\ + \frac{\tau}{N} \sum_{m \in \mathcal{N} \setminus \mathcal{R} \cup \{k\}} \left(\frac{1}{|\mathcal{N}_m|} \hat{\theta}_4 m_{mk} G_{mk} + \frac{1}{|\mathcal{N}_m|} \sum_{j \neq m} \hat{\theta}_5 m_{mj} G_{mj}(\underline{\sigma}_j - \underline{\sigma}'_j) \right) \\ + \frac{\tau}{N} \left(\hat{\theta}_1 + X_i^\top \hat{\theta}_3 + \frac{1}{|\mathcal{N}_k|} \sum_{j \neq k} \hat{\theta}_5 m_{kj} G_{kj}(\underline{\sigma}_j - \underline{\sigma}'_j) + \frac{1}{|\mathcal{N}_k|} \sum_{j \neq k} \hat{\theta}_6 D_j m_{kj} G_{kj} \underline{\sigma}_j \right),$$

where $|\theta|$ denotes the element-wise absolute value of θ . Since $(\underline{\sigma}_i - \underline{\sigma}'_i) \in [0, 1]$ and $D_i, \underline{\sigma}_i \in \{0, 1\}$ for all $i \in \mathcal{N}$, and recall $\bar{m} := \max_{ij} |m_{ij}|$, we can further upper bound the above equation by:

$$W_n(\mathcal{R} \cup \{k\}) - W_n(\mathcal{R}) \\ \leq \frac{\tau}{N} \sum_{i \in \mathcal{R}} \left(\frac{1}{|\mathcal{N}_i|} \sum_{j \neq i} \hat{\theta}_5 \bar{m} G_{ij} + \frac{1}{|\mathcal{N}_i|} \sum_{j \neq i} \hat{\theta}_6 \bar{m} G_{ij} + \frac{\hat{\theta}_4 \bar{m} G_{ik}}{\underline{N}} \right) \\ + \frac{\tau}{N} \sum_{m \in \mathcal{N} \setminus \mathcal{R} \cup \{k\}} \left(\frac{\hat{\theta}_4 \bar{m} G_{mk}}{\underline{N}} + \frac{1}{|\mathcal{N}_m|} \sum_{j \neq m} \hat{\theta}_5 \bar{m} G_{mj} \right) \\ + \frac{\tau}{N} \left(\hat{\theta}_1 + X_i^\top \hat{\theta}_3 + \frac{1}{|\mathcal{N}_k|} \sum_{j \neq k} \hat{\theta}_5 \bar{m} G_{kj} + \frac{1}{|\mathcal{N}_k|} \sum_{j \neq k} \hat{\theta}_6 \bar{m} G_{kj} \right).$$

Given $\frac{1}{|\mathcal{N}_i|} \sum_{j \neq i} G_{ij} = 1$ for all $i, j \in \mathcal{N}$, we can further upper bound the empirical welfare increase by:

$$W_n(\mathcal{R} \cup \{k\}) - W_n(\mathcal{R}) \\ \leq \frac{\tau}{N} \sum_{i \in \mathcal{R}} (\hat{\theta}_5 + \hat{\theta}_6) \bar{m} + \frac{\tau}{N} \sum_{m \in \mathcal{N} \setminus \mathcal{R} \cup \{k\}} \hat{\theta}_5 \bar{m} + \frac{\tau}{N} \left(\hat{\theta}_1 + X_i^\top \hat{\theta}_3 + (\hat{\theta}_5 + \hat{\theta}_6) \bar{m} \right) + \frac{\tau}{N} \frac{\hat{\theta}_4 \bar{m} \bar{N}}{\underline{N}}.$$

Summarizing all the units together, we have:

$$W_n(\mathcal{R} \cup \{k\}) - W_n(\mathcal{R}) \leq \tau(\hat{\theta}_5 + \hat{\theta}_6) \bar{m} + \frac{\tau}{N} (\hat{\theta}_1 + \bar{X}^\top \hat{\theta}_3 + \frac{\hat{\theta}_4 \bar{m} \bar{N}}{\underline{N}}),$$

where $\bar{X} = \{X_i : \max_{X_i \in \mathcal{X}} \bar{X}^\top \hat{\theta}_3\}$. The lower bound of $W(\mathcal{R} \cup \{k\}) - W(\mathcal{R})$ is:

$$\begin{aligned}
& W_n(\mathcal{R} \cup \{k\}) - W_n(\mathcal{R}) \\
& \geq \frac{F_\varepsilon}{N} \sum_{i \in \mathcal{R}} \left(\frac{1}{|\mathcal{N}_i|} \sum_{j \neq i} \hat{\theta}_5 m_{ij} G_{ij}(\underline{\sigma}_j - \underline{\sigma}'_j) + \frac{1}{|\mathcal{N}_i|} \sum_{j \neq i, k} \hat{\theta}_6 D_j m_{ij} G_{ij}(\underline{\sigma}_j - \underline{\sigma}'_j) + \frac{1}{|\mathcal{N}_i|} (\hat{\theta}_4 + \hat{\theta}_6 \underline{\sigma}_k) m_{ik} G_{ik} \right) \\
& + \frac{F_\varepsilon}{N} \sum_{m \in \mathcal{N} \setminus \mathcal{R} \cup \{k\}} \left(\frac{1}{|\mathcal{N}_m|} \hat{\theta}_4 m_{mk} G_{mk} + \frac{1}{|\mathcal{N}_m|} \sum_{j \neq m} \hat{\theta}_5 m_{mj} G_{mj}(\underline{\sigma}_j - \underline{\sigma}'_j) \right) \\
& + \frac{F_\varepsilon}{N} \left(\hat{\theta}_1 + X_i^\top \hat{\theta}_3 + \frac{1}{|\mathcal{N}_k|} \sum_{j \neq k} \hat{\theta}_5 m_{kj} G_{kj}(\underline{\sigma}_j - \underline{\sigma}'_j) + \frac{1}{|\mathcal{N}_k|} \sum_{j \neq k} \hat{\theta}_6 D_j m_{kj} G_{kj} \underline{\sigma}_j \right).
\end{aligned} \tag{A.56}$$

There are three different effects of assigning treatment to unit k . The first effect is the direct treatment effect on unit k , which is the third term in Eq.A.56:

$$\underline{\sigma}_k - \underline{\sigma}'_k \geq F_\varepsilon \left(\hat{\theta}_1 + X_i^\top \hat{\theta}_3 + \frac{1}{|\mathcal{N}_k|} \sum_{j \neq k} \hat{\theta}_5 m_{kj} G_{kj}(\underline{\sigma}_j - \underline{\sigma}'_j) + \frac{1}{|\mathcal{N}_k|} \sum_{j \neq k} \hat{\theta}_6 D_j m_{kj} G_{kj} \underline{\sigma}_j \right).$$

Given $\underline{\sigma}_j - \underline{\sigma}'_j \geq 0$ and $\underline{\sigma}_j \geq 0$ for all $j \in \mathcal{N}$, we further bounds the direct effect from below by:

$$\underline{\sigma}_k - \underline{\sigma}'_k \geq F_\varepsilon (\hat{\theta}_1 + X_k^\top \hat{\theta}_3). \tag{A.57}$$

For the units in the treated and untreated groups, the indirect treatment effects manifest differently. Specifically, for units i in the treated group, their indirect treatment effects are given by the first term in Eq.A.56:

$$\underline{\sigma}_i - \underline{\sigma}'_i \geq F_\varepsilon \left(\frac{1}{|\mathcal{N}_i|} \sum_{j \neq i} \hat{\theta}_5 m_{ij} G_{ij}(\underline{\sigma}_j - \underline{\sigma}'_j) + \frac{1}{|\mathcal{N}_i|} \sum_{j \neq i, k} \hat{\theta}_6 D_j m_{ij} G_{ij}(\underline{\sigma}_j - \underline{\sigma}'_j) + \frac{1}{|\mathcal{N}_i|} (\hat{\theta}_4 + \hat{\theta}_6 \underline{\sigma}_k) m_{ik} G_{ik} \right). \tag{A.58}$$

We then further bound the indirect effects in Eq.A.58 by:

$$\underline{\sigma}_i - \underline{\sigma}'_i \geq F_\varepsilon \frac{1}{|\mathcal{N}_i|} \hat{\theta}_4 m_{ik} G_{ik} \geq \frac{F_\varepsilon}{N} \hat{\theta}_4 m_{ik} G_{ik}. \tag{A.59}$$

For units m not in the treated group, the indirect treatment effects, which is given by the second term of Eq.A.56, can be quantified as follows:

$$\underline{\sigma}_m - \underline{\sigma}'_m \geq F_\varepsilon \left(\frac{1}{|\mathcal{N}_m|} \hat{\theta}_4 m_{mk} G_{mk} + \frac{1}{|\mathcal{N}_m|} \sum_{j \neq m} \hat{\theta}_5 m_{mj} G_{mj}(\underline{\sigma}_j - \underline{\sigma}'_j) \right).$$

This is further bounded below by:

$$\underline{\sigma}_m - \underline{\sigma}'_m \geq \frac{F_\varepsilon}{N} \hat{\theta}_4 m_{mk} G_{mk}. \tag{A.60}$$

Combining Eq.A.57, Eq.A.59 and Eq.A.60 with Eq.A.56 leads to

$$W_n(\mathcal{R} \cup \{k\}) - W_n(\mathcal{R}) \geq \frac{F_\varepsilon}{N\bar{N}} \sum_{i \in \mathcal{N} \setminus \{k\}} \left(\hat{\theta}_4 m_{ik} G_{ik} \right) + \frac{F_\varepsilon}{N} (\hat{\theta}_1 + X_i^\top \hat{\theta}_3).$$

Given that unit k has at least \underline{N} neighbors, we have:

$$W_n(\mathcal{R} \cup \{k\}) - W_n(\mathcal{R}) \geq \frac{F_\varepsilon m \hat{\theta}_4 N}{N\bar{N}} + \frac{F_\varepsilon}{N} (\hat{\theta}_1 + X_i^\top \hat{\theta}_3).$$

Therefore,

$$\frac{W_n(R \cup \{k\}) - W_n(R)}{W_n(S \cup \{k\}) - W_n(S)} \geq \frac{\frac{F_\varepsilon m \hat{\theta}_4 N}{N\bar{N}} + \frac{F_\varepsilon}{N} (\hat{\theta}_1 + X_i^\top \hat{\theta}_3)}{\tau(\hat{\theta}_5 + \hat{\theta}_6) \bar{m} + \frac{\tau}{N} (\hat{\theta}_1 + \bar{X}^\top \hat{\theta}_3 + \frac{\hat{\theta}_4 \bar{m} N}{N})},$$

which ranges between (0,1). As a consequence, the submodularity ratio

$$\xi = \max_{S \subseteq R \subset \mathcal{N}, k \in \mathcal{N} \setminus R} 1 - \frac{W_n(R \cup \{k\}) - W_n(R)}{W_n(S \cup \{k\}) - W_n(S)} \in (0, 1).$$

In addition, the curvature

$$\gamma = \min_{S \neq R} \frac{\sum_{k \in R \setminus S} [W_n(S \cup \{k\}) - W_n(S)]}{W_n(S \cup R) - W_n(S)} \in (0, 1).$$

Combining with [Bian et al. \(2017, Theorem 1\)](#), we finish the proof. \square

A.4 Preliminary Lemmas

A.4.1 Lemma A.4.1

Lemma A.4.1. *The $\|\nabla_\sigma I(\sigma, \theta)\|_\infty^{-1}$ is a continuous function with respect to any entries of $\nabla_\sigma I(\sigma, \theta)$.*

Proof. Let A denote $\nabla_\sigma I(\sigma, \theta)$ and A^{-1} denote $(\nabla_\sigma I(\sigma, \theta))^{-1}$. By the definition of the Uniform norm,

$$\|\nabla_\sigma I(\sigma, \theta)\|_\infty^{-1} = \max_{i \in \mathcal{N}} \sum_{j=1}^N |A_{ij}^{-1}|_1.$$

To prove this maximum is a continuous function, two conditions of Berge's Maximum Theorem must be satisfied:

- **Continuous Function:** Let $f_i(A) = \sum_{j=1}^N |A_{ij}^{-1}|$. Then, $f_i(A) = \|A_{i,:}^{-1}\|_1$, where $A_{i,:}$ denotes the i -th row of matrix A . $f_i(A)$ is continuous with respect to the entries of A by matrix inverse operation and the continuity of the ℓ_1 norm calculation with respect to the vector entries.

- **Compact Parameter Space:** The set over which the maximum is taken (the set of row indices i) is trivially compact as it is finite.

Therefore, by Berge's Maximum Theorem, $\|\nabla_{\sigma} I(\sigma, \theta)\|_{\infty}^{-1}$ is a continuous function with respect to any entry of $\nabla_{\sigma} I(\sigma, \theta)$. \square

A.4.2 Lemma A.4.2

Lemma A.4.2. Under Assumption 12, the smallest and largest eigenvalues of $\frac{1}{n} \sum_{i=1}^n \hat{Z}_i \hat{Z}_i^{\top}$ and $\frac{1}{n} \sum_{i=1}^n Z_i Z_i^{\top}$ are continuous functions of any element of Z_i and \hat{Z}_i , for any $i = 1, \dots, n$.

Proof. Denote the matrix $\frac{1}{N} \sum_{i=1}^N Z_i Z_i^{\top}$ as B . B is a symmetric matrix. Therefore, its smallest (η_{min}) and largest (η_{max}) eigenvalues are given by

$$\eta_{min}(B) = \min_{\theta \neq 0 \in \mathbb{R}^{d_{\theta}}} \frac{\theta^{\top} B \theta}{\theta^{\top} \theta}, \quad (\text{A.61})$$

$$\eta_{max}(B) = \max_{\theta \neq 0 \in \mathbb{R}^{d_{\theta}}} \frac{\theta^{\top} B \theta}{\theta^{\top} \theta}. \quad (\text{A.62})$$

Since we are studying the continuous property of η_{min} (η_{max}) to Z given a θ , we restrict θ to be a unit vector (i.e., $\|\theta\|_2 = 1$) without loss of generality. Rewrite above equations as:

$$\eta_{min}(B) = \min_{\theta: \|\theta\|_2=1} \theta^{\top} B \theta.$$

$$\eta_{max}(B) = \max_{\theta: \|\theta\|_2=1} \theta^{\top} B \theta.$$

Now, we apply the Berge Maximum Theorem to Eq.A.61 and Eq.A.62 to study the continuous property. Two conditions of the Berge Maximum Theorem must be satisfied:

- **Continuous Function:** Given $\theta^{\top} B \theta$ is a quadratic function, and $\|\theta\|_2 = 1$ (i.e., $\theta \neq 0$), it must be a continuous function w.r.t. θ and Z_i for all $i \in \mathcal{N}$.
- **Compact Parameter Space:** Given $\|\theta\|_2 = 1$, the parameter space is compact.

Therefore, by Berge's Maximum Theorem, the largest and smallest eigenvalues are continuous functions of any element of Z_i . By employing a symmetric argument to $\frac{1}{n} \sum_{i=1}^n \hat{Z}_i \hat{Z}_i^{\top}$, we finish the proof. \square

A.4.3 Lemma A.4.3

Lemma A.4.3. Under Assumptions 12, 13, and 15:

$$\mathbb{E}_{\varepsilon^n} \left[\sup_{\theta \in \Theta} |\mathbb{A}_k(\theta)| |S, \sigma^{data} \right] \leq C_A \sqrt{\frac{1 + \ln(2)}{n}},$$

where C_A is a constant that depends only on the support of covariates, the distribution of ε and C_{σ} .

Proof. Recall $\mathbb{A}_k(\theta) = \nabla_k \hat{\mathbb{M}}(\theta) - \nabla_k \mathbb{M}(\theta)$, and

$$\begin{aligned}\nabla_k \hat{\mathbb{M}}(\theta) &= \frac{1}{n} \sum_{i=1}^n \left[Y_i \frac{F'_\varepsilon(\hat{Z}_i^\top \theta)}{F_\varepsilon(\hat{Z}_i^\top \theta)} - (1 - Y_i) \frac{F'_\varepsilon(\hat{Z}_i^\top \theta)}{1 - F_\varepsilon(\hat{Z}_i^\top \theta)} \right] Z_{ik}, \\ \nabla_k \mathbb{M}(\theta) &= \frac{1}{n} \sum_{i=1}^n \left[Y_i \frac{F'_\varepsilon(Z_i^\top \theta)}{F_\varepsilon(Z_i^\top \theta)} - (1 - Y_i) \frac{F'_\varepsilon(Z_i^\top \theta)}{1 - F_\varepsilon(Z_i^\top \theta)} \right] Z_{ik}.\end{aligned}$$

Let $\tilde{\theta} := \arg \sup_{\theta \in \Theta} |\mathbb{A}_k(\theta)|$. Therefore,

$$\mathbb{A}_k(\tilde{\theta}) = \frac{1}{n} \sum_{i=1}^n \left[Y_i \left(\frac{F'_\varepsilon(\hat{Z}_i^\top \tilde{\theta})}{F_\varepsilon(\hat{Z}_i^\top \tilde{\theta})} - \frac{F'_\varepsilon(Z_i^\top \tilde{\theta})}{F_\varepsilon(Z_i^\top \tilde{\theta})} \right) - (1 - Y_i) \left(\frac{F'_\varepsilon(\hat{Z}_i^\top \tilde{\theta})}{1 - F_\varepsilon(\hat{Z}_i^\top \tilde{\theta})} - \frac{F'_\varepsilon(Z_i^\top \tilde{\theta})}{1 - F_\varepsilon(Z_i^\top \tilde{\theta})} \right) \right] Z_{ik}.$$

Applying the Mean value theorem,

$$\frac{F'_\varepsilon(\hat{Z}_i^\top \tilde{\theta})}{F_\varepsilon(\hat{Z}_i^\top \tilde{\theta})} - \frac{F'_\varepsilon(Z_i^\top \tilde{\theta})}{F_\varepsilon(Z_i^\top \tilde{\theta})} = (\hat{Z}_i - Z_i)^\top \nabla_Z \frac{F'_\varepsilon(\tilde{Z}_i^\top \tilde{\theta})}{F_\varepsilon(\tilde{Z}_i^\top \tilde{\theta})},$$

and

$$\frac{F'_\varepsilon(\hat{Z}_i^\top \tilde{\theta})}{1 - F_\varepsilon(\hat{Z}_i^\top \tilde{\theta})} - \frac{F'_\varepsilon(Z_i^\top \tilde{\theta})}{1 - F_\varepsilon(Z_i^\top \tilde{\theta})} = (\hat{Z}_i - Z_i)^\top \nabla_Z \frac{F'_\varepsilon(\tilde{Z}_i^\top \tilde{\theta})}{1 - F_\varepsilon(\tilde{Z}_i^\top \tilde{\theta})},$$

for some $\tilde{Z}_i \in \mathbb{R}^{d_\theta}$ on the segment from \hat{Z}_i to Z_i . Therefore,

$$\begin{aligned}|\mathbb{A}_k(\tilde{\theta})| &= \frac{1}{n} \sum_{i=1}^n |Y_i \omega_0(\tilde{Z}_i^\top \tilde{\theta})(\hat{Z}_i - Z_i)^\top \tilde{Z}_i - (1 - Y_i) \omega_1(\tilde{Z}_i^\top \tilde{\theta})(\hat{Z}_i - Z_i)^\top \tilde{Z}_i| Z_{ik} \\ &\leq \frac{1}{n} \sum_{i=1}^n \left[Y_i |\omega_0(\tilde{Z}_i^\top \tilde{\theta})| |(\hat{Z}_i - Z_i)^\top \tilde{Z}_i| + (1 - Y_i) |\omega_1(\tilde{Z}_i^\top \tilde{\theta})| |(\hat{Z}_i - Z_i)^\top \tilde{Z}_i| \right] Z_{ik} \\ &\leq \frac{1}{n} \sum_{i=1}^n |\underline{\omega}| |(\hat{Z}_i - Z_i)^\top \tilde{Z}_i| Z_{ik}.\end{aligned}$$

By the Cauchy–Schwarz inequality, we have:

$$\begin{aligned}|\mathbb{A}_k(\tilde{\theta})| &\leq \frac{1}{n} \sum_{i=1}^n |\underline{\omega}| \|\hat{Z}_i - Z_i\|_1 \|\tilde{Z}_i\|_\infty Z_{ik} \\ &\leq |\underline{\omega}| \bar{z}^2 \frac{1}{n} \sum_{i=1}^n \|\hat{Z}_i - Z_i\|_1,\end{aligned}\tag{A.63}$$

where $\bar{z} := \max_{i=1, \dots, n} \|Z_i\|_\infty$. Recall the definition of \hat{Z}_i from Eq. 3.17:

$$\hat{Z}_i = \left(1, \mathbf{D}_i, \mathbf{X}_i^\top, \mathbf{X}_i^\top \mathbf{D}_i, \frac{1}{|\mathcal{N}_i|} \sum_{j \neq i} m_{ij} \mathbf{G}_{ij} \mathbf{D}_j, \frac{1}{|\mathcal{N}_i|} \sum_{j \neq i} m_{ij} \mathbf{G}_{ij} \hat{\sigma}_j^{data}, \frac{1}{|\mathcal{N}_i|} \sum_{j \neq i} m_{ij} \mathbf{G}_{ij} \hat{\sigma}_j^{data} \mathbf{D}_i \mathbf{D}_j \right)^\top.$$

Therefore, we rewrite $\|\hat{Z}_i - Z_i\|_1$ as:

$$\|\hat{Z}_i - Z_i\|_1 = \left| \frac{1}{|\mathcal{N}_i|} \sum_{j \neq i} m_{ij} \mathbf{G}_{ij}(\hat{\sigma}_j^{data} - \sigma_j^{data}) \right| + \left| \frac{1}{|\mathcal{N}_i|} \sum_{j \neq i} m_{ij} \mathbf{G}_{ij}(\hat{\sigma}_j^{data} - \sigma_j^{data}) \mathbf{D}_i \mathbf{D}_j \right|.$$

By triangle inequality,

$$\|\hat{Z}_i - Z_i\|_1 \leq \frac{1}{|\mathcal{N}_i|} \sum_{j \neq i} |m_{ij}| \mathbf{G}_{ij}(1 + \mathbf{D}_i \mathbf{D}_j) |\hat{\sigma}_j^{data} - \sigma_j^{data}|.$$

Applying Lemma A.4.5, and defining $\bar{m} := \max_{i,j \in \mathcal{N}} |m_{ij}|$, we have

$$\mathbb{E}_{\varepsilon^n} [\|\hat{Z}_i - Z_i\|_1 | S, \sigma^{data}] \leq 3\bar{m} C_\sigma \sqrt{\frac{1 + \ln(2)}{n}}. \quad (\text{A.64})$$

Plug Eq.A.64 into Eq.A.63,

$$\mathbb{E}_{\varepsilon^n} \left[\sup_{\theta \in \Theta} |\mathbb{A}_k(\theta)| | S, \sigma^{data} \right] \leq 3|\underline{\omega}| \bar{z}^2 \bar{m} C_\sigma \sqrt{\frac{1 + \ln(2)}{n}}.$$

Setting $C_A = 3|\underline{\omega}| \bar{z}^2 \bar{m} C_\sigma$, completes the proof. \square

A.4.4 Lemma A.4.4

Lemma A.4.4. Define $\bar{z} := \max_{i=1,\dots,n} \|Z_i\|_\infty$. Under Assumption 12 to 16, we have:

$$\mathbb{E}_{\varepsilon^n} \left[\sup_{\theta \in \Theta} |\mathbb{B}_k(\theta)| | S, \sigma^{data} \right] \leq \frac{1}{\sqrt{n}} \left(1 + C_{B1} \sqrt{d_\theta \log(1 + C_{B2} \sqrt{n})} \right),$$

where C_{B1} is a universal constant that only depend on the distribution F_{ε^n} and $C_{B2} = 4|\underline{\omega}| \bar{z}^2$.

Proof. For a given $\delta \geq 0$ and associated covering number $H = N_c(\delta, \Theta, L_1)$, let $\mathbb{U} := \{\theta^1, \dots, \theta^H\}$ be a δ -cover of Θ . For any $\theta \in \Theta$, we can find some θ^ℓ such that $\|\theta - \theta^\ell\|_1 \leq \delta$. Let $\tilde{\theta} := \arg \sup_{\theta \in \Theta} |\mathbb{B}_k(\theta)|$. Therefore,

$$\begin{aligned} |\mathbb{B}_k(\tilde{\theta})| &= |\mathbb{B}_k(\tilde{\theta}) - \mathbb{B}_k(\theta^\ell) + \mathbb{B}_k(\theta^\ell)| \\ &\leq |\mathbb{B}_k(\tilde{\theta}) - \mathbb{B}_k(\theta^\ell)| + |\mathbb{B}_k(\theta^\ell)| \\ &\leq \sup_{\substack{\gamma, \gamma' \in \Theta \\ \|\gamma - \gamma'\|_1 \leq \delta}} |\mathbb{B}_k(\gamma) - \mathbb{B}_k(\gamma')| + \max_{\ell=1,\dots,H} |\mathbb{B}_k(\theta^\ell)|. \end{aligned} \quad (\text{A.65})$$

Apply Lemma A.4.6 to bound the first term in Eq.A.65:

$$\mathbb{E}_{\varepsilon^n} \left[\sup_{\substack{\gamma, \gamma' \in \Theta \\ \|\gamma - \gamma'\|_1 \leq \delta}} |\mathbb{B}_k(\gamma) - \mathbb{B}_k(\gamma')| | S, \sigma^{data} \right] \leq |\underline{\omega}| \bar{z}^2 \delta, \quad (\text{A.66})$$

where $\bar{z} := \max_{i=1,\dots,n} \|Z_i\|_\infty$. To bound the second term in Eq.A.65, we introduce $\{\tilde{\varepsilon}_i\}_{i=1}^n$, an independent copy of ε^n that follows the same distribution F_{ε^n} . Hence, the associated $\{\tilde{Y}_i\}_{i=1}^n$ (i.e., $\tilde{Y}_i = \mathbb{1}\{\alpha_i + \sum_{j \neq i} \beta_{ij} \sigma_j^{data} - \tilde{\varepsilon}_i \geq 0\}$) has the same distribution as $\{Y_i\}_{i=1}^N$ conditional on the S and σ^{data} . We denote the expectation with respect to $\tilde{\varepsilon}$ as $\mathbb{E}_{\tilde{\varepsilon}}(\cdot)$. Recall that the criterion function is:

$$m_{Y_i, Z_i}(\theta) := Y_i \log(F_\varepsilon(Z_i^\top \theta)) + (1 - Y_i) \log(1 - F_\varepsilon(Z_i^\top \theta)).$$

Denote the empirical measure of our criterion function with $\{\tilde{Y}_i\}_{i=1}^N$ as $\tilde{\mathbb{M}}(\theta)$:

$$\tilde{\mathbb{M}}(\theta) := \frac{1}{n} \sum_{i=1}^n m_{\tilde{Y}_i, Z_i}(\theta).$$

By definition of \tilde{Y}_i , we have $M(\theta) := \mathbb{E}_{\varepsilon^n}[\mathbb{M}(\theta)|S, \sigma^{data}] = \mathbb{E}_{\tilde{\varepsilon}}[\tilde{\mathbb{M}}(\theta)|S, \sigma^{data}]$. Therefore,

$$\begin{aligned} & \mathbb{E}_{\varepsilon^n} \left[\max_{\ell=1,\dots,H} |\mathbb{B}_k(\theta^\ell)| \middle| S, \sigma^{data} \right] \\ &= \mathbb{E}_{\varepsilon^n} \left[\max_{\ell=1,\dots,H} \left| \nabla_k \mathbb{M}(\theta^\ell) - \nabla_k \mathbb{E}_{\tilde{\varepsilon}}[\tilde{\mathbb{M}}(\theta^\ell)|S, \sigma^{data}] \right| \middle| S, \sigma^{data} \right] \\ &= \mathbb{E}_{\varepsilon^n} \left[\max_{\ell=1,\dots,H} \left| \frac{1}{n} \sum_{i=1}^n [\nabla_k m_{Y_i, Z_i}(\theta^\ell) - \nabla_k \mathbb{E}_{\tilde{\varepsilon}}[m_{\tilde{Y}_i, Z_i}(\theta^\ell)|S, \sigma^{data}]] \right| \middle| S, \sigma^{data} \right] \\ &= \mathbb{E}_{\varepsilon^n} \left[\max_{\ell=1,\dots,H} \left| \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{\tilde{\varepsilon}}[\nabla_k m_{Y_i, Z_i}(\theta^\ell) - \nabla_k m_{\tilde{Y}_i, Z_i}(\theta^\ell)|S, \sigma^{data}] \right| \middle| S, \sigma^{data} \right] \\ & \quad (\text{By Leibniz rule}) \\ &\leq \mathbb{E}_{\varepsilon^n, \tilde{\varepsilon}} \left[\max_{\ell=1,\dots,H} \left| \frac{1}{n} \sum_{i=1}^n [\nabla_k m_{Y_i, Z_i}(\theta^\ell) - \nabla_k m_{\tilde{Y}_i, Z_i}(\theta^\ell)] \right| \middle| S, \sigma^{data} \right]. \end{aligned}$$

Define i.i.d Rademacher variables $\nu := (\nu_1, \dots, \nu_n)$ such that $\Pr(\nu_i = 1) = \Pr(\nu_i = -1) = \frac{1}{2}$. Since $m_{Y_i, Z_i}(\theta^\ell) - m_{\tilde{Y}_i, Z_i}(\theta^\ell) \sim \nu_i [m_{Y_i, Z_i}(\theta^\ell) - m_{\tilde{Y}_i, Z_i}(\theta^\ell)]$, we have

$$\begin{aligned} & \mathbb{E}_{\varepsilon^n, \tilde{\varepsilon}} \left[\max_{\ell=1,\dots,H} \left| \frac{1}{n} \sum_{i=1}^n [\nabla_k m_{Y_i, Z_i}(\theta^\ell) - \nabla_k m_{\tilde{Y}_i, Z_i}(\theta^\ell)] \right| \middle| S, \sigma^{data} \right] \\ &= \mathbb{E}_{\varepsilon^n, \tilde{\varepsilon}, \nu} \left[\max_{\ell=1,\dots,H} \left| \frac{1}{n} \sum_{i=1}^n \nu_i [\nabla_k m_{Y_i, Z_i}(\theta^\ell) - \nabla_k m_{\tilde{Y}_i, Z_i}(\theta^\ell)] \right| \middle| S, \sigma^{data} \right] \\ &\leq \mathbb{E}_{\varepsilon^n, \tilde{\varepsilon}, \nu} \left[\max_{\ell=1,\dots,H} \left| \frac{1}{n} \sum_{i=1}^n \nu_i \nabla_k m_{Y_i, Z_i}(\theta^\ell) \right| + \max_{\ell=1,\dots,H} \left| \frac{1}{n} \sum_{i=1}^n \nu_i \nabla_k m_{\tilde{Y}_i, Z_i}(\theta^\ell) \right| \middle| S, \sigma^{data} \right] \\ &= 2\mathbb{E}_{\varepsilon^n, \nu} \left[\max_{\ell=1,\dots,H} \left| \frac{1}{n} \sum_{i=1}^n \nu_i \nabla_k m_{Y_i, Z_i}(\theta^\ell) \right| \middle| S, \sigma^{data} \right] \\ &= 2\mathbb{E}_{\varepsilon^n} \left[\mathbb{E}_\nu \left[\max_{\ell=1,\dots,H} \left| \frac{1}{n} \sum_{i=1}^n \nu_i \nabla_k m_{Y_i, Z_i}(\theta^\ell) \right| \middle| S, \sigma^{data} \right] \middle| S, \sigma^{data} \right]. \end{aligned}$$

By Lemma A.4.7, $\frac{1}{n} \sum_{i=1}^n \nu_i \nabla_k m_{Y_i, Z_i}(\theta^\ell)$ is a sub-Gaussian process with parameter $\tau/\sqrt{nv^2}$. Therefore, by the upper bound of sub-Gaussian maxima (Lemma A.5.5), we have:

$$\mathbb{E} \left[\max_{\ell=1, \dots, H} |\mathbb{B}_k(\theta^\ell)| \middle| S, \sigma^{data} \right] \leq \frac{2\tau}{\sqrt{nv}} \sqrt{\log(N_c(\delta, \Theta, L_1))}. \quad (\text{A.67})$$

Now, apply Lemma A.4.8 to bound the L_1 -metric entropy $\log(N_c(\delta, \Theta, L_1))$:

$$\log(N_c(\delta, \Theta, L_1)) \leq d_\theta \log \left(1 + \frac{2}{\delta} \right). \quad (\text{A.68})$$

Combining Eq.A.67 with Eq.A.68, we have:

$$\mathbb{E} \left[\max_{\ell=1, \dots, H} |\mathbb{B}_k(\theta^\ell)| \middle| S, \sigma^{data} \right] \leq \frac{2\tau\sqrt{d_\theta}}{\sqrt{nv}} \sqrt{\log \left(1 + \frac{2}{\delta} \right)}. \quad (\text{A.69})$$

Combining Eq.A.65 with Eq.A.66 and Eq.A.69, we have:

$$\mathbb{E} \left[|\mathbb{B}_k(\tilde{\theta})| \middle| S, \sigma^{data} \right] \leq \underline{\omega} \bar{z}^2 \delta + \frac{2\tau\sqrt{d_\theta}}{\sqrt{nv}} \sqrt{\log \left(1 + \frac{2}{\delta} \right)}.$$

By choosing $\delta = \frac{1}{\underline{\omega} \bar{z}^2 \sqrt{n}}$, we conclude:

$$\mathbb{E} \left[\sup_{\theta \in \Theta} |\mathbb{B}_k(\theta)| \middle| S, \sigma^{data} \right] \leq \frac{1}{\sqrt{n}} \left(1 + \frac{2\tau}{v} \sqrt{d_\theta \log(1 + 2|\underline{\omega}| \bar{z}^2 \sqrt{n})} \right).$$

To finish the proof, define $C_{B1} = 2\tau/v$, and $C_{B2} = 2|\underline{\omega}| \bar{z}^2$. □

A.4.5 Lemma A.4.5

Lemma A.4.5. Under Assumption 13, for all $i = 1, \dots, n$,

$$\mathbb{E}_{\varepsilon^n} [|\hat{\sigma}_i^{data} - \sigma_i^{data}| \middle| S, \sigma^{data}] \leq C_\sigma \sqrt{\frac{1 + \ln(2)}{n}}.$$

Proof. This proof follows the same proof strategy as Lemma 5.1 in Kitagawa and Wang (2023b). Recall that for any nonnegative random variable Y , $\mathbb{E}(Y) = \int_0^\infty \Pr(Y \geq t) dt$. Hence, for any $a > 0$,

$$\begin{aligned} \mathbb{E}(|\hat{\sigma}_i^{data} - \sigma_i^{data}|^2) &= \int_0^\infty \Pr(|\hat{\sigma}_i^{data} - \sigma_i^{data}|^2 \geq t) dt \\ &= \int_0^a \Pr(|\hat{\sigma}_i^{data} - \sigma_i^{data}|^2 \geq t) dt + \int_a^\infty \Pr(|\hat{\sigma}_i^{data} - \sigma_i^{data}|^2 \geq t) dt \\ &\leq a + \int_a^\infty \Pr(|\hat{\sigma}_i^{data} - \sigma_i^{data}|^2 \geq t) dt. \end{aligned}$$

Assumption 13 implies that $\Pr(|\hat{\sigma}_i^{data} - \sigma_i^{data}| \geq \sqrt{t}) \leq 2e^{-Nt/C_\sigma^2}$. Hence,

$$\begin{aligned} \mathbb{E}(|\hat{\sigma}_i^{data} - \sigma_i^{data}|^2) &\leq a + \int_a^\infty \Pr(|\hat{\sigma}_i^{data} - \sigma_i^{data}|^2 \geq t) dt \\ &= a + \int_a^\infty \Pr(|\hat{\sigma}_i^{data} - \sigma_i^{data}| \geq \sqrt{t}) dt \\ &\leq a + 2 \int_a^\infty e^{-nt/C_\sigma^2} dt \\ &= a + 2 \frac{C_\sigma^2}{n} e^{-Na/C_\sigma^2}. \end{aligned}$$

Set $a = C_\sigma^2 \ln(2)/n$ and we have

$$\mathbb{E}(|\hat{\sigma}_i^{data} - \sigma_i^{data}|^2) \leq \frac{\ln(2)C_\sigma^2}{n} + \frac{C_\sigma^2}{n} = C_\sigma^2 \frac{1 + \ln(2)}{n}.$$

Therefore,

$$\mathbb{E}(|\hat{\sigma}_i^{data} - \sigma_i^{data}|) \leq \sqrt{\mathbb{E}(|\hat{\sigma}_i^{data} - \sigma_i^{data}|^2)} \leq C_\sigma \sqrt{\frac{1 + \ln(2)}{n}}.$$

□

A.4.6 Lemma A.4.6: Lipschitz Property

Lemma A.4.6. (Lipschitz Property) Define $\bar{z} := \max_{i=1,\dots,n} \|Z_i\|_\infty$. The following condition on $\tilde{\mathbb{G}}_n(\cdot)$ is satisfied:

$$\sup_{\substack{\gamma, \gamma' \in \Theta \\ \|\gamma - \gamma'\|_1 \leq \delta}} |\mathbb{B}_k(\gamma) - \mathbb{B}_k(\gamma')| \leq |\omega| \bar{z}^2 \delta.$$

Proof. First, we have:

$$\sup_{\substack{\gamma, \gamma' \in \Theta \\ \|\gamma - \gamma'\|_1 \leq \delta}} |\mathbb{B}_k(\gamma) - \mathbb{B}_k(\gamma')| \leq \sup_{\substack{\gamma, \gamma' \in \Theta \\ \|\gamma - \gamma'\|_1 \leq \delta}} |\nabla_k \mathbb{M}(\gamma) - \nabla_k \mathbb{M}(\gamma')| + \sup_{\substack{\gamma, \gamma' \in \Theta \\ \|\gamma - \gamma'\|_1 \leq \delta}} |\nabla_k M(\gamma') - \nabla_k M(\gamma)|. \quad (\text{A.70})$$

The first term in Eq. A.70 is:

$$\begin{aligned} \nabla_k \mathbb{M}(\gamma) - \nabla_k \mathbb{M}(\gamma') &= \frac{1}{n} \sum_{i=1}^n Y_i \left[\frac{F'_\varepsilon(Z_i^\top \gamma)}{F_\varepsilon(Z_i^\top \gamma)} - \frac{F'_\varepsilon(Z_i^\top \gamma')}{F_\varepsilon(Z_i^\top \gamma')} \right] Z_{ik} \\ &\quad - \frac{1}{n} \sum_{i=1}^n (1 - Y_i) \left[\frac{F'_\varepsilon(Z_i^\top \gamma)}{1 - F_\varepsilon(Z_i^\top \gamma)} - \frac{F'_\varepsilon(Z_i^\top \gamma')}{1 - F_\varepsilon(Z_i^\top \gamma')} \right] Z_{ik}. \end{aligned}$$

The second term in Eq.A.70 is:

$$\begin{aligned}\nabla_k M(\gamma) - \nabla_k M(\gamma') &= \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{\varepsilon^n} \left[Y_i \left[\frac{F'_\varepsilon(Z_i^\top \gamma)}{F_\varepsilon(Z_i^\top \gamma)} - \frac{F'_\varepsilon(Z_i^\top \gamma')}{F_\varepsilon(Z_i^\top \gamma')} \right] Z_{ik} \middle| S, \sigma^{data} \right] \\ &\quad - \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{\varepsilon^n} \left[(1 - Y_i) \left[\frac{F'_\varepsilon(Z_i^\top \gamma)}{1 - F_\varepsilon(Z_i^\top \gamma)} - \frac{F'_\varepsilon(Z_i^\top \gamma')}{1 - F_\varepsilon(Z_i^\top \gamma')} \right] Z_{ik} \middle| S, \sigma^{data} \right].\end{aligned}$$

Applying the Mean Value Theorem to both, we have:

$$\begin{aligned}\nabla_k \mathbb{M}(\gamma) - \nabla_k \mathbb{M}(\gamma') &= \frac{1}{n} \sum_{i=1}^n \left[Y_i \omega_0(Z_i^\top \hat{\gamma}) - (1 - Y_i) \omega_1(Z_i^\top \hat{\gamma}) \right] Z_{ik} Z_i^\top (\gamma - \gamma'), \\ \nabla_k M(\gamma) - \nabla_k M(\gamma') &= \frac{1}{n} \sum_{i=1}^n \left[\sigma_i^n \omega_0(Z_i^\top \hat{\gamma}) - (1 - \sigma_i) \omega_1(Z_i^\top \hat{\gamma}) \right] Z_{ik} Z_i^\top (\gamma - \gamma').\end{aligned}$$

where $\sigma_i^n = \mathbb{E}_{\varepsilon^n}[Y_i | S, \sigma^{data}]$, and for some $\hat{\gamma} \in \mathbb{R}^{d_\theta}$, $\hat{\gamma} \in \mathbb{R}^{d_\theta}$ on the segment from γ to γ' . Then,

$$|\nabla_k \mathbb{M}(\gamma) - \nabla_k \mathbb{M}(\gamma')| \leq \frac{1}{n} \sum_{i=1}^n |Y_i \omega_0(Z_i^\top \hat{\gamma}) - (1 - Y_i) \omega_1(Z_i^\top \hat{\gamma})| \cdot |Z_{ik}| \cdot |Z_i^\top (\gamma - \gamma')|. \quad (\text{A.71})$$

By Assumption 15, $\omega_0(a) < 0$ and $\omega_1(a) > 0$ for all $a \in \mathbb{R}$. Recall that

$$\underline{\omega}_0 := \min_{x \in \Xi} \omega_0(x), \quad \underline{\omega}_1 := \min_{x \in \Xi} -\omega_1(x), \quad \underline{\omega} := \min\{\underline{\omega}_0, \underline{\omega}_1\}. \quad (\text{A.72})$$

Combining Eq.A.71 with Eq.A.72, we have

$$\begin{aligned}|\nabla_k \mathbb{M}(\gamma) - \nabla_k \mathbb{M}(\gamma')| &\leq \frac{|\underline{\omega}|}{n} \sum_{i=1}^n |Z_i^\top (\gamma - \gamma')| \cdot |Z_{ik}| \\ &\leq \frac{|\underline{\omega}|}{n} \sum_{i=1}^n \|Z_i\|_2^2 \|\gamma - \gamma'\|_1 \\ &\quad (\text{By Holder's Inequality}) \\ &\leq |\underline{\omega}| \bar{z}^2 \|\gamma - \gamma'\|_1,\end{aligned}$$

where $\bar{z} := \max_{i=1, \dots, n} \|Z_i\|_\infty$. By the same argument,

$$|\nabla_k M(\gamma) - \nabla_k M(\gamma')| \leq |\underline{\omega}| \bar{z}^2 \|\gamma - \gamma'\|_1.$$

Combining above two equations with Eq.A.70 gives:

$$\sup_{\substack{\gamma, \gamma' \in \Theta \\ \|\gamma - \gamma'\|_1 \leq \delta}} |\mathbb{B}_k(\gamma) - \mathbb{B}_k(\gamma')| \leq |\underline{\omega}| \bar{z}^2 \delta.$$

□

A.4.7 Lemma A.4.7: Sub-Gaussian Process

Lemma A.4.7. (Sub-Gaussian Process) Define $\bar{z} := \max_{i=1, \dots, n} \|Z_i\|_\infty \cdot \frac{1}{n} \sum_{i=1}^n \nu_i \nabla_k m_{Y_i, Z_i}(\theta^\ell)$ is a sub-Gaussian process with parameter $\tau / \sqrt{nv^2}$.

Proof. We start from the expectation of the moment-generating function of $1/n \sum_{i=1}^n \nu_i \nabla_k m_{Y_i, Z_i}(\cdot)$, which is

$$\mathbb{E}_\nu \left[\exp \left[\frac{1}{n} \sum_{i=1}^n s \nu_i \nabla_k m_{Y_i, Z_i}(\theta^\ell) \right] \middle| S, \sigma^{data} \right] = \prod_{i=1}^n \mathbb{E}_\nu \left[\exp \left[\frac{s}{n} \nu_i \nabla_k m_{Y_i, Z_i}(\theta^\ell) \right] \middle| S, \sigma^{data} \right],$$

where the equality holds as $\{\varepsilon_i\}_{i=1}^n, \{v_i\}_{i=1}^n$ are i.i.d. In addition, the gradient of $m_{Y_i, Z_i}(\theta^\ell)$ is:

$$\nabla_\theta m_{Y_i, Z_i}(\theta^\ell) = \left[Y_i \frac{F'_\varepsilon(Z_i^\top \theta^\ell)}{F_\varepsilon(Z_i^\top \theta^\ell)} - (1 - Y_i) \frac{F'_\varepsilon(Z_i^\top \theta^\ell)}{1 - F_\varepsilon(Z_i^\top \theta^\ell)} \right] Z_i.$$

Therefore,

$$\begin{aligned} & \mathbb{E}_\nu \left[\exp \left[\frac{1}{n} \sum_{i=1}^n s \nu_i \nabla_k m_{Y_i, Z_i}(\theta^\ell) \right] \middle| S, \sigma^{data} \right] \\ &= \prod_{i=1}^n \mathbb{E}_\nu \left[\exp \left[\frac{s \nu_i}{n} \left(Y_i \frac{F'_\varepsilon(Z_i^\top \theta^\ell)}{F_\varepsilon(Z_i^\top \theta^\ell)} - (1 - Y_i) \frac{F'_\varepsilon(Z_i^\top \theta^\ell)}{1 - F_\varepsilon(Z_i^\top \theta^\ell)} \right) Z_{ik} \right] \middle| S, \sigma^{data} \right]. \end{aligned}$$

By Hoeffding's Lemma (Lemma A.5.4),

$$\begin{aligned} & \mathbb{E}_\nu \left[\exp \left[\frac{1}{n} \sum_{i=1}^n s \nu_i \nabla_k m_{Y_i, Z_i}(\theta^\ell) \right] \middle| S, \sigma^{data} \right] \\ &\leq \prod_{i=1}^n \exp \left[\frac{s^2}{2n^2} \left(Y_i \frac{F'_\varepsilon(Z_i^\top \theta^\ell)}{F_\varepsilon(Z_i^\top \theta^\ell)} - (1 - Y_i) \frac{F'_\varepsilon(Z_i^\top \theta^\ell)}{1 - F_\varepsilon(Z_i^\top \theta^\ell)} \right)^2 Z_{ik}^2 \right] \\ &\leq \prod_{i=1}^n \exp \left[\frac{s^2}{2n^2} \left(Y_i \frac{\tau}{v} + (1 - Y_i) \frac{\tau}{v} \right)^2 Z_{ik}^2 \right] \\ &= \exp \left[\frac{s^2 \tau^2}{2n^2 v^2} \sum_{i=1}^n Z_{ik}^2 \right] \\ &\leq \exp \left[\frac{s^2 \tau^2}{2nv^2} \bar{z}^2 \right]. \end{aligned}$$

Recall $v := \min\{\underline{F}_\varepsilon, 1 - \bar{F}_\varepsilon\}$, where $\underline{F}_\varepsilon := \min_{\substack{\theta \in \Theta \\ z \in \mathcal{Z}}} F_\varepsilon(z^\top \theta)$, and $\bar{F}_\varepsilon := \max_{\substack{\theta \in \Theta \\ z \in \mathcal{Z}}} F_\varepsilon(z^\top \theta)$. Therefore, $\frac{1}{n} \sum_{i=1}^n \nu_i \nabla_k m_{Y_i, Z_i}(\theta^\ell)$ is a sub-Gaussian process with parameter $\tau / \sqrt{nv^2}$. \square

A.4.8 Lemma A.4.8: Covering Number

Lemma A.4.8. (Covering Number) The δ -covering number of a compact parameter space $\Theta \in \mathbb{R}^{d_\theta}$ with L_1 metric $N_c(\delta, \Theta, L_1)$ is upper bounded by $(1 + \frac{1}{\delta})^{d_\theta}$.

Proof. As parameter space Θ is compact, there exists a constant C_θ such that $\sup_{\theta \in \Theta} \|\theta\|_1 \leq C_\theta < \infty$. Let us denote C_θ -ball as $B := \{\theta \in \mathbb{R}^{d_\theta} \mid \|\theta\|_1 \leq C_\theta\}$. Then, the covering number of the parameter space $N_c(\delta, \Theta, L_1)$ is bounded by the covering number of the C_θ -ball $N_c(\delta, B, L_1)$. Applying Lemma A.5.3, we have:

$$N_c(\delta, B, L_1) \leq \frac{\text{vol}((1 + \frac{2}{\delta})B)}{\text{vol}(B)} = (1 + \frac{2}{\delta})^{d_\theta},$$

where the first inequality holds as the C_θ -ball is defined using the same metric as the covering number. Therefore,

$$N_c(\delta, \Theta, L_1) \leq (1 + \frac{2}{\delta})^{d_\theta}.$$

□

A.5 Results from Previous Literature

Lemma A.5.1. (Extreme Value Theorem) If f is continuous on a closed interval $[a, b]$, then f attains both an absolute maximum value and an absolute minimum value at some numbers in $[a, b]$.

Lemma A.5.2. (Berge's Maximum Theorem (Berge, 1963)) Let $X \subseteq \mathbb{R}^L$ and $Y \subseteq \mathbb{R}^K$, let $f : X \times Y \rightarrow \mathbb{R}$ be a continuous function and $\Gamma : X \rightarrow Y$ be a compact-valued and continuous correspondence. Then the function $v : X \rightarrow \mathbb{R}$ such that $v(x) = \sup_{y \in \Gamma(x)} f(x, y)$ is continuous.

Lemma A.5.3. (Volume ratios and Metric Entropy (Wainwright, 2019, §Lemma 5.7)) Consider a pair of norms $\|\cdot\|$ and $\|\cdot\|'$ on \mathbb{R}^d , and let B and B' be their corresponding unit balls (i.e., $B = \{\theta \in \mathbb{R}^d \mid \|\theta\| \leq 1\}$, with B' similarly defined). Then the δ -covering number of B in the $\|\cdot\|'$ -norm obeys the bounds

$$\left(\frac{1}{\delta}\right)^d \frac{\text{vol}(B)}{\text{vol}(B')} \leq N_c(\delta; B, \|\cdot\|') \leq \frac{\text{vol}(\frac{2}{\delta}B + B')}{\text{vol}(B')}.$$

Lemma A.5.4. (Hoeffding's Lemma) Let X be a random variable with $\mathbb{E}X = 0$, $a \leq X \leq b$. Then, for $s > 0$,

$$\mathbb{E}(e^{sX}) \leq e^{s^2(b-a)^2/8}.$$

Lemma A.5.5. (Upper bounds for Sub-Gaussian maxima) Let $\lambda > 0$, $n \geq 2$, and let Y_1, \dots, Y_n be real-valued random variables such that, for all $s > 0$ and $1 \leq i \leq n$, $\mathbb{E}(e^{sY_i}) \leq e^{\lambda^2 s^2/2}$ holds. Then,

$$(i) \quad \mathbb{E}(\max_{i \leq n} Y_i) \leq \lambda \sqrt{2 \ln n},$$

$$(ii) \quad \mathbb{E}(\max_{i \leq n} |Y_i|) \leq \lambda \sqrt{2 \ln(2n)} \leq 2\lambda \sqrt{\ln(n)}.$$

Lemma A.5.6 (Hoeffding's inequality ([Hoeffding, 1963](#))). *Let X_1, \dots, X_n be independent bounded random variables such that X_i falls in the interval $[a_i, b_i]$ with probability one. Denote their sum by $S_n = \sum_{i=1}^n X_i$. Then for any $\varepsilon > 0$ we have*

$$\Pr\{S_n - \mathbb{E} S_n \geq \varepsilon\} \leq e^{-2\varepsilon^2 / \sum_{i=1}^n (b_i - a_i)^2},$$

and

$$\Pr\{S_n - \mathbb{E} S_n \leq -\varepsilon\} \leq e^{-2\varepsilon^2 / \sum_{i=1}^n (b_i - a_i)^2}.$$

Bibliography

- ACEMOGLU, D., V. CHERNOZHUKOV, I. WERNING, AND M. D. WHINSTON (2020): “A multi-risk SIR model with optimally targeted lockdown,” Tech. rep., National Bureau of Economic Research.
- ADAMCZAK, R., M. KOTOWSKI, B. POLACZYK, AND M. STRZELECKI (2019): “A note on concentration for polynomials in the Ising model,” *Electron. J. Probab.*, 24, 1–22.
- ADJAH, C. AND T. CHRISTENSEN (2022): “Externally Valid Treatment Choice,” *arXiv preprint arXiv:2205.05561*.
- AGUIRREGABIRIA, V. AND P. MIRA (2007): “Sequential estimation of dynamic discrete games,” *Econometrica*, 75, 1–53.
- (2010): “Dynamic discrete choice structural models: A survey,” *Journal of Econometrics*, 156, 38–67.
- (2019): “Identification of games of incomplete information with multiple equilibria and unobserved heterogeneity,” *Quantitative Economics*, 10, 1659–1701.
- AKBARPOUR, M., C. COOK, A. MARZUOLI, S. MONGEY, A. NAGARAJ, M. SACCAROLA, P. TEBALDI, S. VASSERMAN, AND H. YANG (2020a): “Socioeconomic network heterogeneity and pandemic policy response,” Tech. rep., National Bureau of Economic Research.
- AKBARPOUR, M., S. MALLADI, AND A. SABERI (2020b): “Just a few seeds more: value of network information for diffusion,” *Available at SSRN 3062830*.
- ALVAREZ, F. E., D. ARGENTE, F. LIPPI, E. MÉNDEZ, AND D. VAN PATTEN (2023): “Strategic

- Complementarities in a Dynamic Model of Technology Adoption: P2P Digital Payments,” Tech. rep., National Bureau of Economic Research.
- ALVAREZ, F. E., F. LIPPI, AND T. SOUGANIDIS (2022): “Price setting with strategic complementarities as a mean field game,” Tech. rep., National Bureau of Economic Research.
- ANANTH, A. (2020a): “Optimal Treatment Assignment Rules on Networked Populations,” *working paper*.
- (2020b): “Optimal treatment assignment rules on networked populations,” Tech. rep., *working paper*.
- ANDREWS, I., T. KITAGAWA, AND A. MCCLOSKEY (2020): “Inference on winners,” *cemmap working paper CWP43/20*.
- ARCIDIACONO, P. AND R. A. MILLER (2011): “Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity,” *Econometrica*, 79, 1823–1867.
- ARONOW, P. M., C. SAMII, ET AL. (2017): “Estimating average causal effects under general interference, with application to a social network experiment,” *The Annals of Applied Statistics*, 11, 1912–1947.
- ATHEY, S., D. ECKLES, AND G. W. IMBENS (2018): “Exact p-values for network interference,” *Journal of the American Statistical Association*, 113, 230–240.
- ATHEY, S. AND S. WAGER (2020): “Policy learning with observational data,” *Econometrica*.
- (2021): “Policy learning with observational data,” *Econometrica*, 89, 133–161.
- ATKESON, A. (2020): “What will be the economic impact of covid-19 in the us? rough estimates of disease scenarios,” Tech. rep., National Bureau of Economic Research.
- BABICHENKO, Y. AND O. TAMUZ (2016): “Graphical potential games,” *Journal of Economic Theory*, 163, 889–899.
- BACH, F. (2011): “Learning with submodular functions: A convex optimization perspective,” *arXiv preprint arXiv:1111.6453*.
- BADEV, A. (2021): “Nash equilibria on (un) stable networks,” *Econometrica*, 89, 1179–1206.

- BAJARI, P., H. HONG, J. KRAINER, AND D. NEKIPELOV (2010a): “Estimating static models of strategic interactions,” *Journal of Business & Economic Statistics*, 28, 469–482.
- BAJARI, P., H. HONG, AND S. P. RYAN (2010b): “Identification and estimation of a discrete game of complete information,” *Econometrica*, 78, 1529–1568.
- BAKSHY, E., J. M. HOFMAN, W. A. MASON, AND D. J. WATTS (2011): “Everyone’s an influencer: quantifying influence on twitter,” in *Proceedings of the fourth ACM international conference on Web search and data mining*, 65–74.
- BALLESTER, C., A. CALVÓ-ARMENGOL, AND Y. ZENOU (2006): “Who’s who in networks. Wanted: The key player,” *Econometrica*, 74, 1403–1417.
- BANERJEE, A., A. G. CHANDRASEKHAR, E. DUFLO, AND M. O. JACKSON (2013): “The diffusion of microfinance,” *Science*, 341.
- BAQAEE, D., E. FARHI, M. J. MINA, AND J. H. STOCK (2020): “Reopening Scenarios,” Tech. rep., National Bureau of Economic Research.
- BASSE, G., A. FELLER, AND P. TOULIS (2019): “Randomization tests of causal effects under interference,” *Biometrika*, 106, 487–494.
- BERGE, C. (1963): *Topological Spaces*, translated by E. M. Patterson (Oliver & Boyd).
- BERGER, J. O. (1994): *The Robust Bayesian Viewpoint*, Robustness of Bayesian Analysis.
- BERTSEKAS, D. (2016): *Nonlinear Programming*, Athena scientific optimization and computation series, Athena Scientific.
- BESAG, J. (1974): “Spatial interaction and the statistical analysis of lattice systems,” *Journal of the Royal Statistical Society: Series B (Methodological)*, 36, 192–225.
- BHAMIDI, S., G. BRESLER, AND A. SLY (2008): “Mixing time of exponential random graphs,” in *2008 49th Annual IEEE Symposium on Foundations of Computer Science*, IEEE, 803–812.
- BHATTACHARYA, D. AND P. DUPAS (2012): “Inferring welfare maximizing treatment assignment under budget constraints,” *Journal of Econometrics*, 167, 168–196.
- BIAN, A. A., J. M. BUHMANN, A. KRAUSE, AND S. TSCHIATSCHEK (2017): “Guarantees for

- greedy maximization of non-submodular functions with applications,” in *International conference on machine learning*, PMLR, 498–507.
- BISHOP, C. M. AND N. M. NASRABADI (2006): *Pattern recognition and machine learning*, vol. 4, Springer.
- BLOOM, H. S., L. L. ORR, S. H. BELL, G. CAVE, F. DOOLITTLE, W. LIN, AND J. M. BOS (1997): “The benefits and costs of JTPA Title II-A programs: Key findings from the National Job Training Partnership Act study,” *Journal of human resources*, 549–576.
- BLUME, L. E. (1993): “The statistical mechanics of strategic interaction,” *Games and economic behavior*, 5, 387–424.
- BOUCHER, V. AND I. MOURIFIÉ (2017): “My friend far, far away: a random field approach to exponential random graph models,” *The econometrics journal*, 20, S14–S46.
- BRAMOULLÉ, Y. AND R. KRANTON (2007): “Public goods in networks,” *Journal of Economic theory*, 135, 478–494.
- BRAMOULLÉ, Y., R. KRANTON, AND M. D’AMOURS (2014): “Strategic interaction and networks,” *American Economic Review*, 104, 898–930.
- BRÉMAUD, P. (2013): *Markov chains: Gibbs fields, Monte Carlo simulation, and queues*, vol. 31, Springer Science & Business Media.
- BRESNAHAN, T. F. AND P. C. REISS (1991): “Empirical models of discrete games,” *Journal of Econometrics*, 48, 57–81.
- BROCK, W. A. AND S. N. DURLAUF (2001): “Discrete choice with social interactions,” *The Review of Economic Studies*, 68, 235–260.
- BROOKS, B. AND S. DU (2024): “On the structure of informationally robust optimal mechanisms,” *Econometrica*, 92, 1391–1438.
- BROUWER, L. E. J. (1911): “Über abbildung von mannigfaltigkeiten,” *Mathematische annalen*, 71, 97–115.
- BU, F., A. E. AIELLO, J. XU, AND A. VOLFOVSKY (2020): “Likelihood-based Inference for

- Partially Observed Epidemics on Dynamic Networks,” *Journal of the American Statistical Association*, 1–17.
- BULOW, J. I., J. D. GEANAKOPOLOS, AND P. D. KLEMPERER (1985): “Multimarket oligopoly: Strategic substitutes and complements,” *Journal of Political economy*, 93, 488–511.
- CANEN, N. AND K. SONG (2020): “A decomposition approach to counterfactual analysis in game-theoretic models,” *arXiv preprint arXiv:2010.08868*.
- CELA, E. (2013): *The quadratic assignment problem: theory and algorithms*, vol. 1, Springer Science & Business Media.
- CHAMBERLAIN, G. (2000a): “Econometric applications of maxmin expected utility,” *Journal of Applied Econometrics*, 15, 625–644.
- (2000b): “Econometrics and decision theory,” *Journal of Econometrics*, 95, 255–283.
- (2011): “Bayesian aspects of treatment choice,” *The Oxford Handbook of Bayesian Econometrics*, 11–39.
- (2020): “Robust decision theory and econometrics,” *Annual Review of Economics*, 12, 239–271.
- CHANDRASEKHAR, A. G. AND M. O. JACKSON (2014): “Tractable and consistent random graph models,” Tech. rep., National Bureau of Economic Research.
- CHATTERJEE, S. AND A. DEMBO (2016): “Nonlinear large deviations,” *Advances in Mathematics*, 299, 396–450.
- CHATTERJEE, S. AND P. DIACONIS (2013): “Estimating and understanding exponential random graph models,” *The Annals of Statistics*, 41, 2428–2461.
- CHEN, W., C. WANG, AND Y. WANG (2010): “Scalable influence maximization for prevalent viral marketing in large-scale social networks,” in *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*, 1029–1038.
- CHEN, X., M. LI, D. SIMCHI-LEVI, AND T. ZHAO (2020): “Allocation of COVID-19 Vaccines Under Limited Supply,” *Available at SSRN 3678986*.
- CHEN, Y. AND R. GAZZALE (2004): “When does learning in games generate convergence

- to Nash equilibria? The role of supermodularity in an experimental setting,” *American Economic Review*, 94, 1505–1535.
- CHERNOZHUKOV, V., J. C. ESCANCIANO, H. ICHIMURA, W. K. NEWEY, AND J. M. ROBINS (2022): “Locally robust semiparametric estimation,” *Econometrica*, 90, 1501–1535.
- CHESHER, A. AND A. M. ROSEN (2017): “Generalized instrumental variable models,” *Econometrica*, 85, 959–989.
- (2020): “Structural modeling of simultaneous discrete choice,” Tech. rep., cemmap Working paper.
- CHRISTAKIS, N., J. FOWLER, G. W. IMBENS, AND K. KALYANARAMAN (2020): “An empirical model for strategic network formation,” in *The Econometric Analysis of Network Data*, Elsevier, 123–148.
- CHRISTENSEN, T. AND B. CONNAULT (2023): “Counterfactual sensitivity and robustness,” *Econometrica*, 91, 263–298.
- CHRISTENSEN, T., H. R. MOON, AND F. SCHORFHEIDE (2022): “Optimal Discrete Decisions when Payoffs are Partially Identified,” *arXiv preprint arXiv:2204.11748*.
- CILIBERTO, F. AND E. TAMER (2009): “Market structure and multiple equilibria in airline markets,” *Econometrica*, 77, 1791–1828.
- CONFORTI, M. AND G. CORNUÉJOLS (1984): “Submodular set functions, matroids and the greedy algorithm: tight worst-case bounds and some generalizations of the Rado-Edmonds theorem,” *Discrete applied mathematics*, 7, 251–274.
- CUNNINGHAM, W. H. (1985): “Minimum cuts, modular functions, and matroid polyhedra,” *Networks*, 15, 205–215.
- DAS, A. AND D. KEMPE (2011): “Submodular meets spectral: Greedy algorithms for subset selection, sparse approximation and dictionary selection,” *arXiv preprint arXiv:1102.3975*.
- DE PAULA, Á. (2013): “Econometric analysis of games with multiple equilibria,” *Annual Review of Economics*, 5, 107–131.

- (2020): “Econometric models of network formation,” *Annual Review of Economics*, 12, 775–799.
- DE PAULA, A., I. RASUL, AND P. SOUZA (2024): “Identifying Network Ties from Panel Data: Theory and an Application to Tax Competition,” *Review of Economic Studies*, forthcoming.
- DE PAULA, Á., S. RICHARDS-SHUBIK, AND E. TAMER (2018): “Identifying preferences in networks with bounded degree,” *Econometrica*, 86, 263–288.
- DE PAULA, A. AND X. TANG (2012): “Inference of signs of interaction effects in simultaneous games with incomplete information,” *Econometrica*, 80, 143–172.
- (2020): “Testable implications of multiple equilibria in discrete games with correlated types,” *arXiv preprint arXiv:2012.00787*.
- DEHEJIA, R. H. (2005): “Program evaluation as a decision problem,” *Journal of Econometrics*, 125, 141–173.
- DENNEBERG, D. (1994): *Non-additive measure and integral*, vol. 27, Springer Science & Business Media.
- DOBRUSHIN, R. L. (1970): “Prescribing a system of random variables by conditional distributions,” *Theory of Probability & Its Applications*, 15, 458–486.
- DOMINGOS, P. AND M. RICHARDSON (2001): “Mining the network value of customers,” in *Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining*, 57–66.
- ECHENIQUE, F. AND I. KOMUNJER (2009): “Testing models with multiple equilibria by quantile methods,” *Econometrica*, 77, 1281–1297.
- EL HALABI, M. AND S. JEGELKA (2020): “Optimal approximation for unconstrained non-submodular minimization,” in *International Conference on Machine Learning*, PMLR, 3961–3972.
- ELLISON, G. (2020): “Implications of heterogeneous SIR models for analyses of COVID-19,” Tech. rep., National Bureau of Economic Research.
- EMANUEL, E. J., G. PERSAD, A. KERN, A. BUCHANAN, C. FABRE, D. HALLIDAY, J. HEATH,

- L. HERZOG, R. LELAND, E. T. LEMANGO, ET AL. (2020): “An ethical framework for global vaccine allocation,” *Science*, 369, 1309–1312.
- FERGUSON, N., D. LAYDON, G. NEDJATI-GILANI, N. IMAI, K. AINSLIE, M. BAGUELIN, S. BHATTIA, A. BOONYASIRI, Z. CUCUNUBÁ, G. CUOMO-DANNENBURG, ET AL. (2020): “Report 9: Impact of non-pharmaceutical interventions (NPIs) to reduce COVID19 mortality and healthcare demand,” *Imperial College London*, 10, 77482.
- FERNANDEZ, A. A., J. L. M. OLEA, C. QIU, J. STOYE, AND S. TINDA (2024): “Robust Bayes Treatment Choice with Partial Identification,” *working paper*.
- FERNÁNDEZ-VILLAVERDE, J. AND C. I. JONES (2020): “Estimating and Simulating a SIRD Model of COVID-19 for Many Countries, States, and Cities,” Tech. rep., National Bureau of Economic Research.
- FISHER, M. L., G. L. NEMHAUSER, AND L. A. WOLSEY (1978): “An analysis of approximations for maximizing submodular set functions—II,” in *Polyhedral combinatorics*, Springer, 73–87.
- FU, H., N. HAGHPANAH, J. HARTLINE, AND R. KLEINBERG (2021): “Full surplus extraction from samples,” *Journal of Economic Theory*, 193, 105230.
- FUDENBERG, D. AND D. K. LEVINE (1998): *The theory of learning in games*, vol. 2, MIT press.
- GALEOTTI, A., B. GOLUB, AND S. GOYAL (2020): “Targeting interventions in networks,” *Econometrica*, 88, 2445–2471.
- GALEOTTI, A., S. GOYAL, M. O. JACKSON, F. VEGA-REDONDO, AND L. YARIV (2010): “Network games,” *The review of economic studies*, 77, 218–244.
- GALEOTTI, A. AND B. W. ROGERS (2013): “Strategic immunization and group structure,” *American Economic Journal: Microeconomics*, 5, 1–32.
- GEMAN, S. AND D. GEMAN (1984): “Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images,” *IEEE Transactions on pattern analysis and machine intelligence*, 721–741.
- GEYER, C. J. AND E. A. THOMPSON (1992): “Constrained Monte Carlo maximum likelihood

- for dependent data,” *Journal of the Royal Statistical Society: Series B (Methodological)*, 54, 657–683.
- GIACOMINI, R. AND T. KITAGAWA (2021): “Robust Bayesian inference for set-identified models,” *Econometrica*, 89, 1519–1556.
- GIACOMINI, R., T. KITAGAWA, AND M. READ (2021): “Robust Bayesian analysis for econometrics,” *CEPR Discussion Paper No. DP16488*.
- GILBOA, I. (2009): *Theory of decision under uncertainty*, vol. 45, Cambridge university press.
- GOLDENBERG, J., B. LIBAI, AND E. MULLER (2001): “Talk of the network: A complex systems look at the underlying process of word-of-mouth,” *Marketing letters*, 12, 211–223.
- GOLDSMITH-PINKHAM, P. AND G. W. IMBENS (2013): “Social networks and the identification of peer effects,” *Journal of Business & Economic Statistics*, 31, 253–264.
- GONÇALVES, D. AND B. FURTADO (2020): “Statistical mechanism design: robust pricing and reliable projections,” *working paper*.
- GÖTZE, F., H. SAMBALE, AND A. SINULIS (2019): “Higher order concentration for functions of weakly dependent random variables,” *Electron. J. Probab*, 24, 1–19.
- GRAHAM, B. AND A. DE PAULA (2020): *The Econometric Analysis of Network Data*, Academic Press.
- GRAHAM, B. S. AND A. PELICAN (2023): “Scenario Sampling for Large Supermodular Games,” Tech. rep., National Bureau of Economic Research.
- GRANOVETTER, M. (1978): “Threshold models of collective behavior,” *American journal of sociology*, 83, 1420–1443.
- GRIECO, P. L. (2014): “Discrete games with flexible information structures: An application to local grocery markets,” *The RAND Journal of Economics*, 45, 303–340.
- GU, J., T. RUSSELL, AND T. STRINGHAM (2022): “Counterfactual Identification and Latent Space Enumeration in Discrete Outcome Models,” *Available at SSRN 4188109*.
- GUREN, A. M. (2018): “House price momentum and strategic complementarity,” *Journal of Political Economy*, 126, 1172–1218.

- HAMMERSLEY, J. AND P. CLIFFORD (1971): "Markov fields on finite graphs and lattices," *Unpublished manuscript*.
- HANNAN, J. (1957): "APPROXIMATION TO BAYES RISK IN REPEATED PLAY," *Contributions to the Theory of Games*, 21, 97.
- HANSEN, L. P. AND T. J. SARGENT (2001): "Robust control and model uncertainty," *American Economic Review*, 91, 60–66.
- (2008): *Robustness*, Princeton university press.
- HARSANYI, J. C. (1967): "Games with incomplete information played by "Bayesian" players, I–III Part I. The basic model," *Management science*, 14, 159–182.
- HIRANO, K. AND J. R. PORTER (2009): "Asymptotics for statistical treatment rules," *Econometrica*, 77, 1683–1701.
- (2020): "Asymptotic analysis of statistical decision rules in econometrics," in *Handbook of Econometrics*, Elsevier, vol. 7, 283–354.
- HOEFFDING, W. (1963): "Probability inequalities for sums of bounded random variables," *Journal of the American Statistical Association*, 58, 13–30.
- HOTZ, V. J. AND R. A. MILLER (1993): "Conditional choice probabilities and the estimation of dynamic models," *The Review of Economic Studies*, 60, 497–529.
- HSIAO, A. (2022): "Educational Investment in Spatial Equilibrium: Evidence from Indonesia," Tech. rep., Working Paper.
- ISING, E. (1925): "Beitrag zur Theorie des Ferromagnetismus," *Z. Phys*, 31, 253–258.
- JACKSON, M. O. (2010): *Social and economic networks*, Princeton university press.
- JACKSON, M. O. AND A. WATTS (2002): "The evolution of social and economic networks," *Journal of economic theory*, 106, 265–295.
- JACKSON, M. O. AND Y. ZENOU (2015): "Games on networks," in *Handbook of game theory with economic applications*, Elsevier, vol. 4, 95–163.
- JACKSON, M. O. ET AL. (2008): *Social and economic networks*, vol. 3, Princeton university press Princeton.

- JAGALUR-MOHAN, J. AND Y. M. MARZOUK (2021): “Batch greedy maximization of non-submodular functions: Guarantees and applications to experimental design.” *J. Mach. Learn. Res.*, 22, 252–1.
- JIA, P. (2008): “What happens when Wal-Mart comes to town: An empirical analysis of the discount retailing industry,” *Econometrica*, 76, 1263–1316.
- JOSEPH, R. AND S. H. MARK (2022): “On Approximate Variational Estimation for a Model of Network Formation,” *working paper*.
- JOVANOVIĆ, B. (1989): “Observable implications of models with multiple equilibria,” *Econometrica: Journal of the Econometric Society*, 1431–1437.
- KAIDO, H. AND Y. ZHANG (2023): “Applications of Choquet expected utility to hypothesis testing with incompleteness,” *The Japanese Economic Review*, 1–22.
- KAKADE, S., M. KEARNS, J. LANGFORD, AND L. ORTIZ (2003): “Correlated equilibria in graphical games,” in *Proceedings of the 4th ACM Conference on Electronic Commerce*, 42–47.
- KASHAEV, N., N. LAZZATI, AND R. XIAO (2023): “Peer Effects in Consideration and Preferences,” *arXiv preprint arXiv:2310.12272*.
- KASY, M. AND A. SAUTMANN (2019): “Adaptive treatment assignment in experiments for policy choice,” *CESifo Working Paper*.
- (2021): “Adaptive treatment assignment in experiments for policy choice,” *Econometrica*, 89, 113–132.
- KEANE, M. AND N. WASI (2013): “Comparing alternative models of heterogeneity in consumer choice behavior,” *Journal of Applied Econometrics*, 28, 1018–1045.
- KEARNS, M., M. L. LITTMAN, AND S. SINGH (2013): “Graphical models for game theory,” *arXiv preprint arXiv:1301.2281*.
- KEELING, M. J. AND P. ROHANI (2011): *Modeling infectious diseases in humans and animals*, Princeton University Press.
- KEMPE, D., J. KLEINBERG, AND É. TARDOS (2003): “Maximizing the spread of influence

- through a social network,” in *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, 137–146.
- KERMACK, W. O. AND A. G. MCKENDRICK (1927): “A contribution to the mathematical theory of epidemics,” *Proceedings of the royal society of london. Series A, Containing papers of a mathematical and physical character*, 115, 700–721.
- KITAGAWA, T., H. LOPEZ, AND J. ROWLEY (2022a): “Stochastic treatment choice with empirical welfare updating,” *arXiv preprint arXiv:2211.01537*.
- KITAGAWA, T., S. SAKAGUCHI, AND A. TETENOV (2021): “Constrained Classification and Policy Learning,” *arXiv preprint arXiv:2106.12886*.
- KITAGAWA, T. AND A. TETENOV (2018): “Who should be treated? empirical welfare maximization methods for treatment choice,” *Econometrica*, 86, 591–616.
- KITAGAWA, T. AND G. WANG (2023a): “Individualized Treatment Allocation in Sequential Network Games,” *arXiv preprint arXiv:2302.05747*.
- (2023b): “Who should get vaccinated? Individualized allocation of vaccines over SIR network,” *Journal of Econometrics*, 232, 109–131.
- KITAGAWA, T., W. WANG, AND M. XU (2022b): “Policy choice in time series by empirical welfare maximization,” *arXiv preprint arXiv:2205.03970*.
- KLINE, B., A. PAKES, AND E. TAMER (2021): “Moment inequalities and partial identification in industrial organization,” in *Handbook of Industrial Organization, Volume 4*, Elsevier, 345–431.
- KLINE, B. AND E. TAMER (2020): “Econometric analysis of models with social interactions,” in *The Econometric Analysis of Network Data*, Elsevier, 149–181.
- KOCK, A. B., D. PREINERSTORFER, AND B. VELIYEV (2020): “Functional sequential treatment allocation,” *Journal of the American Statistical Association*, 1–36.
- KOJEVNIKOV, D., V. MARMER, AND K. SONG (2021): “Limit theorems for network dependent random variables,” *Journal of Econometrics*, 222, 882–908.

- KOOPMANS, T. C. AND M. BECKMANN (1957): “Assignment problems and the location of economic activities,” *Econometrica: journal of the Econometric Society*, 53–76.
- KOR, R. AND J. ZHOU (2022): “Welfare and distributional effects of joint intervention in networks,” *arXiv preprint arXiv:2206.03863*.
- KOROLEV, I. (2021): “Identification and Estimation of the SEIRD Epidemic Model for COVID-19,” *Journal of econometrics*, 220, 63–85.
- KRAUSE, A. AND D. GOLOVIN (2014): “Submodular function maximization.” .
- KRAUSE, A., A. SINGH, AND C. GUESTRIN (2008): “Near-optimal sensor placements in Gaussian processes: Theory, efficient algorithms and empirical studies.” *Journal of Machine Learning Research*, 9.
- LAWRENCE, N., M. SEEGER, AND R. HERBRICH (2002): “Fast sparse Gaussian process methods: The informative vector machine,” *Advances in neural information processing systems*, 15.
- LAZZATI, N. (2015): “Treatment response with social interactions: Partial identification via monotone comparative statics,” *Quantitative Economics*, 6, 49–83.
- LEE, L.-F., X. LIU, E. PATACCHINI, AND Y. ZENOU (2020): “Who is the key player? A network analysis of juvenile delinquency,” *Journal of Business & Economic Statistics*, 1–9.
- (2021): “Who is the key player? A network analysis of juvenile delinquency,” *Journal of Business & Economic Statistics*, 39, 849–857.
- LEE, R. S. AND A. PAKES (2009): “Multiple equilibria and selection by learning in an applied setting,” *Economics Letters*, 104, 13–16.
- LEUNG, M. P. (2015): “Two-step estimation of network-formation models with incomplete information,” *Journal of Econometrics*, 188, 182–195.
- (2019): “Inference in models of discrete choice with social interactions using network data,” *arXiv preprint arXiv:1911.07106*.
- (2020): “Equilibrium computation in discrete network games,” *Quantitative Economics*, 11, 1325–1347.

- LEUNG, M. P. AND H. R. MOON (2019): “Normal Approximation in Large Network Models,” *Available at SSRN 3377709*.
- LI, M. Y. AND J. S. MULDOWNNEY (1995): “Global stability for the SEIR model in epidemiology,” *Mathematical biosciences*, 125, 155–164.
- LI, S. AND S. WAGER (2020): “Random Graph Asymptotics for Treatment Effect Estimation under Network Interference,” *arXiv preprint arXiv:2007.13302*.
- LOIOLA, E. M., N. M. M. DE ABREU, P. O. BOAVENTURA-NETTO, P. HAHN, AND T. QUERIDO (2007): “A survey for the quadratic assignment problem,” *European journal of operational research*, 176, 657–690.
- LOVÁSZ, L. (1983): “Submodular functions and convexity,” *Mathematical Programming The State of the Art: Bonn 1982*, 235–257.
- LUGOSI, G. (2002): “Pattern classification and learning theory,” in *Principles of nonparametric learning*, Springer, 1–56.
- MAIER, B. F., A. BURDINSKI, A. H. ROSE, F. SCHLOSSER, D. HINRICHS, C. BETSCH, L. KORN, P. SPRENGHOLZ, M. MEYER-HERMANN, T. MITRA, ET AL. (2021): “Potential benefits of delaying the second mRNA COVID-19 vaccine dose,” *arXiv preprint arXiv:2102.13600*.
- MANSKI, C. F. (1993): “Identification of endogenous social effects: The reflection problem,” *The review of economic studies*, 60, 531–542.
- (2003): *Partial identification of probability distributions*, Springer Science & Business Media.
- (2004): “Statistical treatment rules for heterogeneous populations,” *Econometrica*, 72, 1221–1246.
- (2010): “Vaccination with partial knowledge of external effectiveness,” *Proceedings of the National Academy of Sciences*, 107, 3953–3960.
- (2013): “Identification of treatment response with social interactions,” *The Econometrics Journal*, 16, S1–S23.

- (2017): “Mandating vaccination with unknown indirect effects,” *Journal of Public Economic Theory*, 19, 603–619.
- (2019): “Treatment choice with trial data: statistical decision theory should supplant hypothesis testing,” *The American Statistician*, 73, 296–304.
- MANSKI, C. F. AND F. MOLINARI (2020): “Estimating the COVID-19 infection rate: Anatomy of an inference problem,” *Journal of Econometrics*.
- MANSKI, C. F. AND A. TETENOV (2021): “Statistical Decision Properties of Imprecise Trials Assessing Coronavirus Disease 2019 (COVID-19) Drugs,” *Value in Health*, 24, 641–647.
- MATHEVET, L. (2010): “Supermodular mechanism design,” *Theoretical Economics*, 5, 403–443.
- MTAKOP, E. AND M. TABORD-MEEHAN (2021): “Model selection for treatment choice: Penalized welfare maximization,” *Econometrica*, 89, 825–848.
- MELE, A. (2017): “A structural model of dense network formation,” *Econometrica*, 85, 825–850.
- MELE, A. AND L. ZHU (2022): “Approximate Variational Estimation for a Model of Network Formation,” *The Review of Economics and Statistics*, 1–12.
- MENZEL, K. (2016): “Inference for games with many players,” *The Review of Economic Studies*, 83, 306–337.
- METROPOLIS, N. AND S. ULAM (1949): “The monte carlo method,” *Journal of the American statistical association*, 44, 335–341.
- MIAO, Y. (2010): “Concentration inequality of maximum likelihood estimator,” *Applied mathematics letters*, 23, 1305–1309.
- MILGROM, P. AND J. ROBERTS (1990): “Rationalizability, learning, and equilibrium in games with strategic complementarities,” *Econometrica: Journal of the Econometric Society*, 1255–1277.
- (1991): “Adaptive and sophisticated learning in normal form games,” *Games and economic Behavior*, 3, 82–100.

- MILGROM, P. AND C. SHANNON (1994): “Monotone comparative statics,” *Econometrica: Journal of the Econometric Society*, 157–180.
- MOLINARI, F. (2020): “Microeconometrics with partial identification,” *Handbook of econometrics*, 7, 355–486.
- MOLINARI, F. AND A. ROSEN (2008): “The identification power of equilibrium in games: The supermodular case,” *Journal of Business & Economic Statistics*, 26, 297–302.
- MONDERER, D. AND L. S. SHAPLEY (1996): “Potential games,” *Games and economic behavior*, 14, 124–143.
- MORRIS, S. (2000): “Contagion,” *The Review of Economic Studies*, 67, 57–78.
- MORRIS, S., D. OYAMA, AND S. TAKAHASHI (2024): “Implementation via Information Design in Binary-Action Supermodular Games,” *Econometrica*, 92, 775–813.
- MUNRO, E. (2024): “Treatment allocation with strategic agents,” *Management Science*.
- MUNRO, E., S. WAGER, AND K. XU (2021): “Treatment effects in market equilibrium,” *arXiv preprint arXiv:2109.11647*.
- (2023): “Treatment effects in market equilibrium,” *arXiv preprint arXiv:2109.11647*.
- NAKAJIMA, R. (2007): “Measuring peer effects on youth smoking behaviour,” *The Review of Economic Studies*, 74, 897–935.
- NÅSELL, I. (1996): “The quasi-stationary distribution of the closed endemic SIS model,” *Advances in Applied Probability*, 895–932.
- NEMHAUSER, G. L., L. A. WOLSEY, AND M. L. FISHER (1978): “An analysis of approximations for maximizing submodular set functions—I,” *Mathematical programming*, 14, 265–294.
- PARENTE, S. L. AND E. C. PRESCOTT (1994): “Barriers to technology adoption and development,” *Journal of political Economy*, 102, 298–321.
- PARISE, F. AND O. ASUMAN (2023): “Graphon games: A statistical framework for network games and interventions,” *Econometrica*, 91, 191–225.
- PASTOR-SATORRAS, R. AND A. VESPIGNANI (2002): “Immunization of complex networks,” *Physical review E*, 65, 036104.

- PESARAN, H. AND C. F. YANG (2020): “Matching Theory and Evidence on COVID-19 using a Stochastic Model of Epidemics on Networks,” *IAAE Webinar*.
- RICHARDSON, M. AND P. DOMINGOS (2002): “Mining knowledge-sharing sites for viral marketing,” in *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*, 61–70.
- RIDDER, G. AND S. SHENG (2020): “Two-step estimation of a strategic network formation model with clustering,” *arXiv preprint arXiv:2001.03838*.
- ROSENTHAL, R. W. (1973): “A class of games possessing pure-strategy Nash equilibria,” *International Journal of Game Theory*, 2, 65–67.
- ROUGHGARDEN, T. (2010): “Algorithmic game theory,” *Communications of the ACM*, 53, 78–86.
- RUBIN, D. B. (1974): “Estimating causal effects of treatments in randomized and nonrandomized studies,” *Journal of educational Psychology*, 66, 688.
- SAKAGUCHI, S. (2021): “Estimation of Optimal Dynamic Treatment Assignment Rules under Policy Constraint,” *arXiv preprint arXiv:2106.05031*.
- SAVAGE, L. J. (1951): “The theory of statistical decision,” *Journal of the American Statistical association*, 46, 55–67.
- SÄVJE, F., P. M. ARONOW, AND M. G. HUDGENS (2017): “Average treatment effects in the presence of unknown interference,” *arXiv preprint arXiv:1711.06399*.
- SEGAL, I. (2003): “Coordination and discrimination in contracting with externalities: Divide and conquer?” *Journal of Economic Theory*, 113, 147–181.
- SHENG, S. (2020): “A structural econometric analysis of network formation games through subnetworks,” *Econometrica*, 88, 1829–1858.
- SNIJDERS, T. A. ET AL. (2002): “Markov chain Monte Carlo estimation of exponential random graph models,” *Journal of Social Structure*, 3, 1–40.
- STOCK, J. H. (2020): “Data gaps and the policy response to the novel coronavirus,” Tech. rep., National Bureau of Economic Research.

- STOYE, J. (2009): “Minimax regret treatment choice with finite samples,” *Journal of Econometrics*, 151, 70–81.
- (2012): “Minimax regret treatment choice with covariates or with limited validity of experiments,” *Journal of Econometrics*, 166, 138–156.
- (2021): “Bounding Infection Prevalence by Bounding Selectivity and Accuracy of Tests: With Application to Early COVID-19,” *arXiv preprint arXiv:2008.06178*.
- SUN, L. (2021): “Empirical welfare maximization with constraints,” *arXiv preprint arXiv:2103.15298*.
- SUN, Y., W. ZHAO, AND J. ZHOU (2023): “Structural interventions in networks,” *International economic review*, 64, 1533–1563.
- TAKAHASHI, S. (2008): “The number of pure Nash equilibria in a random game with nondecreasing best responses,” *Games and Economic Behavior*, 63, 328–340.
- TAMER, E. (2003): “Incomplete simultaneous discrete response model with multiple equilibria,” *The Review of Economic Studies*, 70, 147–165.
- TARSKI, A. (1955): “A LATTICE-THEORETICAL FIXPOINT THEOREM AND ITS APPLICATIONS,” *Pacific J. Math*, 5, 285–309.
- TETENOV, A. (2012): “Statistical treatment choice based on asymmetric minimax regret criteria,” *Journal of Econometrics*, 166, 157–165.
- TOPKIS, D. M. (1978): “Minimizing a submodular function on a lattice,” *Operations research*, 26, 305–321.
- (1979): “Equilibrium points in nonzero-sum n-person submodular games,” *Siam Journal on control and optimization*, 17, 773–787.
- (1998): *Supermodularity and complementarity*, Princeton university press.
- TRAIN, K. E., D. L. MCFADDEN, AND M. BEN-AKIVA (1987): “The demand for local telephone service: A fully discrete model of residential calling patterns and service choices,” *The RAND Journal of Economics*, 109–123.
- TUITE, A. R., L. ZHU, D. N. FISMAN, AND J. A. SALOMON (2021): “Alternative dose alloca-

- tion strategies to increase benefits from constrained COVID-19 vaccine supply,” *Annals of internal medicine*.
- VAN DUIJN, M. A., K. J. GILE, AND M. S. HANDCOCK (2009): “A framework for the comparison of maximum pseudo-likelihood and maximum likelihood estimation of exponential family random graph models,” *Social networks*, 31, 52–62.
- VAZQUEZ-BARE, G. (2020): “Causal spillover effects using instrumental variables,” *arXiv preprint arXiv:2003.06023*.
- VIVIANO, D. (2019): “Policy targeting under network interference,” *arXiv preprint arXiv:1906.10258*.
- (2020): “Policy design in experiments with unknown interference,” *arXiv preprint arXiv:2011.08174*.
- (2024): “Policy targeting under network interference,” *Review of Economic Studies*, rdae041.
- WAINWRIGHT, M. J. (2019): *High-dimensional statistics: A non-asymptotic viewpoint*, vol. 48, Cambridge university press.
- WAINWRIGHT, M. J., M. I. JORDAN, ET AL. (2008): “Graphical models, exponential families, and variational inference,” *Foundations and Trends® in Machine Learning*, 1, 1–305.
- WANG, C., N. KOMODAKIS, AND N. PARAGIOS (2013): “Markov random field modeling, inference & learning in computer vision & image understanding: A survey,” *Computer Vision and Image Understanding*, 117, 1610–1627.
- WANG, X., Z. DU, K. JOHNSON, S. FOX, M. LACHMANN, J. S. MCLELLAN, AND L. A. MEYERS (2021): “The impacts of COVID-19 vaccine timing, number of doses, and risk prioritization on mortality in the US,” *medRxiv*.
- WANG, Y., C. GU, AND T. OTSU (2024): “Graph Neural Networks: Theory for Estimation with Application on Network Heterogeneity,” *arXiv preprint arXiv:2401.16275*.
- XU, H. (2018): “Social interactions in large networks: A game theoretic approach,” *International Economic Review*, 59, 257–284.

- ZHAO, Y., D. ZENG, A. J. RUSH, AND M. R. KOSOROK (2012): “Estimating individualized treatment rules using outcome weighted learning,” *Journal of the American Statistical Association*, 107, 1106–1118.
- ZHAO, Y.-Q., D. ZENG, E. B. LABER, AND M. R. KOSOROK (2015): “New statistical learning methods for estimating optimal dynamic treatment regimes,” *Journal of the American Statistical Association*, 110, 583–598.
- ZHOU, Z., S. ATHEY, AND S. WAGER (2018): “Offline multi-action policy learning: Generalization and optimization,” *arXiv preprint arXiv:1810.04778*.