

Robust Network Targeting with Multiple Nash Equilibria[†]

Guanyi Wang[‡]

December 2, 2024

[Latest version here](#)

Abstract

Many policy problems involve designing individualized treatment allocation rules to maximize the equilibrium social welfare of interacting agents. Focusing on large-scale simultaneous decision games with strategic complementarities, we develop a method to estimate an optimal treatment allocation rule that is robust to the presence of multiple equilibria. Our approach remains agnostic about changes in the equilibrium selection mechanism under counterfactual policies, and we provide a closed-form expression for the boundary of the identified set of equilibrium outcomes. To address the incompleteness that arises when an equilibrium selection mechanism is not specified, we use a maximin welfare criterion to select a policy based on the “least favourable” equilibrium outcome, and implement this policy using a greedy algorithm. We establish performance guarantees for our method by deriving a welfare regret bound, which accounts for sampling uncertainty and the use of a greedy algorithm. We demonstrate our method with an application to the microfinance dataset of [Banerjee et al. \(2013\)](#).

Keywords: Treatment choice, simultaneous games, incomplete model, strategic complementarity, multiple Nash equilibria.

[†]I am especially grateful to Tim Christensen, Aureo de Paula, Toru Kitagawa, and Andrei Zelenev for their continuous support, and Abhishek Ananth, Dmitry Arkhangelsky, Karun Adusumilli, Miguel Ballester, Andrew Chesher, Pedro Carneiro, Ben Deaner, James Duffy, Hugo Freeman, Andrea Galeotti, Duarte Gonçalves, Francois Gerard, Joao Granja, Raffaella Giacomini, Wayne Gao, Chen-Wei Hsiang, Hu Hao, Keisuke Hirano, Sukjin Han, Stephen Hansen, Guido Imbens, Hidehiko Ichimura, Hua Jin, Ziyu Jiang, Dennis Kristensen, Ivana Komunjer, Frank Kleibergen, Maximilian Kasy, Nathan Canen, Soonwoo Kwon, Daniel Lewis, Jessie Li, Michael Leung, Yikai Li, Anna Mikusheva, Chuck Manski, Konrad Menzel, Robert Miller, Sophocles Mavroeidis, Lars Nesheim, Whitney Newey, Guillaume Pouliot, Martin Pesendorfer, Adam Rosen, Imran Rasul, Jonathan Roth, Jeff Rowley, Liyang Sun, Jörg Stoye, Shuyang Sheng, Xiaoxia Shi, Yuya Sasaki, Gabriel Ulyssea, Davide Viviano, Kaspar Wuthrich, Yike Wang, Mengshan Xu, Tian Xie, Weisheng Zhang, Yanziyi Zhang, and Yuanqi Zhang for beneficial comments and discussions.

[‡]CeMMAP and Department of Economics, University College London. Email: guanyi.wang.17@ucl.ac.uk

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 pro-

¹This incomplete information setting (Brock and Durlauf, 2001; Bajari et al., 2010a; de Paula and Tang, 2012), is our primary focus. Section 7 extends our results to the complete information setting.

vision 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. Without 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 allocation, which requires optimizing an objective function dependent on a system of

²See the detailed surveys by [De Paula \(2013\)](#); [Molinari \(2020\)](#); [Kline and Tamer \(2020\)](#).

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 4.1.2.

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 performance, 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.

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 incomplete 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 (2020), 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. (2022), 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 (2020), 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 (2020) 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 2 introduces the game setting and the solution concept. Section 3 discusses counterfactual analysis. Section 4 focuses on treatment allocation and implementation. Section 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 6. Section 7 extends our analysis to the complete information setting. Section 8 concludes. All proofs and derivations are shown in Appendix B to Appendix E.

2 Model

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 2.3 and Section 4.1.1.

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 1. *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*:

$$W_{X,G}(D) = \frac{1}{N} \sum_{i=1}^N \mathbb{E}_\varepsilon [g_i(Y, X, D, G)|X, G], \quad (1)$$

subject to the capacity constraint κ (i.e., $\sum_{i=1}^N D_i \leq \kappa$). The expectation in Eq.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 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 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:

$$Y_i = \mathbb{1} \left\{ U_i(1, Y_{-i}, X, D, G) \geq 0 \right\}, \quad \forall i \in \mathcal{N}.$$

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 2.2.

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,

$$Y_i = \mathbf{1} \left\{ \mathbb{E}_\varepsilon [U_i(1, Y_{-i}, X, D, G) | X, D, G, \varepsilon_i] \geq 0 \right\}, \quad \forall i \in \mathcal{N}. \quad (2)$$

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 1, this can be simplified to:

$$\mathbb{E}_\varepsilon [U_i(1, Y_{-i}, X, D, G) | X, D, G, \varepsilon_i] = \alpha_i + \sum_{j \neq i} \beta_{ij} \mathbb{E}_\varepsilon [Y_j | X, D, G] - \varepsilon_i. \quad (3)$$

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* \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.

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 2.1. (Pure Strategy Bayesian Nash equilibrium) Let \mathcal{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

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 $y \leq y''$.

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 $y'_i \in \mathbf{Y}$, where $\mathbb{E}_\varepsilon[U_i(\cdot)|X, D, G, \varepsilon_i]$ is defined as in Eq.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:

$$\sigma_i(X, D, G) := \mathbb{E}_\varepsilon[Y_i|X, D, G]. \quad (4)$$

Combining the specification of Y_i (Eq.2 and Eq.3) with Eq.4, we have:

$$\sigma_i(X, D, G) = \int \mathbb{1}\left\{\alpha_i + \sum_{j \neq i} \beta_{ij} \sigma_j(X, D, G) \geq \varepsilon_i\right\} dF_\varepsilon. \quad (5)$$

Let Ω be a mapping from $[0, 1]^N$ to $[0, 1]^N$ that collects Eq.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:

$$\sigma^*(X, D, G) = \Omega(\sigma^*(X, D, G)). \quad (6)$$

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:

$$y_i^* = \mathbb{1}\left\{\alpha_i + \sum_{j \neq i} \beta_{ij} \sigma_j^*(X, D, G) \geq \varepsilon_i\right\}, \quad \forall i \in \mathcal{N}. \quad (7)$$

As the right hand side of Eq.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 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

$$\Pr[Y_i = 1|X, D, G, \lambda] = \sum_{\sigma^* \in \Sigma} \lambda(\sigma^*|X, D, G)\sigma_i^*, \quad \forall i \in \mathcal{N}. \quad (8)$$

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. 8. To improve the tractability of counterfactual analysis, we focus on a game with strategic complementarity.

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 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 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 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.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, we show how strategic complementarity simplifies counterfactual analysis.

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:

$$W_{X,G}(D) = \frac{1}{N} \sum_{i=1}^N \mathbb{E}_\varepsilon [g_i(Y, X, D, G) | X, G]. \quad (9)$$

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.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

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.

(Bloom et al., 1997). In such scenarios, social welfare is defined as:

$$W_{X,G,\lambda}(D) = \frac{1}{N} \sum_{i=1}^N \mathbb{E}_{\varepsilon} [U_i(Y, X, D, G) + \varepsilon_i Y_i | X, D, G, \lambda], \quad (10)$$

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 information settings). However, exact bounds of social welfare cannot be directly obtained

⁷Brock and Durlauf (2001, Section 4) show that introducing a shock term in Eq.10 would render the model analytically intractable.

from the bounds of the CCP. This is because:

$$\inf_{\lambda \in \Lambda} \sum_{i=1}^N \Pr(Y_i = 1 | X, D, G, \lambda) \leq \sum_{i=1}^N \inf_{\lambda \in \Lambda} \Pr(Y_i = 1 | X, D, G, \lambda), \quad (11)$$

and

$$\sup_{\lambda \in \Lambda} \sum_{i=1}^N \Pr(Y_i = 1 | X, D, G, \lambda) \geq \sum_{i=1}^N \sup_{\lambda \in \Lambda} \Pr(Y_i = 1 | X, D, G, \lambda). \quad (12)$$

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.11 and Eq.12 hold with equality in a supermodular game. Formally,

Theorem 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 3.1 is provided in Appendix C.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 3.1. (Utilitarian Welfare at Equilibrium) *Under Assumption 1, 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 3.1 is provided in Appendix B.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 3.1. 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 4.1.3. 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.

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:

$$D^* = \arg \max_{D \in \mathcal{D}_\kappa} \min_{\lambda \in \Lambda} W_{X,G,\lambda}(D). \quad (13)$$

Recall that, by Theorem 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.

4.1 Implementation

4.1.1 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 4.1.2.

The discussion of counterfactual analysis in Section 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.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:

$$\sigma_i^{data} = \int \mathbb{1}\left\{\varepsilon_i^n \leq \alpha_i + \sum_{j \neq i} \beta_{ij} \sigma_j^{data}\right\} dF_{\varepsilon^n} = F_{\varepsilon^n}\left[\alpha_i + \sum_{j \neq i} \beta_{ij} \sigma_j^{data}\right]. \quad (14)$$

Taking the inverse of the CDF of ε on both sides in Eq.14 yields:

$$F_{\varepsilon^n}^{-1}(\sigma_i^{data}) = \alpha_i + \sum_{j \neq i} \beta_{ij} \sigma_j^{data}. \quad (15)$$

Even assuming that the equilibrium CCP profile in the training data is observable, identifying all the parameters in Eq.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.15, only provides information about N scalars. Given these limitations, we define our utility function as follows to ensure identifiability and allow for the

⁸This approach builds on [Hotz and Miller \(1993\)](#) and [Aguirregabiria and Mira \(2007\)](#).

analysis of general treatment effects:

$$\begin{aligned}
U_i(y, X, D, G) = & \overbrace{y_i(\theta_0 + \theta_1 D_i + X_i^\top \theta_2 + X_i^\top \theta_3 D_i)}^{\alpha_i} + \frac{1}{|\mathcal{N}_i|} \sum_{j \neq i} \theta_4 m_{ij} G_{ij} D_j - \varepsilon_i \\
& + \sum_{j \neq i} \frac{1}{|\mathcal{N}_i|} \underbrace{(\theta_5 + \theta_6 D_i D_j) m_{ij} G_{ij}}_{\beta_{ij}} y_i y_j,
\end{aligned} \tag{16}$$

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. \(2010a\)](#).

4.1.2 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 $\{\hat{\sigma}_i^{data}\}_{i=1}^N$ 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 2. (Sub-Gaussian CCP Estimator) *There exists a positive constant C_σ such that for every $t \geq 0$, we have*

$$\Pr(|\hat{\sigma}_i^{data} - \sigma_i^{data}| \geq t | S, \sigma^{data}) \leq 2 \exp(-nt^2/C_\sigma^2), \quad \forall i = 1, \dots, n.$$

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 $\hat{Q}_n(\hat{\sigma}^{data}, \theta)$.

$$\hat{Q}_n(\hat{\sigma}^{data}, \theta) = \frac{1}{n} \sum_{i=1}^n Y_i \log(F_\varepsilon(\hat{Z}_i^\top \theta)) + (1 - Y_i) \log(1 - F_\varepsilon(\hat{Z}_i^\top \theta)),$$

where

$$\hat{Z}_i = \left(1, D_i, X_i^\top, X_i^\top D_i, \frac{1}{|\mathcal{N}_i|} \sum_{j \neq i} m_{ij} G_{ij} D_j, \frac{1}{|\mathcal{N}_i|} \sum_{j \neq i} m_{ij} G_{ij} \hat{\sigma}_j^{data}, \frac{1}{|\mathcal{N}_i|} \sum_{j \neq i} m_{ij} G_{ij} \hat{\sigma}_j D_i D_j\right)^\top. \quad (17)$$

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} .

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.

¹⁰By Hoeffding's inequality, the frequency estimator easily satisfies the Assumption 2. The probit/logit estimator satisfies the Assumption 2 by the same argument as the second-stage MLE estimator in Section 5.

4.1.3 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:

$$\begin{aligned}\Pr(Y_i = 1|X, D, G, \bar{\lambda}; \hat{\theta}) &= \int \mathbb{1}\left\{\hat{\alpha}_i + \sum_{j \neq i} \hat{\beta}_{ij} \Pr(Y_j = 1|X, D, G, \bar{\lambda}; \hat{\theta}) \geq \varepsilon_i\right\} dF_\varepsilon, \\ \Pr(Y_i = 1|X, D, G, \underline{\lambda}; \hat{\theta}) &= \int \mathbb{1}\left\{\hat{\alpha}_i + \sum_{j \neq i} \hat{\beta}_{ij} \Pr(Y_j = 1|X, D, G, \underline{\lambda}; \hat{\theta}) \geq \varepsilon_i\right\} dF_\varepsilon.\end{aligned}\quad (18)$$

From Theorem 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^* , $\bar{\sigma}^M = \Omega^M(\bar{\sigma}^0) \geq \Omega^M(\sigma^*) = \sigma^*$.

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^* , $\underline{\sigma}^M = \Omega^M(\underline{\sigma}^0) \leq \Omega^M(\sigma^*) = \sigma^*$.

4.1.4 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:

$$\tilde{D} = \arg \max_{D \in \mathcal{D}_\kappa} W_{X,G,\lambda}^n(D) = \arg \max_{D \in \mathcal{D}_\kappa} W_{X,G,\underline{\sigma}^*}^n(D). \quad (19)$$

As shown in Eq.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.19 is a complicated combinatorial optimization problem. We propose a greedy algorithm¹¹ (Algorithm 1) 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 1 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 1: Maximizing Over Treatment Allocation Rules

Input: Weighted 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

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:

$$\mathbb{E}_{\varepsilon^n} [R(\hat{D}_G) | S, \sigma^{data}] := \max_{D \in \mathcal{D}} W(D) - \mathbb{E}_{\varepsilon^n} [W(\hat{D}_G) | S, \sigma^{data}],$$

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

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.

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 5.3. The final term also measures regret introduced by using the empirical social welfare function. This term is bounded by:

$$W_n(\hat{D}_G) - W(\hat{D}_G) \leq \sup_{D \in \mathcal{D}} |W_n(D) - W(D)|.$$

Combining all the above results, we conclude that expected regret is bounded by:

$$\mathbb{E}_{\varepsilon^n} \left[R(\hat{D}_G) | S, \sigma^{data} \right] \leq 2\mathbb{E}_{\varepsilon^n} \left[\max_{D \in \mathcal{D}} |W_n(D) - W(D)| | S, \sigma^{data} \right] + \mathbb{E}_{\varepsilon^n} \left[W_n(\tilde{D}) - W_n(\hat{D}_G) | S, \sigma^{data} \right]. \quad (20)$$

In the remainder of this section, we provide a non-asymptotic upper bound for expected regret.

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 3. (Compactness) *The parameter θ lies in a compact set $\Theta \subseteq \mathbb{R}^{d_\theta}$.*

Assumption 3 is standard. We now proceed to characterize the sampling uncertainty associated with using the empirical welfare function.

Lemma 5.1. *Under Assumptions 1 and 3,*

$$\mathbb{E}_{\varepsilon^n} \left[\max_{D \in \mathcal{D}} |W_n(D) - W(D)| \middle| S, \sigma^{data} \right] \leq C_1 \mathbb{E}_{\varepsilon^n} \left[\|\hat{\theta} - \theta_0\|_1 \middle| S, \sigma^{data} \right],$$

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 5.1 is provided in Appendix B.2. Lemma 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

$$\mathbb{E}_{\varepsilon^n} [|\mathbb{G}_n(\hat{\theta}) - \mathbb{G}_n(\theta_0)| | S, \sigma^{data}], \quad (21)$$

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.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.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:

$$\mathbb{G}_n(\hat{\theta}) - \mathbb{G}_n(\theta_0) = \hat{\mathbb{M}}(\hat{\theta}) - \hat{\mathbb{M}}(\theta_0) + M(\theta_0) - M(\hat{\theta}). \quad (22)$$

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.22:

$$\hat{\mathbb{M}}(\theta_0) - \hat{\mathbb{M}}(\hat{\theta}) = \frac{1}{2}(\hat{\theta} - \theta_0)^\top \nabla_{\hat{\theta}}^2 \hat{\mathbb{M}}(\hat{\theta})(\hat{\theta} - \theta_0), \quad (23)$$

$$M(\hat{\theta}) - M(\theta_0) = \frac{1}{2}(\hat{\theta} - \theta_0)^\top \nabla_{\hat{\theta}}^2 M(\hat{\theta})(\hat{\theta} - \theta_0), \quad (24)$$

for some $\hat{\theta} \in \mathbb{R}^{d_\theta}$ and $\tilde{\theta} \in \mathbb{R}^{d_\theta}$ 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 1 (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:

$$\hat{\mathbb{M}}(\theta_0) - \hat{\mathbb{M}}(\hat{\theta}) \leq \frac{1}{2} \eta_{max}^0 \|\hat{\theta} - \theta_0\|_2^2. \quad (25)$$

$$M(\hat{\theta}) - M(\theta_0) \leq \frac{1}{2} \eta_{max}^1 \|\hat{\theta} - \theta_0\|_2^2. \quad (26)$$

Combining Eq.22 with Eq.25 and Eq.26 yields

$$-(\mathbb{G}_n(\hat{\theta}) - \mathbb{G}_n(\theta_0)) \leq \frac{1}{2}(\eta_{max}^0 + \eta_{max}^1) \|\hat{\theta} - \theta_0\|_2^2. \quad (27)$$

Applying the mean value Theorem to the left-hand side of Eq.27, we have

$$\mathbb{G}_n(\hat{\theta}) - \mathbb{G}_n(\theta_0) = (\hat{\theta} - \theta_0)^\top \nabla_{\tilde{\theta}} \mathbb{G}_n(\tilde{\theta}), \quad (28)$$

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.28 must be positive given the definition of \mathbb{G}_n . By the Cauchy–Schwarz inequality,

$$\mathbb{G}_n(\hat{\theta}) - \mathbb{G}_n(\theta_0) = (\hat{\theta} - \theta_0)^\top \nabla_{\tilde{\theta}} \mathbb{G}_n(\tilde{\theta}) = |(\hat{\theta} - \theta_0)^\top \nabla_{\tilde{\theta}} \mathbb{G}_n(\tilde{\theta})| \leq \|\hat{\theta} - \theta_0\|_2 \|\nabla_{\tilde{\theta}} \mathbb{G}_n(\tilde{\theta})\|_2. \quad (29)$$

Combining Eq.27 and Eq.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., $\eta_{max} = \max_{A \in \mathbb{R}^C \setminus \{0\}} \frac{A^\top M A}{A^\top A}$).

$\nabla_{\theta}^2 \hat{M}(\hat{\theta})$ 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,

$$\begin{aligned} \|\hat{\theta} - \theta_0\|_1 &\leq d_{\theta} \|\hat{\theta} - \theta_0\|_2 \leq -\frac{2d_{\theta}}{\eta_{max}^0 + \eta_{max}^1} \|\nabla_{\theta} \mathbb{G}_n(\tilde{\theta})\|_2 \\ &\leq -\frac{2d_{\theta}}{\eta_{max}^0 + \eta_{max}^1} \|\nabla_{\theta} \mathbb{G}_n(\tilde{\theta})\|_1. \end{aligned} \quad (30)$$

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 4. (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 5. (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 6. (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 4 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 6 guarantees that the average number of treated units in the training data is non-zero for any network size. Building on Eq.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 5.2. (Sampling Uncertainty) *Under Assumption 1, 4, 5, and 6, 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 5.2 is provided in Appendix B.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 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.

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 5.3. *Under Assumption 1 to 5, we have*

$$\mathbb{E}_{\varepsilon^n} \left[\sup_{\theta \in \Theta} \|\nabla_{\theta} \mathbb{G}_n(\theta)\|_1 | S, \sigma^{data} \right] \leq \frac{C_3 + C_4 \sqrt{\log(n)}}{\sqrt{n}},$$

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 5.3 is provided in Appendix B.4. Lemma 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 5.1, Lemma 5.2 and Lemma 5.3, we have our main theorem. This theorem characterizes the sampling uncertainty of using the empirical welfare:

Theorem 5.1. (Sampling Uncertainty of Regret) *Under Assumption 1 to 6, 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 5.1 is provided in Appendix C.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.

5.3 Regret due to our Greedy Algorithm

Now, we evaluate the second term of Eq.20, which is the regret introduced by our greedy algorithm,

$$\mathbb{E}_{\varepsilon^n} [W_n(\tilde{D}) - W_n(\hat{D}_G) | S, \sigma^{data}].$$

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 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 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 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 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 $(\hat{\theta}_1 + X_i^T \hat{\theta}_3)$ 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. \(2017, Theorem 1\)](#) leads to:

Proposition 5.1. *Under Assumption 1 and Assumption 3, 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.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 1.

A proof is provided in Appendix C.3. This proof is similar to [Kitagawa and Wang \(2023a\)](#). The first part of Proposition 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,

$$\mathbb{E}_{\varepsilon^n} \left[W_n(\tilde{D}) - W_n(\hat{D}_G) \mid S, \sigma^{data} \right] \leq \mathcal{O}(1) \left(1 - \frac{1}{\xi} (1 - e^{-\xi\gamma}) \right), \quad (31)$$

where $\mathcal{O}(1)$ captures the $\mathbb{E}_{\varepsilon^n} [W_n(\tilde{D}) \mid S, \sigma^{data}]$. Combining Eq.31 with Theorem 5.1, we obtain our main theorem:

Theorem 5.2. (Regret Bound) *Let D^* denote the maximizer of $W(D)$ and D_G be the assignment vector obtained by Algorithm 1. Under Assumptions 1 to 5, given curvature ξ and submodularity ratio γ , the regret is bounded from above by:*

$$\mathbb{E}_{\varepsilon^n} \left[R(\hat{D}_G) \mid S, \sigma^{data} \right] \leq C_1 C_2 \frac{C_3 + C_4 \log(n)}{\sqrt{n}} + \mathcal{O}(1) \left(1 - \frac{1}{\xi} (1 - e^{-\xi\gamma}) \right). \quad (32)$$

Theorem 5.2 is our key result. The first term in Eq.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 5.1, C_5 and C_6 in Lemma D.4, and C_4 in Lemma D.3. The second term comes from the use of a greedy algorithm, and converges to a constant.

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). 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 $m(X_i, X_j) = \frac{1}{1+|X_i-X_j|}$.

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.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 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 7¹⁵.

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.

¹⁵The goodness of fit test indicates that the estimation for Village 7, referred to as Village 17 in Table 2, may not adequately fit the data, potentially due to inaccuracies in the first-stage Conditional Choice Probability (CCP) estimation.

Village	Sample Avg.	Welfare under			Welfare Gain*	
		Original	Random	Robust	Level	Percentage
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 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.

Second, we find that the equilibrium average share of households purchasing microfinance 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 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 1: Comparisons between three approaches

7 Extension

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 3.1 under an incomplete information setting, we propose the following:

Proposition 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:

$$\begin{aligned} \inf_{\lambda \in \Lambda} \sum_{i=1}^N \Pr(Y_i = 1 | X, D, G, \lambda) &= \sum_{i=1}^N \inf_{\lambda \in \Lambda} \Pr(Y_i = 1 | X, D, G, \lambda), \\ \sup_{\lambda \in \Lambda} \sum_{i=1}^N \Pr(Y_i = 1 | X, D, G, \lambda) &= \sum_{i=1}^N \sup_{\lambda \in \Lambda} \Pr(Y_i = 1 | X, D, G, \lambda). \end{aligned}$$

The proof of Proposition 7.1 mirrors that of Theorem 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:

$$\Pr(Y_i = 1|X, D, G, \bar{\lambda}) = \int 1 \left\{ \max_{y_{-i}: (1, y_{-i}) \in \Sigma(\varepsilon)} U_i(1, y_{-i}, X, D, G) \geq 0 \right\} dF_\varepsilon, \quad (33)$$

$$\Pr(Y_i = 0|X, D, G, \bar{\lambda}) = \int 1 \left\{ \min_{y_{-i}: (1, y_{-i}) \in \Sigma(\varepsilon)} U_i(1, y_{-i}, X, D, G) < 0 \right\} dF_\varepsilon. \quad (34)$$

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.33 (maximization) and Eq.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.33 to:

$$\max_{y_{\mathcal{N}_i}, y_{i^C}} U_i(1, y_{\mathcal{N}_i}, X, D, G)$$

with constraints:

$$y_j = 1\{U_j(1, y_{\mathcal{N}_j \setminus \{i\}}, X, D, G) \geq 0\}, \forall j \in \mathcal{N}_i, \quad (35)$$

$$y_k = 1\{U_k(1, y_{\mathcal{N}_k}, X, D, G) \geq 0\}, \forall k \in i^C. \quad (36)$$

These constraints (Eq.35 and Eq.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.

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.

References

- ADJAHO, 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.
- 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. (2020): “Optimal treatment assignment rules on networked populations,” Tech. rep., working paper.
- ARCIDIACONO, P. AND R. A. MILLER (2011): “Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity,” *Econometrica*, 79, 1823–1867.

- ATHEY, S. AND S. WAGER (2021): “Policy learning with observational data,” *Econometrica*, 89, 133–161.
- 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.
- 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, 1236–1249.
- BERGE, C. (1963): *Topological Spaces*, translated by E. M. Patterson (Oliver & Boyd).
- BERGER, J. O. (1994): *The Robust Bayesian Viewpoint*, Robustness of Bayesian Analysis.
- BIAN, A. A., J. M. BUHMANN, A. KRAUSE, AND S. TSCHIATSCHKEK (2017): “Guarantees for greedy maximization of non-submodular functions with applications,” in *International conference on machine learning*, PMLR, 498–507.
- 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.
- 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.
- 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*.
- 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*.
- (2020): “Robust decision theory and econometrics,” *Annual Review of Economics*, 12, 239–271.
- 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.
- 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, A. (2013): “Econometric analysis of games with multiple equilibria,” *Annu. Rev. Econ.*, 5, 107–131.
- 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.
- ECHENIQUE, F. AND I. KOMUNJER (2009): “Testing models with multiple equilibria by quantile methods,” *Econometrica*, 77, 1281–1297.
- FERNANDEZ, A. A., J. L. M. OLEA, C. QIU, J. STOYE, AND S. TINDA (2024): “Robust Bayes Treatment Choice with Partial Identification,” *working paper*.
- 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.
- 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.

- 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*.
- GRAHAM, B. S. AND A. PELICAN (2023): “Scenario Sampling for Large Supermodular Games,” Tech. rep., National Bureau of Economic Research.
- 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.
- 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.
- 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.
- JIA, P. (2008): “What happens when Wal-Mart comes to town: An empirical analysis of the discount retailing industry,” *Econometrica*, 76, 1263–1316.
- JOVANOVIC, 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.
- KEANE, M. AND N. WASI (2013): “Comparing alternative models of heterogeneity in consumer choice behavior,” *Journal of Applied Econometrics*, 28, 1018–1045.
- 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 (2022): “Policy choice in time series by empirical welfare maximization,” *arXiv preprint arXiv:2205.03970*.
- KLINE, B. AND E. TAMER (2020): “Econometric analysis of models with social interactions,” in *The Econometric Analysis of Network Data*, Elsevier, 149–181.
- KOR, R. AND J. ZHOU (2022): “Welfare and distributional effects of joint intervention in networks,” *arXiv preprint arXiv:2206.03863*.
- 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, 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*.
- LOVÁSZ, L. (1983): “Submodular functions and convexity,” *Mathematical Programming The State of the Art: Bonn 1982*, 235–257.
- 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.
- MATHEVET, L. (2010): “Supermodular mechanism design,” *Theoretical Economics*, 5, 403–443.
- MBAKOP, 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.
- MENZEL, K. (2016): “Inference for games with many players,” *The Review of Economic Studies*, 83, 306–337.
- 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.
- 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*.
- NAKAJIMA, R. (2007): “Measuring peer effects on youth smoking behaviour,” *The Review of Economic Studies*, 74, 897–935.
- 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.
- RIDDER, G. AND S. SHENG (2020): “Two-step estimation of a strategic network formation model with clustering,” *arXiv preprint arXiv:2001.03838*.
- 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.
- 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.
- 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.
- 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.
- VIVIANO, D. (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.
- 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.

Supplementary Materials of Robust Network Targeting with Multiple Nash Equilibria

Appendix A Chi-Square Goodness of Fit Test

Village	Number of Rooms				Village	Number of Rooms			
	1-2	3-4	5-6	≥ 7		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 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.

Appendix B Lemmas

We first introduce notation. We define $\bar{m} := \max_{ij} |m_{ij}|$, and $\underline{F}_\varepsilon := \min_{\theta \in \Theta} F_\varepsilon(z^\top \theta)$. 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: $\|A\|_\infty = \max_{k=1, \dots, K} \sum_{l=1}^L |A_{kl}|$.

B.1 Proof of Corollary 3.1

Corollary B.1. (Utilitarian Welfare at Equilibrium) Under Assumption 1, 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,

$$W_{X,G,\lambda}(D) = \frac{1}{N} \sum_{i=1}^N \alpha_i \Pr(Y_i = 1 | X, D, G, \lambda) + \frac{1}{N} \sum_{i=1}^N \sum_{j \neq i} \beta_{ij} \Pr(Y_i Y_j = 1 | X, D, G, \lambda). \quad (37)$$

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 1. 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 3.1, $\{\Pr(Y_j = 1 | X, D, G, \sigma^*)\}_{i=1}^N$ 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 $i, j \in \mathcal{N}$,

$$\beta_{ij} \Pr(Y_i Y_j = 1 | X, D, G, \lambda) \in [\beta_{ij} \underline{\sigma}_i^* \underline{\sigma}_j^*, \beta_{ij} \bar{\sigma}_i^* \bar{\sigma}_j^*]. \quad (38)$$

Plugging Eq.38 into Eq.37 completes the proof. \square

B.2 Proof of Lemma 5.1

Lemma B.1. *Under Assumptions 1 and 3,*

$$\mathbb{E}_{\varepsilon^n} \left[\max_{D \in \mathcal{D}} |W_n(D) - W(D)| \middle| S, \sigma^{data} \right] \leq C_1 \mathbb{E}_{\varepsilon^n} \left[\|\hat{\theta} - \theta\|_1 \middle| S, \sigma^{data} \right]$$

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:

$$\begin{aligned}\underline{\sigma}_i(\theta) &= \int \mathbb{1} \left\{ \alpha_i + \sum_{j \neq i} \beta_{ij} \underline{\sigma}_j(\theta) - \varepsilon_i \geq 0 \right\} dF(\varepsilon) \\ &= F_\varepsilon \left(\theta_0 + \theta_1 D_i + \theta_2^\top X_i + \theta_3^\top X_i D_i + \frac{\theta_4}{N_i} \sum_{j \neq i} G_{ij} m_{ij} D_j + \frac{1}{N_i} \sum_{j \neq i} (\theta_5 + \theta_6 D_i D_j) G_{ij} m_{ij} \underline{\sigma}_j(\theta) \right).\end{aligned}$$

Let $r(i, \theta) := \theta_0 + \theta_1 D_i + \theta_2^\top X_i + \theta_3^\top X_i D_i + \frac{\theta_4}{N_i} \sum_{j \neq i} G_{ij} m_{ij} D_j + \frac{1}{N_i} \sum_{j \neq i} (\theta_5 + \theta_6 D_i D_j) G_{ij} m_{ij} \underline{\sigma}_j(\theta)$. 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:

$$|\underline{\sigma}_i(\hat{\theta}) - \underline{\sigma}_i(\theta)| \leq \tau \|\nabla_{\theta} r(i, \tilde{\theta})\|_2 \|\hat{\theta} - \theta\|_2 \leq \tau \|\nabla_{\theta} r(i, \tilde{\theta})\|_1 \|\hat{\theta} - \theta\|_1. \quad (39)$$

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,

$$\|\nabla_{\theta} r(i, \theta)\|_1 \leq \|\nabla_{\theta} r(\theta)\|_{\infty}, \quad \forall i \in \mathcal{N}, \quad (40)$$

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 $a_i = \theta_0 + \theta_1 D_i - \theta_2^T X_i - \theta_3^T X_i D_i$. By the Implicit Function Theorem, $\nabla_{\theta} r(\theta)$ is given by:

$$\nabla_{\theta} r(\theta) = -(\nabla_r I(r, \theta))^{-1} \nabla_{\theta} I(r, \theta), \quad (41)$$

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.41 is bounded by

$$\|\nabla_\theta r(\theta)\|_\infty = \|(\nabla_r I(r, \theta))^{-1} \nabla_\theta I(r, \theta)\|_\infty \leq \|(\nabla_r I(r, \theta))^{-1}\|_\infty \|\nabla_\theta I(r, \theta)\|_\infty. \quad (42)$$

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 D.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 3), and $F'_\varepsilon(\cdot) \in [0, \tau]$, the Extreme Value Theorem (Lemma E.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.42, we know:

$$\begin{aligned} \|\nabla_\theta I(r, \tilde{\theta})\|_\infty &= \max_{i \in \mathcal{N}} 1 + D_i + \|X_i\|_1 + D_i \|X_i\|_1 + \frac{1}{|\mathcal{N}_i|} \sum_{j \neq i} D_j G_{ij} |m_{ij}| \\ &\quad + \frac{1}{|\mathcal{N}_i|} \sum_{j \neq i} F'_\varepsilon(r(j, \tilde{\theta})) (1 + D_i D_j) G_{ij} |m_{ij}|. \end{aligned} \quad (43)$$

Furthermore, Eq.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

$$\|\nabla_{\theta} I(r, \tilde{\theta})\|_{\infty} \leq 2 + 2\|X\|_{\infty} + \bar{m} + 2\bar{m}\tau. \quad (44)$$

Combining Eq.39, Eq.40 and Eq.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 $C_1 = \zeta\tau(2 + 2\|X\|_{\infty} + \bar{m} + 2\bar{m}\tau)$. □

B.3 Proof of Lemma 5.2

Lemma B.2. (Sampling Uncertainty) Under Assumption 1, 4, 5, and 6, 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.23 and Eq.24, we have:

$$\hat{\mathbb{M}}(\theta_0) - \hat{\mathbb{M}}(\hat{\theta}) = \frac{1}{2n} \sum_{i=1}^n [Y_i \omega_0(\hat{Z}_i^{\top} \hat{\theta}) - (1 - Y_i) \omega_1(\hat{Z}_i^{\top} \hat{\theta})] [(\hat{\theta} - \theta_0)^{\top} \hat{Z}_i]^2 \leq 0. \quad (45)$$

$$M(\hat{\theta}) - M(\theta_0) = \frac{1}{2n} \sum_{i=1}^n [\sigma_i \omega_0(Z_i^{\top} \hat{\theta}) - (1 - \sigma_i) \omega_1(Z_i^{\top} \hat{\theta})] [(\hat{\theta} - \theta_0)^{\top} Z_i]^2 \leq 0, \quad (46)$$

where $\sigma_i = \mathbb{E}_{\varepsilon^n} [Y_i | S, \sigma^{data}]$. Although Eq.45 and Eq.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 4, we have $\omega_0(a) < 0$ and $\omega_1(a) > 0$ for all $a \in \mathbb{R}$. Furthermore, Assumption 3 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 1 and 4, the probability density function of ε^n and its derivative are bounded. In addition, under Assumption 1 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.45 and Eq.46:

$$\mathbb{G}_n(\hat{\theta}) - \mathbb{G}_n(\theta_0) \geq -\bar{\omega}(\hat{\theta} - \theta_0)^\top \frac{1}{2n} \sum_{i=1}^n (Z_i Z_i^\top + \hat{Z}_i \hat{Z}_i^\top)(\hat{\theta} - \theta_0). \quad (47)$$

Under Assumption 5, $\frac{1}{n} \sum_{i=1}^n \hat{Z}_i \hat{Z}_i^\top$ 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 D.2 addresses this by guaranteeing the existence of a uniform lower bound on the smallest eigenvalue. Lemma D.2 is a corollary of *Berge's Maximum Theorem* (Lemma E.2). A proof is provided in Appendix D.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 E.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 6 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,

$$(\hat{\theta} - \theta_0)^\top \frac{1}{2n} \sum_{i=1}^n (Z_i Z_i^\top + \hat{Z}_i \hat{Z}_i^\top)(\hat{\theta} - \theta_0) \geq \underline{\varsigma}_{min} \|\hat{\theta} - \theta_0\|_2^2 > 0. \quad (48)$$

Combining Eq.30 with Eq.47 and Eq.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 $C_2 = \frac{d_\theta}{-\underline{\omega}_{\min}}$. □

B.4 Proof of Lemma 5.3

Lemma B.3. *Under Assumption 1 to 5, we have*

$$\mathbb{E}_{\varepsilon^n} \left[\sup_{\theta \in \Theta} \|\nabla_\theta \mathbb{G}_n(\theta)\|_1 \mid \mathcal{S}, \sigma^{data} \right] \leq \frac{C_3 + C_4 \sqrt{\log(n)}}{\sqrt{n}},$$

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:

$$\|\nabla_\theta \mathbb{G}_n(\tilde{\theta})\|_1 \leq \|\nabla_\theta \hat{\mathbb{M}}(\tilde{\theta}) - \nabla_\theta \mathbb{M}(\tilde{\theta})\|_1 + \|\nabla_\theta \mathbb{M}(\tilde{\theta}) - \nabla_\theta M(\tilde{\theta})\|_1. \quad (49)$$

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.49 is then

$$\|\nabla_\theta \hat{\mathbb{M}}(\tilde{\theta}) - \nabla_\theta \mathbb{M}(\tilde{\theta})\|_1 = \sum_{k=1}^{d_\theta} |\mathbb{A}_k(\tilde{\theta})|, \quad (50)$$

and the second term in Eq.49 is

$$\|\nabla_{\theta}\mathbb{M}(\tilde{\theta}) - \nabla_{\theta}M(\tilde{\theta})\|_1 = \sum_{k=1}^{d_{\theta}} |\mathbb{B}_k(\tilde{\theta})|. \quad (51)$$

Combining Eq.49 with Eq.50 and Eq.51, we conclude

$$\mathbb{E}_{\varepsilon^n} [\|\nabla_{\theta}\mathbb{G}_n(\tilde{\theta})\|_1 | S, \sigma^{data}] \leq \sum_{k=1}^{d_{\theta}} \mathbb{E}_{\varepsilon^n} [|\mathbb{A}_k(\tilde{\theta})| | S, \sigma^{data}] + \sum_{k=1}^{d_{\theta}} \mathbb{E}_{\varepsilon^n} [|\mathbb{B}_k(\tilde{\theta})| | S, \sigma^{data}]. \quad (52)$$

An upper bound for the first term in Eq.52 is provided by Lemma D.3. An upper bound for the second term in Eq.52 is provided by Lemma D.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

Appendix C Theorems and Propositions

C.1 Proof of Theorem 3.1

Theorem C.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.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:

$$\Pr(Y_i = 1|X, D, G, \underline{\lambda}) \leq \Pr(Y_i = 1|X, D, G, \lambda), \quad \forall \lambda \in \Lambda, \quad (53)$$

$$\Pr(Y_i = 1|X, D, G, \bar{\lambda}) \geq \Pr(Y_i = 1|X, D, G, \lambda), \quad \forall \lambda \in \Lambda, \quad (54)$$

Therefore, we must have:

$$\sum_{i=1}^N \Pr(Y_i = 1|X, D, G, \underline{\lambda}) \leq \sum_{i=1}^N \Pr(Y_i = 1|X, D, G, \lambda), \quad \forall \lambda \in \Lambda,$$

$$\sum_{i=1}^N \Pr(Y_i = 1|X, D, G, \bar{\lambda}) \geq \sum_{i=1}^N \Pr(Y_i = 1|X, D, G, \lambda), \quad \forall \lambda \in \Lambda.$$

As a consequence,

$$\underline{\lambda} = \arg \min_{\lambda \in \Lambda} \sum_{i=1}^N \Pr(Y_i = 1|X, D, G, \lambda).$$

$$\bar{\lambda} = \arg \max_{\lambda \in \Lambda} \sum_{i=1}^N \Pr(Y_i = 1|X, D, G, \lambda).$$

In addition, given Eq.53 and Eq.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,

$$\inf_{\lambda \in \Lambda} \sum_{i=1}^N \Pr(Y_i = 1|X, D, G, \lambda) = \sum_{i=1}^N \inf_{\lambda \in \Lambda} \Pr(Y_i = 1|X, D, G, \lambda),$$

$$\sup_{\lambda \in \Lambda} \sum_{i=1}^N \Pr(Y_i = 1 | X, D, G, \lambda) = \sum_{i=1}^N \sup_{\lambda \in \Lambda} \Pr(Y_i = 1 | X, D, G, \lambda).$$

□

C.2 Proof of Theorem 5.1

Theorem C.2. (Sampling Uncertainty of Regret) Under Assumption 1 to 6, 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

$$\mathbb{E}_{\varepsilon^n} [\|\hat{\theta} - \theta_0\|_1 | S, \sigma^{data}] \leq C_2 \mathbb{E}_{\varepsilon^n} [\|\nabla_{\theta} \mathbb{G}_n(\hat{\theta})\|_1 | S, \sigma^{data}]. \quad (55)$$

Plugging Lemma 5.3 into Eq.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 5.1 with Eq.55, we conclude:

$$\mathbb{E}_{\varepsilon^n} \left[\max_{D \in \mathcal{D}} |W_n(D) - W(D)| \middle| S, \sigma^{data} \right] \leq C_1 C_2 \frac{C_3 + C_4 \log(n)}{\sqrt{n}}.$$

□

C.3 Proof of Proposition 5.1

Proposition C.1. Under Assumptions 1 and Assumptions 3, 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.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 1.

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.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 1, $W_n(\mathcal{R} \cup \{k\}) - W_n(\mathcal{R})$ 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:

$$W_n(\mathcal{R} \cup \{k\}) - W_n(\mathcal{R}) \\ \leq \frac{\tau}{N} \sum_{i \in \mathcal{R}} \left(\frac{1}{|\mathcal{N}_i|} \sum_{j \neq i} \hat{\theta}_5 \bar{m} G_{ij} + \frac{1}{|\mathcal{N}_i|} \sum_{j \neq i} \hat{\theta}_6 \bar{m} G_{ij} + \frac{\hat{\theta}_4 \bar{m} G_{ik}}{N} \right) \\ + \frac{\tau}{N} \sum_{m \in \mathcal{N} \setminus \mathcal{R} \cup \{k\}} \left(\frac{\hat{\theta}_4 \bar{m} G_{mk}}{N} + \frac{1}{|\mathcal{N}_m|} \sum_{j \neq m} \hat{\theta}_5 \bar{m} G_{mj} \right) \\ + \frac{\tau}{N} \left(\hat{\theta}_1 + X_i^\top \hat{\theta}_3 + \frac{1}{|\mathcal{N}_k|} \sum_{j \neq k} \hat{\theta}_5 \bar{m} G_{kj} + \frac{1}{|\mathcal{N}_k|} \sum_{j \neq k} \hat{\theta}_6 \bar{m} G_{kj} \right).$$

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:

$$W_n(\mathcal{R} \cup \{k\}) - W_n(\mathcal{R}) \\ \leq \frac{\tau}{N} \sum_{i \in \mathcal{R}} (\hat{\theta}_5 + \hat{\theta}_6) \bar{m} + \frac{\tau}{N} \sum_{m \in \mathcal{N} \setminus \mathcal{R} \cup \{k\}} \hat{\theta}_5 \bar{m} + \frac{\tau}{N} \left(\hat{\theta}_1 + X_i^\top \hat{\theta}_3 + (\hat{\theta}_5 + \hat{\theta}_6) \bar{m} \right) + \frac{\tau}{N} \frac{\hat{\theta}_4 \bar{m} N}{N}.$$

Summarizing all the units together, we have:

$$W_n(\mathcal{R} \cup \{k\}) - W_n(\mathcal{R}) \leq \tau (\hat{\theta}_5 + \hat{\theta}_6) \bar{m} + \frac{\tau}{N} \left(\hat{\theta}_1 + \bar{X}^\top \hat{\theta}_3 + \frac{\hat{\theta}_4 \bar{m} N}{N} \right),$$

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:

$$\begin{aligned}
& W_n(\mathcal{R} \cup \{k\}) - W_n(\mathcal{R}) \\
& \geq \frac{F_\varepsilon}{N} \sum_{i \in \mathcal{R}} \left(\frac{1}{|\mathcal{N}_i|} \sum_{j \neq i} \hat{\theta}_5 m_{ij} G_{ij}(\underline{\sigma}_j - \underline{\sigma}'_j) + \frac{1}{|\mathcal{N}_i|} \sum_{j \neq i, k} \hat{\theta}_6 D_j m_{ij} G_{ij}(\underline{\sigma}_j - \underline{\sigma}'_j) + \frac{1}{|\mathcal{N}_i|} (\hat{\theta}_4 + \hat{\theta}_6 \underline{\sigma}_k) m_{ik} G_{ik} \right) \\
& + \frac{F_\varepsilon}{N} \sum_{m \in \mathcal{N} \setminus \mathcal{R} \cup \{k\}} \left(\frac{1}{|\mathcal{N}_m|} \hat{\theta}_4 m_{mk} G_{mk} + \frac{1}{|\mathcal{N}_m|} \sum_{j \neq m} \hat{\theta}_5 m_{mj} G_{mj}(\underline{\sigma}_j - \underline{\sigma}'_j) \right) \\
& + \frac{F_\varepsilon}{N} \left(\hat{\theta}_1 + X_i^\top \hat{\theta}_3 + \frac{1}{|\mathcal{N}_k|} \sum_{j \neq k} \hat{\theta}_5 m_{kj} G_{kj}(\underline{\sigma}_j - \underline{\sigma}'_j) + \frac{1}{|\mathcal{N}_k|} \sum_{j \neq k} \hat{\theta}_6 D_j m_{kj} G_{kj} \underline{\sigma}_j \right).
\end{aligned} \tag{56}$$

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.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:

$$\underline{\sigma}_k - \underline{\sigma}'_k \geq F_\varepsilon (\hat{\theta}_1 + X_k^\top \hat{\theta}_3). \tag{57}$$

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.56:

$$\underline{\sigma}_i - \underline{\sigma}'_i \geq F_\varepsilon \left(\frac{1}{|\mathcal{N}_i|} \sum_{j \neq i} \hat{\theta}_5 m_{ij} G_{ij}(\underline{\sigma}_j - \underline{\sigma}'_j) + \frac{1}{|\mathcal{N}_i|} \sum_{j \neq i, k} \hat{\theta}_6 D_j m_{ij} G_{ij}(\underline{\sigma}_j - \underline{\sigma}'_j) + \frac{1}{|\mathcal{N}_i|} (\hat{\theta}_4 + \hat{\theta}_6 \underline{\sigma}_k) m_{ik} G_{ik} \right). \tag{58}$$

We then further bound the indirect effects in Eq.58 by:

$$\underline{\sigma}_i - \underline{\sigma}'_i \geq F_\varepsilon \frac{1}{|\mathcal{N}_i|} \hat{\theta}_4 m_{ik} G_{ik} \geq \frac{F_\varepsilon}{N} \hat{\theta}_4 m_{ik} G_{ik}. \tag{59}$$

For units m not in the treated group, the indirect treatment effects, which is given by the second term of Eq.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:

$$\underline{\sigma}_m - \underline{\sigma}'_m \geq \frac{F_\varepsilon}{N} \hat{\theta}_4 m_{mk} G_{mk}. \tag{60}$$

Combining Eq.57, Eq.59 and Eq.60 with Eq.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:

$$W_n(\mathcal{R} \cup \{k\}) - W_n(\mathcal{R}) \geq \frac{F_\varepsilon m \hat{\theta}_4 \underline{N}}{N\bar{N}} + \frac{F_\varepsilon}{N} (\hat{\theta}_1 + X_i^\top \hat{\theta}_3).$$

Therefore,

$$\frac{W_n(R \cup \{k\}) - W_n(R)}{W_n(S \cup \{k\}) - W_n(S)} \geq \frac{\frac{F_\varepsilon m \hat{\theta}_4 \underline{N}}{N\bar{N}} + \frac{F_\varepsilon}{N} (\hat{\theta}_1 + X_i^\top \hat{\theta}_3)}{\tau(\hat{\theta}_5 + \hat{\theta}_6) \bar{m} + \frac{\tau}{N} (\hat{\theta}_1 + \bar{X}^\top \hat{\theta}_3 + \frac{\hat{\theta}_4 \bar{m} \bar{N}}{N})},$$

which ranges between (0,1). As a consequence, the submodularity ratio

$$\xi = \max_{S \subseteq RCN, k \in \mathcal{N} \setminus R} 1 - \frac{W_n(R \cup \{k\}) - W_n(R)}{W_n(S \cup \{k\}) - W_n(S)} \in (0, 1).$$

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

Appendix D Preliminary Lemmas

D.1 Lemma D.1

Lemma D.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, $f_i(A) = \|A_{i,:}^{-1}\|$, 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

D.2 Lemma D.2

Lemma D.2. Under Assumption 1, 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

$$\eta_{min}(B) = \min_{\theta \neq 0 \in \mathbb{R}^{d_{\theta}}} \frac{\theta^{\top} B \theta}{\theta^{\top} \theta}, \quad (61)$$

$$\eta_{max}(B) = \max_{\theta \neq 0 \in \mathbb{R}^{d_{\theta}}} \frac{\theta^{\top} B \theta}{\theta^{\top} \theta}. \quad (62)$$

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.61 and Eq.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

D.3 Lemma D.3

Lemma D.3. Under Assumptions 1, 2, and 4:

$$\mathbb{E}_{\varepsilon^n} \left[\sup_{\theta \in \Theta} |\Lambda_k(\theta)| \mid S, \sigma^{data} \right] \leq C_A \sqrt{\frac{1 + \ln(2)}{n}},$$

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:

$$\begin{aligned}|\mathbb{A}_k(\tilde{\theta})| &\leq \frac{1}{n} \sum_{i=1}^n |\underline{\omega}| \|\hat{Z}_i - Z_i\|_1 \|\tilde{Z}_i\|_\infty Z_{ik} \\ &\leq |\underline{\omega}| \bar{z}^2 \frac{1}{n} \sum_{i=1}^n \|\hat{Z}_i - Z_i\|_1,\end{aligned}\tag{63}$$

where $\bar{z} := \max_{i=1, \dots, n} \|Z_i\|_\infty$. Recall the definition of \hat{Z}_i from Eq. 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 D.5, and defining $\bar{m} := \max_{i,j \in \mathcal{N}} |m_{ij}|$, we have

$$\mathbb{E}_{\varepsilon^n} [\|\hat{Z}_i - Z_i\|_1 | S, \sigma^{data}] \leq 3\bar{m} C_\sigma \sqrt{\frac{1 + \ln(2)}{n}}. \quad (64)$$

Plug Eq.64 into Eq.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

D.4 Lemma D.4

Lemma D.4. Define $\bar{z} := \max_{i=1,\dots,n} \|Z_i\|_\infty$. Under Assumption 1 to 5, we have:

$$\mathbb{E}_{\varepsilon^n} \left[\sup_{\theta \in \Theta} |\mathbb{B}_k(\theta)| | S, \sigma^{data} \right] \leq \frac{1}{\sqrt{n}} \left(1 + 2C_4 \sqrt{d_\theta \log(1 + C_5 \sqrt{n})} \right),$$

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,

$$\begin{aligned} |\mathbb{B}_k(\tilde{\theta})| &= |\mathbb{B}_k(\tilde{\theta}) - \mathbb{B}_k(\theta^\ell) + \mathbb{B}_k(\theta^\ell)| \\ &\leq |\mathbb{B}_k(\tilde{\theta}) - \mathbb{B}_k(\theta^\ell)| + |\mathbb{B}_k(\theta^\ell)| \\ &\leq \sup_{\substack{\gamma, \gamma' \in \Theta \\ \|\gamma - \gamma'\|_1 \leq \delta}} |\mathbb{B}_k(\gamma) - \mathbb{B}_k(\gamma')| + \max_{\ell=1,\dots,H} |\mathbb{B}_k(\theta^\ell)|. \end{aligned} \quad (65)$$

Apply Lemma D.6 to bound the first term in Eq.65:

$$\mathbb{E}_{\varepsilon^n} \left[\sup_{\substack{\gamma, \gamma' \in \Theta \\ \|\gamma - \gamma'\|_1 \leq \delta}} |\mathbb{B}_k(\gamma) - \mathbb{B}_k(\gamma')| | S, \sigma^{data} \right] \leq |\underline{\omega}| \bar{z}^2 \delta, \quad (66)$$

where $\bar{z} := \max_{i=1, \dots, n} \|Z_i\|_\infty$. To bound the second term in Eq.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 D.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 E.5), we have:

$$\mathbb{E} \left[\max_{\ell=1, \dots, H} |\mathbb{B}_k(\theta^\ell)| \middle| S, \sigma^{data} \right] \leq \frac{2\tau}{\sqrt{nv}} \sqrt{\log(N_c(\delta, \Theta, L_1))}. \quad (67)$$

Now, apply Lemma D.8 to bound the L_1 -metric entropy $\log(N_c(\delta, \Theta, L_1))$:

$$\log(N_c(\delta, \Theta, L_1)) \leq d_\theta \log \left(1 + \frac{2}{\delta} \right). \quad (68)$$

Combining Eq.67 with Eq.68, we have:

$$\mathbb{E} \left[\max_{\ell=1, \dots, H} |\mathbb{B}_k(\theta^\ell)| \middle| S, \sigma^{data} \right] \leq \frac{2\tau \sqrt{d_\theta}}{\sqrt{nv}} \sqrt{\log \left(1 + \frac{2}{\delta} \right)}. \quad (69)$$

Combining Eq.65 with Eq.66 and Eq.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 $C_{B2} = 2|\underline{\omega}| \bar{z}^2$. □

D.5 Lemma D.5

Lemma D.5. Under Assumption 2, 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 2 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,

$$\mathbb{E}(|\hat{\sigma}_i^{data} - \sigma_i^{data}|) \leq \sqrt{\mathbb{E}(|\hat{\sigma}_i^{data} - \sigma_i^{data}|^2)} \leq C_\sigma \sqrt{\frac{1 + \ln(2)}{n}}.$$

□

D.6 Lemma D.6: Lipschitz Property

Lemma D.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:

$$\sup_{\substack{\gamma, \gamma' \in \Theta \\ \|\gamma - \gamma'\|_1 \leq \delta}} |\mathbb{B}_k(\gamma) - \mathbb{B}_k(\gamma')| \leq |\underline{\omega}| \bar{z}^2 \delta.$$

Proof. First, we have:

$$\sup_{\substack{\gamma, \gamma' \in \Theta \\ \|\gamma - \gamma'\|_1 \leq \delta}} |\mathbb{B}_k(\gamma) - \mathbb{B}_k(\gamma')| \leq \sup_{\substack{\gamma, \gamma' \in \Theta \\ \|\gamma - \gamma'\|_1 \leq \delta}} |\nabla_k \mathbb{M}(\gamma) - \nabla_k \mathbb{M}(\gamma')| + \sup_{\substack{\gamma, \gamma' \in \Theta \\ \|\gamma - \gamma'\|_1 \leq \delta}} |\nabla_k M(\gamma') - \nabla_k M(\gamma)|. \quad (70)$$

The first term in Eq.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.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,

$$|\nabla_k \mathbb{M}(\gamma) - \nabla_k \mathbb{M}(\gamma')| \leq \frac{1}{n} \sum_{i=1}^n |Y_i \omega_0(Z_i^\top \hat{\gamma}) - (1 - Y_i) \omega_1(Z_i^\top \hat{\gamma})| \cdot |Z_{ik}| \cdot |Z_i^\top (\gamma - \gamma')|. \quad (71)$$

By Assumption 4, $\omega_0(a) < 0$ and $\omega_1(a) > 0$ for all $a \in \mathbb{R}$. Recall that

$$\underline{\omega}_0 := \min_{x \in \Xi} \omega_0(x), \quad \underline{\omega}_1 := \min_{x \in \Xi} -\omega_1(x), \quad \underline{\omega} := \min\{\underline{\omega}_0, \underline{\omega}_1\}. \quad (72)$$

Combining Eq.71 with Eq.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.70 gives:

$$\sup_{\substack{\gamma, \gamma' \in \Theta \\ \|\gamma - \gamma'\|_1 \leq \delta}} |\mathbb{B}_k(\gamma) - \mathbb{B}_k(\gamma')| \leq |\underline{\omega}| \bar{z}^2 \delta.$$

□

D.7 Lemma D.7: Sub-Gaussian Process

Lemma D.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 E.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 $\underline{F}_\varepsilon := \min_{z \in \mathcal{Z}} \min_{\theta \in \Theta} F_\varepsilon(z^\top \theta)$, and $\bar{F}_\varepsilon := \max_{z \in \mathcal{Z}} \max_{\theta \in \Theta} F_\varepsilon(z^\top \theta)$. 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

D.8 Lemma D.8: Covering Number

Lemma D.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 E.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,

$$N_c(\delta, \Theta, L_1) \leq (1 + \frac{2}{\delta})^{d_\theta}.$$

□

Appendix E Results from Previous Literature

Lemma E.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 E.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 E.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 E.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 E.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,*

$$(i) \quad \mathbb{E}(\max_{i \leq n} Y_i) \leq \lambda \sqrt{2 \ln n},$$

$$(ii) \mathbb{E}(\max_{i \leq n} |Y_i|) \leq \lambda \sqrt{2 \ln(2n)} \leq 2\lambda \sqrt{\ln(n)}.$$

Lemma E.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

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