Policy Targeting under Network Interference*

Davide Viviano †

First version: June, 2019
This Version: September, 2021

Abstract

This paper studies the problem of optimally allocating treatments in the presence of spillover effects, using information from a quasi-experiment. We introduce a method that maximizes the sample analog of average social welfare when spillovers occur. Our procedure presents several attractive features for applications: (i) it does not necessitate network information of the target population; (ii) it exploits heterogeneity in treatment effects for targeting individuals; (iii) it does not rely on the correct specification of a particular structural model; (iv) it accommodates arbitrary constraints on the policy function. We construct semi-parametric welfare estimators with known and unknown propensity scores and cast the optimization problem into a mixed-integer linear program, which can be solved using off-the-shelf algorithms. We derive a strong set of guarantees on the regret, i.e., the difference between the maximum attainable welfare and the welfare evaluated at the estimated policy. An application for targeting information on social networks illustrates the advantages of the method.

Keywords: Causal Inference, Welfare Maximization, Spillovers, Social Interactions.
JEL Codes: C10, C14, C31, C54.

*Previous versions can be found at https://arxiv.org/abs/1906.10258. The replication code is available on the author's website. I am especially grateful to Graham Elliott, and James Fowler, Paul Niehaus, Yixiao Sun, and Kaspar Wüthrich for their advice and support. I thank Brendan Beare, Jelena Bradic, Guido Imbens, Toru Kitagawa, Craig Mcintosh, Karthik Muralidharan, James Rauch, Aleksey Tetenov for helpful comments. I particularly thank Vikram Jambulapati for invaluable discussions at the beginning of this project. I also thank participants at the EGSC Conference at the University of Washington in St. Louis 2019, at the Causal Inference with Interactions conference at Cemmap, at the Young Economists' Symposium at UPenn 2020, at the Econometric Society World Congress 2020, at the ASSA meeting 2021, at the 6th Annual Conference of Network Science and Economics, at the North American Econometric Society summer meeting 2021, and at the EEA-ESEM 2021 conference. All mistakes are my own.

†Department of Economics, University of California, San Diego. Correspondence: dviviano@ucsd.edu.

1 Introduction

One of the goals of researchers or policy-makers is to estimate the most effective (welfare-maximizing) policy. Consider a policy-maker who must use a *quasi*-experiment, such as an existing experiment or observational study, to choose how to allocate treatments on a target population optimally. The policymaker's goal is to design a decision rule (policy) that, based on observable characteristics, chooses whom to treat in order to maximize welfare. The main challenge is that treating an individual may generate spillovers on her friends or neighbors. This may, in turn, affect the design of the optimal policy. This paper studies the problem of choosing whom to treat (targeting) in the presence of spillover effects using a quasi-experiment. Relevant applications include public policy programs, cash transfer programs, educational programs, information campaigns, to cite some.[1]

We consider a setting where units are connected in a network. Individuals are assigned treatments whose effects are assumed to propagate locally in the network (i.e., to their neighbors). Researchers sample n units from an experiment or a quasi-experiment, collecting information on their covariates, treatment assignments, outcomes, covariates, and assignments of their neighbors. We do not require knowledge of the entire population network. Researchers' goal is to estimate a treatment allocation rule for new applications. For example, consider the problem discussed in Cai et al. (2015) on studying the effect of information sessions on insurance take-up in villages subject to environmental disasters. In this case, our method uses the experimental data collected by Cai et al. (2015) to estimate which individuals the insurance company should target in new villages.

The first challenge for policy targeting is that the network on the target units may be unobserved due to the cost associated with collecting it on large populations.[2] For instance, in the example from Cai et al. (2015), the insurance company has only access to the network information of the experimental participants. It is costly or infeasible to collect network information from all target individuals in a region or country. Considering this, we develop a method that allows for arbitrary constraints

[1] Some examples of relevant applications are Barrera-Osorio et al. (2011); Egger et al. (2019); Opper (2016); Zubcsek and Sarvary (2011); Bond et al. (2012). Spillover effects have been documented in development economics (Banerjee et al., 2013), social economics (Sobel, 2006), medicine (Christakis and Fowler, 2010) among many others.

[2] The reader may refer to Breza et al. (2020) for a discussion on the cost associated with collecting network information.

on the policy space, with the network not necessarily observed on the target sample. We leverage the dependence between covariates and network information on the in-sample units to optimally target individuals.

A second challenge is heterogeneity in treatment effects: an individual may respond differently to her and her neighbors' treatment depending on her type (e.g., her age or education). We impose and leverage the exogenous and anonymous interference assumption[3], often documented in applications[4], to estimate welfare without imposing conditions on the individual heterogeneity in treatment effects.

Our method, entitled Network Empirical Welfare Maximization (NEWM), estimates the welfare of the network as a function of the policy, using arbitrary (machine-learning) estimators; it then solves an optimization procedure over the policy space using an exact optimization algorithm. We interpret the policy targeting as a treatment choice problem (Manski, 2004; Kitagawa and Tetenov, 2018; Athey and Wager, 2021), which we extend to the case of network interference. We evaluate the method's performance based on its maximum regret, i.e., on the difference between the largest achievable welfare and the welfare from deploying the estimated policy function.

From a theoretical perspective, we make three contributions: (i) we derive the first set of guarantees on the regret for treatment rules with spillover effects; (ii) we introduce an estimation procedure to achieve the $n^{-1/2}$-convergence rate for estimation with machine learning estimators and networked units; (iii) we show that for a large class of policy functions, the optimization problem can be written as a mixed-integer linear program, which can be solved using off-the-shelf optimization routines.

First, we discuss the identification of social welfare under network interference for a generic network formation model. Identification relies on the unconfoundedness of treatment assignments and the conditional network exogeneity, i.e., potential

[3]In particular, we assume that potential outcomes are functions of individual and neighbors' treatment assignments (but not neighbors' identity), the number of neighbors, arbitrary individual characteristics which are observable in the experiment but may not be observed on the target sample (e.g., covariates, centrality measures, or summary statistics of neighbors' covariates), and exogenous unobservables. See Manski (2013) for a discussion on exogenous and anonymous interference.

[4]Cai et al. (2015) show that inviting individuals to information sessions generates large effects on their friends' insurance take-up in subsequent sessions. However, an individual's take-up decision does not depend on her friends' take-up decision, consistently with an exogenous interference model. Examples of studies with models consistent with the assumption of anonymous and exogenous interference include Sinclair et al. (2012); Duflo et al. (2011); Muralidharan et al. (2017), where for the second reference, networks can be considered groups of classrooms with units within each classroom being fully connected. Athey et al. (2018) provide a general framework for testing the exogenous and anonymous assumption.

outcomes independent of the neighbors' identity, conditional on individual observable characteristics.[5] We then study semi-parametric (machine-learning) estimators for the welfare and analyze the performance of the estimated policy. We show that under regularity conditions, the regret of the estimated policy scales at the minimax rate $1/\sqrt{n}$, under bounded degree. In the presence of an unbounded degree, we characterize the rate as a function of the maximum degree.

New challenges in our derivations are: (i) individuals' dependence on neighbors' assignments that we control through contraction inequalities; (ii) statistical dependence, which invalidates standard symmetrization arguments, and which is arbitrary unconditionally on the network. Also, with networked units, semi-parametric (machine-learning) estimators may present non-vanishing bias even when cross-fitting methods are employed (Chernozhukov et al., 2018; Athey and Wager, 2021). We introduce a novel cross-fitting algorithm for networked observations to control the bias and achieve the regret's minimax rate.[6]

A condition for these results to hold is that the optimization procedure achieves the in-sample optimum. We guarantee it by casting the problem in a mixed-integer linear program, which can be solved using off-the-shelf algorithms. We show that we can achieve a linear representation of spillover effects in the objective function by introducing an additional set of linear constraints and binary decision variables.

Data from Cai et al. (2015) illustrate the advantages of the method. The NEWM method leads to (out-of-sample) improvements up to twelve percentage points compared to methods that ignore network effect (Akbarpour et al., 2018; Kitagawa and Tetenov, 2018; Athey and Wager, 2021). We obtain these improvements despite network information *not* being accessible on the target villages. This is a crucial advantage whenever policymakers have limited information for targeting.

We organize the paper as follows. In Section 2, we discuss the identification condition, the causal estimands, and the welfare function. Estimation and theoretical analysis is contained in Section 3. Optimization is contained in Section 4; Section 5.1 contains extensions under non compliance; Section 5.2 discusses the case of target

[5]Network exogeneity is often and explicitly stated when inference is performed unconditionally on the network structure (Leung, 2020). The reader may refer to Goldsmith-Pinkham and Imbens (2013) for a discussion. Network exogeneity is attained if, for example, two individuals form a link based on observable individual-level characteristics and unobservables, which are independent of potential outcomes.

[6]The algorithm consists of coloring individuals, assigning the same colors to units far enough in the network, and estimating different nuisance functions using units assigned to the same color.

units drawn from a different distribution than sample units; Section 6 contains an application. Section 7 concludes. In the Appendix we include additional extensions, numerical studies and formal derivations.

1.1 Related literature

This paper builds on the growing literature on statistical treatment choice. Examples include Kitagawa and Tetenov (2018), Kitagawa and Tetenov (2019), Athey and Wager (2021), Mbakop and Tabord-Meehan (2016). A list of additional references on optimal treatment allocation includes Armstrong and Shen (2015), Bhattacharya and Dupas (2012), Hirano and Porter (2009), Stoye (2009), Stoye (2012), Tetenov (2012) among others. Further connections are more broadly related to the literature on classification (Elliott and Lieli, 2013). Our focus is quite different from previous references: we estimate an optimal policy when treatments generate spillovers. This paper is the first to study the properties of targeting under network interference in the context of the empirical welfare maximization literature.

It is important to distinguish our framework from the *i.i.d.* setting of multi-valued treatments of Kitagawa and Tetenov (2018); Zhou et al. (2018). A first conceptual difference is that here individuals depend on neighbors' assignments, while treatments are individual-specific. This structure permits the network on the target units not to be observable. The second difference is that individuals exhibit dependence and hence standard theoretical arguments based on *i.i.d.* sampling fail. This motivates (a) a novel estimation algorithm[7]; (b) theoretical arguments that exploit properties of the network in our derivations.[8] The optimization program differs due to the dependence of individuals on neighbors' assignments.

A contribution of this paper is to bridge the literature on causal inference and statistical treatment choice with the literature on network interference and influence maximization. This latter strand of literature mostly focuses on detecting the most influential "seeds" based on particular notions of centrality (Bloch et al., 2017) but often ignores heterogeneity in treatment effects and constraints on the policy space. Recent advances include Jackson and Storms (2018), Akbarpour et al. (2018),

[7]Namely when using standard cross-fitting for policy learning (Athey and Wager, 2021), the predicted propensity score and conditional mean of individual i is dependent with unobservables and observables of i since it depends on observables of i's friend, j, who is statistically dependent on i.

[8]While spillovers affect the complexity of the function class of interest, dependence invalidates standard theoretical arguments based on i.i.d. sampling (e.g., symmetrization).

Banerjee et al. (2019), Banerjee et al. (2014), Galeotti et al. (2020) among others. Examples in computer science and marketing are Kempe et al. (2003), Eckles et al. (2019) and references therein. A disadvantage of using centrality measures for targeting is that an empirical comparison is possible only between two or a few decision rules through cluster experiments (Banerjee et al., 2019; Kim et al., 2015; Chin et al., 2018). Our contribution is complementary to this literature. Three main differences are: we estimate policies from an infinite (but constrained) set of candidates from quasi-experimental studies; our approach leverages local interference to take into account heterogeneity; we design policies that depend on arbitrary information from observable characteristics (which may also depend on the network when observable).

Spillover effects have been studied in the context of policy intervention from different angles. Bhattacharya et al. (2019) and Wager and Xu (2020) study the effect of a *global* policy intervention on social welfare, using partial identification and a sequential experiment respectively. This paper considers instead *individualized* policy interventions, and point identification due to the local interference assumption. In the presence of spillover effects, Li et al. (2019), Graham et al. (2010), Bhattacharya (2009) consider the problem of optimal *allocation of individuals* across *independepent* groups. Differences are: (i) policy functions denote group assignment mechanisms, instead of binary treatment allocations, inducing a different definition and identification of the welfare function; (ii) the allocation does not allow for constrained environments, and (iii) the authors assume a clustered network structures with small independent clusters. Su et al. (2019) discuss instead a closed-form expression for treatment allocations under interference in an *unconstrained* environment, assuming linearity. In this paper, instead, we do not impose such structural assumptions. The constraints lead to a different estimation strategy, and justify the regret analysis discussed in the current paper. Laber et al. (2018) consider a Bayesian structural model whose estimation relies on computational intensive Monte Carlo methods and the correct model specification.

We build a connection to the econometric and statistical literature on inference on networks, including literature on social interaction (Manski, 2013; Manresa, 2013; Auerbach, 2019), and more generally causal inference under interference (Liu et al., 2019; Li et al., 2019; Hudgens and Halloran, 2008; Goldsmith-Pinkham and Imbens, 2013; Sobel, 2006; Sävje et al., 2020). The exogenous and anonymous interference condition is most closely related to Leung (2020). However, knowledge of treatment effects is not sufficient to construct optimal treatment rules in the presence of either (or both) constraints on the policy functions or treatment effects heterogeneity. We use the notion of non-parametric estimators from the literature on causal inference (Aronow et al., 2017) for the construction of welfare. We study semi-parametric estimators in the context of policy-targeting and devise a novel algorithm to achieve the minimax rate of convergence of the regret. We also contribute to the literature on interference and non-compliance (see, e.g., Kang and Imbens 2016), proposing a framework for policy choice when non compliance occurs.[9]

New papers have followed our work, studying targeting on networks in new directions. In a more recent paper, Kitagawa and Wang (2020) discuss the problem of targeting in a parametric setting in the presence of a SIR network. The advantage of a model-based approach is the lack of a pilot experiment, but its validity depends on the specific application in mind. Ananth (2021) builds on our work, introducing more general treatment configurations arbitrarily dependent on the network information of the target units. A disadvantage of assignments arbitrarily dependent on the target population network is that these are infeasible in applications when the network of the target units is costly or impossible to collect. When feasible, these assignments may be prone to overfitting.[10] In our framework, individuals can be targeted based on observable covariates, which may also include individual-specific network statistics (e.g., centrality measures) when available from the target sample. Finally, in Viviano (2020) I address the different problem of experimental design, which, albeit agnostic on the network of sampled units, uses an adaptive experiment that does not apply to quasi-experimental studies.

2 Set up and identification

This section discusses the main identification conditions, the causal estimands of interest, and defines utilitarian welfare under interference.

[9]Details are included in Section 5.1.

[10]Overfitting can be the result of a violation of the bounded VC-dimension unless the network enters through some statistics.

2.1 Experiment and sampling

The policy targeting exercise considered in this paper consists of two steps. In the first step, researchers collect data from an experiment or quasi-experiment. In the second step, they estimate a policy recommendation to be implemented on an independent sample $j \in \mathcal{I}$.

We start introducing the notation. Each individual is associated with an arbitrary vector of covariates $Z_i \in \mathcal{Z}$ This vector is random, and it may contain individual, neighbors' covariates statistics, and network statistics. Therefore, the entries of Z_i may exhibit dependence. We impose restrictions in Assumption 3.2.

There are E (finitely) many individuals who participate in the experiment and who constitute an independent network. These individuals are connected under an adjacency matrix A. We define

$$A \in \mathcal{A} \subseteq \{0,1\}^{E \times E}, \quad N_i = \Big\{k \in \{1, \cdots, E\} \setminus i : A_{i,j} = 1\Big\}, \quad Z = (Z_i)_{i=1}^{E},$$

where $\mathcal{A}$ denotes the set of symmetric and unweighted adjacency matrices, N_i denotes all neighbors of individual i, and Z the matrix of covariates of all participants.[11] For example, in Cai et al. (2015) $\{1, \cdots, E\}$ denotes all individuals in those villages in the experiment. In full generality, we let edges be drawn as $A_{i,j} \sim \mathcal{P}_{i,j}$, with $\mathcal{P}_{i,j}$ being left unspecified, and arbitrary dependent. Define $|N_i|$ the cardinality of the set N_i.

For some function $f_D(\cdot)$ known in an experiment, and to be estimated in a quasi-experiment, a binary treatment is assigned to each experimental participant as follows:

$$D_i = f_D(Z_i, \varepsilon_{D_i}), \quad \varepsilon_{D_i} \sim_{iid} \mathcal{D}, \quad i \in \{1, \cdots, E\}. \tag{1}$$

We discuss explicit conditions on ε_{D_i} in Section 2.2. Researchers sample $n \leq E$ experimental participants at random, and observe the vector

$$\Big(Y_i, Z_i, Z_{j \in N_i}, D_i, D_{j \in N_i}, N_i\Big)_{i=1}^{n}, \qquad \text{for some } n \leq E$$

where Y_i denotes the post-treatment outcome.[12]

[11] Namely, we say that two individuals (i, j) are connected if $A_{i,j} = 1$, where $A_{i,j}$ denote the edge between individual i and individual j, with $A_{i,j} = A_{j,i} \in \{0, 1\}$ and $A_{i,i} = 0$.

[12] Such a sample can be constructed as follows: researchers sample n random vertices; for each of these vertices, they observe the individual outcome, covariates, treatment assignments, and neighbors' identities; for each individual, they then collect information of neighbors' treatment assignments

2.2 Interference

In the presence of interference, the outcome of individual i depends on its treatment assignment and all other units' treatment assignments. For a given vector $D^E = (D_i)_{i=1}^E$ denoting the vector of treatment assignments of experimental participants, we let

$$Y_i = \tilde{r}_E(i, D^E, A, Z_i, \varepsilon_i), \quad i \in \{1, \cdots, E\} \tag{2}$$

for some (unobserved) random variables ε_i. The above equation states that experimental partipants depend on the assignment of all units in the experiment. In the following condition, we impose that the outcomes only depend on the treatment assigned to their first-degree neighbors, and we make assumptions on interactions being anonymous (Manski, 2013).

Assumption 2.1 (Interference)**.** For all $(v, u) \in \{1, \cdots, E\}^2$, for any $D^E = (D_k)_{k=1}^E \in \{0,1\}^E$, $\tilde{D}^E = (\tilde{D}_k)_{k=1}^E \in \{0,1\}^E$ and $A, \tilde{A} \in \mathcal{A}$,

$$\tilde{r}_E(v, D^E, A, z, e) = \tilde{r}_E(u, \tilde{D}^E, \tilde{A}, z, e)$$

for all $z \in \mathcal{Z}$, $e \in \Big\{\text{supp}(\varepsilon_v) \cup \text{supp}(\varepsilon_u)\Big\}$, if all the following three conditions hold: (i) $\sum_k A_{v,k} = \sum_k \tilde{A}_{u,k}$; (ii) $\sum_k A_{v,k} D_k^E = \sum_k \tilde{A}_{u,k} \tilde{D}_k^E$; (iii) $D_v^E = \tilde{D}_u^E$.

Assumption 2.1 postulates that outcomes only depend on (i) the number of first-degree neighbors, (ii) the number of first-degree treated neighbors (iii) individual's treatment status as well as covariates of interest. Under Assumption 2.1 we can write for each individual $i \in \{1, \cdots, E\}$

$$Y_i = r\Big(D_i, \sum_{k \in N_i} D_k, Z_i, |N_i|, \varepsilon_i\Big), \tag{3}$$

for some possibly unknown function $r(\cdot)$.

Assumption 2.1 states that interactions are anonymous (Manski, 2013), and spillovers occur within neighbors only. Under the above condition, the outcome can depend on the number of treated neighbors, the share of treated neighbors, or whether at least one neighbor is treated. Heterogeneity is allowed and captured by the dependence of $r(\cdot)$ with Z_i and $|N_i|$. The assumption is consistent with findings in Cai et al. (2015),

and covariates.

where the authors document exogenous interference and lack of endogenous spillover effects. The model is closely related to Leung (2020), and Athey et al. (2018) provide a general framework for testing anonymous and local interference.

We complement Equation (3) with the following condition.

Assumption 2.2 (Unconfoundedness)**.** For some arbitrary function $f_D(.)$, let the following hold: for all $i \in \{1, \cdots, E\}$,

(A) $D_i = f_D(Z_i, \varepsilon_{D_i})$ and ε_{D_i} are *i.i.d.*, and $\varepsilon_{D_i} \perp \Big(Z, A, (\varepsilon_j)_{j=1}^E\Big)$;

(B) For each i, $\varepsilon_i \perp \Big(A, Z\Big)\Big| Z_i, |N_i|$;

Condition (A) states that the treatment is randomized in the experiment based on observable Z_i. Since Z_i may contain network information of a given individual, the assumption also accommodates randomization schemes where treatment assignment is based on network information. Condition (B) imposes that the network and others' covariates are independent of unobservables. Network exogeneity is attained if, for example, two individuals form a link based on observable individual-level characteristics and unobservables, which are independent of potential outcomes.[13] We illustrate an example of a network formation model satisfying Assumption 2.2.

Example 2.1 (Network Formation)**.** *Let* $A_{i,j} = g(Z_i, Z_j, \alpha_{i,j}), (Z_i, \varepsilon_i) \sim_{i.i.d.} \Gamma$ *for some unknown function* $g(\cdot)$ *and* $\alpha_{i,j}$ *denoting exogenous unobservables jointly independent of observables and* $(\varepsilon_j)_{j=1}^E$ *(but possibly dependent with each other). Then Condition (B) in Assumption 2.2 holds.*

2.3 Policy targeting: problem description

The planner's goal is to design a treatment mechanism that maximizes social welfare on units $j \in \mathcal{I}$, where $\mathcal{I}$ denotes the target population. For instance, in Cai et al. (2015), $\mathcal{I}$ defines the units in new villages where the experiment has not been conducted. The target population is separated from the experimental participants, and connected under a matrix $G \in \mathcal{G} \subset \{0, 1\}^{|\mathcal{I}| \times |\mathcal{I}|}$ possibly unknown to the researcher.

With an abuse of notation, we will denote for all $j \in \mathcal{I}$, $N_j = \{k \in \mathcal{I} \backslash j : G_{j,k} = 1\}$, the set of neighbors of individual j. The matrix G and covariates $(Z_j)_{j \in \mathcal{I}}$ are realized

[13] Network exogeneity is often stated when inference is performed unconditionally on the network structure (Leung, 2020). The reader may refer to Goldsmith-Pinkham and Imbens (2013) for a discussion.

before treatments are assigned. The policy-maker has limited information on the target sample. Namely, she has only access to the variables

$$X_j \subseteq Z_j, \quad X_j \in \mathcal{X} \subseteq \mathcal{Z}, \quad j \in \mathcal{I},$$

denoting a *subset* of individuals' characteristics, which may also include arbitrary network information which is included in Z_j. She *does not* necessarily observe neighbors' identity of the target participants.

Her goal is to design a policy as follows:

(A) Individuals may be treated differently, depending on observables characteristics;

(B) The assignment mechanism must be easy to implement, without requiring knowledge of the population network and covariates of all other individuals;

(C) The assignment mechanism can be subject to arbitrary constraints (e.g., ethical or economic constraints).

We formalize the above conditions by defining an *individualized* treatment assignment

$$\pi : \mathcal{X} \mapsto \{0, 1\}, \quad \pi \in \Pi, \tag{4}$$

where Π denotes the set of constraints imposed on the decision function.[14] The map amounts to a partition of $\mathcal{X}$, the support of X_j. The policy function may depend on arbitrary observables X_j, such as measures of centrality (i.e., network statistics), or covariates available to the policy-maker. This policy function (i) can be implemented in online fashion, individual by the individual; (ii) it is feasible since it does not require knowledge of the population network and covariates of all units.[15]

The decision rule $\pi(\cdot)$, states that each unit $j \in \mathcal{I}$ is treated according to $D_j = \pi(X_j)$ for all $j \in \mathcal{I}$. The planner's goal is to choose the policy π that maximizes welfare. The utilitarian welfare is defined as the expected outcome once we fix $D_j = \pi(X_j)$ for all $j \in \mathcal{I}$. Namely,

$$W(\pi) = \frac{1}{|\mathcal{I}|} \sum_{j \in \mathcal{I}} \mathbb{E}\Big[Y_j \Big| \Big\{(D_j)_{j\in\mathcal{I}} = (\pi(X_j))_{j \in \mathcal{I}}\Big\}\Big].$$

[14] An example is the capacity constraint of the form $\int_{x \in \mathcal{X}} \pi(x) dF_X(x) \leq K$ for a constant K (Kitagawa and Tetenov, 2018).

[15] It also avoids overfitting with a domain constant in the size of the population.

Importantly, the expectation is not conditional on a particular value of the treatment assignments $(\pi(X_j))_{j\in\mathcal{I}}$, but instead, it is conditional on the event that the treatment assignment mechanism is such that $D_j = \pi(X_j)$ for each unit in the target sample.

Welfare effects depend on two main components: (i) the direct effect of the treatment since each treatment is chosen based on $\pi(X_j)$, and (ii) the spillover effect since each neighbors' exposures also depend on the policy π. For example, welfare on the *experimental* participants, reads as follows

$$\frac{1}{n}\sum_{i=1}^{n}\mathbb{E}\Big[r\Big(\pi(X_i), \sum_{k\in N_i}\pi(X_k), Z_i, |N_i|, \varepsilon_i\Big)\Big], \tag{5}$$

that denotes the effect of π also mediated through neighbors' assignments. Its definition follows from the deterministic nature of the policy $\pi \in \Pi$.

Unfortunately, we can only identify the welfare of the experimental participants. Discrepancies between the expected welfare of sample units and the targeted welfare $W(\pi)$ will reflect in the quality of the estimated policy function. We introduce the following definition.

Definition 2.1 (Sample-Target Discrepancy)**.** Define

$$\mathcal{K}_\Pi(n) = \sup_{\pi\in\Pi}\Big|\frac{1}{n}\sum_{i=1}^{n}\mathbb{E}\Big[r\Big(\pi(X_i), \sum_{k\in N_i}\pi(X_k), Z_i, |N_i|, \varepsilon_i\Big)\Big] - W(\pi)\Big|,$$

for some function $\mathcal{K}_\Pi(.) : \mathbb{Z} \mapsto \mathbb{R}_+$.

Here, $\mathcal{K}_\Pi(.)$ denotes the difference between the expected welfare of in-sample units and the welfare of the target population. Such difference captures the bias induced by estimating the policy function on units that are possibly drawn from a different population than target units.

While we derive our results also as a function of the discrepancy, we can guarantee that $\mathcal{K}_\Pi(\cdot) = 0$ under the following conditions.

Assumption 2.3 (Local interference on the target population)**.** For a given target population $\mathcal{I}$, and for each (deterministic) policy function $\pi : \mathcal{X} \mapsto \{0, 1\}, \pi \in \Pi$ let

$$W(\pi) = \frac{1}{|\mathcal{I}|}\sum_{j\in\mathcal{I}}\mathbb{E}\Big[r\Big(\pi(X_j), \sum_{k\in N_j}\pi(X_k), Z_j, |N_j|, \varepsilon_j\Big)\Big] \tag{6}$$

for some unobservables $(\varepsilon_j)_{j\in\mathcal{I}}$.

Assumption 2.3 states that the local interference model in Assumption 2.1 also holds on the target sample. In the following lemma, we provide conditions for the discrepancy to be equal to zero.

Lemma 2.1. *Let Assumption 2.1, 2.2, 2.3 hold. Suppose in addition that* $(Z_i, Z_{k\in N_i}, |N_i|, \varepsilon_i) \sim \mathcal{P}$ *for all* $i \in \{1, \cdots, E\} \cup \mathcal{I}$*. Then* $\mathcal{K}_\Pi(n) = 0$*, i.e., for all* $\pi \in \Pi$*,*

$$W(\pi) = \frac{1}{n}\sum_{i=1}^{n} \mathbb{E}\Big[r\Big(\pi(X_i), \sum_{k\in N_i} \pi(X_k), Z_i, |N_i|, \varepsilon_i\Big)\Big].$$

Lemma 2.1 shows that $\mathcal{K}_\Pi(n)$ is equal to zero whenever the vector of unobservables ε_i, individual covariates Z_i, number of neighbors, and neighbors' covariates are drawn jointly from the same distribution. To gain further intuition, note that the condition in the above example is satisfied if the conditional distribution of $(Z_i, \varepsilon_i, Z_{N_i})\Big||N_i| = l \sim \mathcal{P}_1(l)$, for some distribution $\mathcal{P}_1(l)$ which depends on the number of connections of an individual, and the *unconditional* distribution of the degree is the same across units.

Remark 1 (Units drawn from a different distribution)**.** In Section 5.2 we allow sampled and target units to have covariates and degrees drawn from different distributions. □

Remark 2 (An extension under spillovers on non-compliance)**.** Consider the case where spillovers also occur over individuals' compliance. Namely, let $D_i \in \{0, 1\}$ denote the assigned treatment and $T_i \in \{0, 1\}$ denotes the selected treatment from individual i. We model non-compliance as follows:

$$Y_i = r\Big(T_i, \sum_{k\in N_i} T_k, Z_i, |N_i|, \varepsilon_i\Big), \quad T_i = h_\theta\Big(D_i, \sum_{k\in N_i} D_k, Z_i, |N_i|, \nu_i\Big), \quad \nu_i \sim_{i.i.d.} \mathcal{V} \quad (7)$$

where ν_i denote an exogenous unobservables, independent from ε_i, and $(r(\cdot), \theta)$ are unknown, with θ denoting the set of parameters indexing h. Intuitively, treatment effects and compliance depend on neighbors' selected and assigned treatments. In Section 5.1 we discuss conditions for identification of welfare effects, assuming that researchers collect information also on second-degree neighbors. □

Remark 3 (Extension with higher-order interference)**.** We can allow for higher-order interference by assuming that researchers also observe second-degree neighbors. The outcome model reads as follows

$$Y_i = r\Big(D_i, \sum_{k \in N_i} D_k, \sum_{k \in N_{i,2}} D_k, Z_i, |N_i|, |N_{i,2}|, \varepsilon_i\Big),$$

where $N_{i,2}$ denotes the second-order degree neighbors. Intuitively, the outcome depends on the number of first-degree neighbors and the number of second-degree neighbors. The welfare reads as follows.

$$W_2(\pi) = \frac{1}{|\mathcal{I}|} \sum_{i \in \mathcal{I}} \mathbb{E}\Big[r\Big(\pi(X_i), \sum_{k \in N_i} \pi(X_k), \sum_{k \in N_{i,2}} \pi(X_k), Z_i, |N_i|, |N_{i,2}|, \varepsilon_i\Big)\Big].$$

□

2.4 Identification

We now turn to the problem of identifying the welfare of the experimental participants, which, under Lemma 2.1 equals the welfare of the target units. To achieve this goal, we first need to identify the conditional mean function and the propensity score for each unit. We introduce two additional conditions.

Assumption 2.4. For all (i, j), $P(\varepsilon_i \leq t | Z_i = z, |N_i| = l) = P(\varepsilon_j \leq t | Z_j = z, |N_j| = l)$ for all $z \in \mathcal{Z}$, $l \in \mathbb{Z}$, $t \in \text{supp}(\varepsilon_i) \cup \text{supp}(\varepsilon_j)$.

The condition states that unobservables are identically (but not necessarily independently) distributed.

Assumption 2.5 (M-Local Dependence)**.** Assume that for any $A \in \mathcal{A}, i \in \{1, \cdots, E\}$ and some possibly unknown $M \geq 2$,

$$\varepsilon_i \perp (\varepsilon_j)_{j \notin N_{i,M-1}} | A, Z,$$

where $N_{i,M-1}$ denote the set of nodes connected to i by at most $M-1$ edges.[16]

Assumption 2.5 states that unobservables are conditionally independent for individuals that are not in a neighborhood up with radius $M-1$.[17]

[16] Formally, such set is defined as the set $\{k : A^{M-1}_{i,k} \neq 0\}$.

[17] Local dependency graphs are often assumed in the presence of network data, see, e.g., Leung (2020).

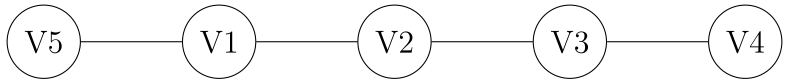

Figure 1: Example of network. Under Assumption 2.5 with $M = 4$, the unobservables ε_i are independent only between V5 and V4, while they can be dependent for any other pair of nodes.

Example 2.2 (Graphical Illustration)**.** *Consider Figure 1: under Assumption 2.5, for $M = 4$, for each node the corresponding vector ε_i are independent only between V5 and V4, while they can be dependent for any other pair of nodes.*

Next, we discuss the causal estimands of interest.

Definition 2.2 (Conditional Mean)**.** Under Assumption 2.4 the conditional mean function for each $i \in \{1, \cdots, E\}$, is defined as

$$m(d, s, z, l) = \mathbb{E}\Big[r(d, s, Z_i, |N_i|, \varepsilon_i)\Big| Z_i = z, |N_i| = l\Big]. \tag{8}$$

Assumption 2.4 is necessary in order for the conditional mean function to be identical across units. The second causal estimand of interest is the propensity score.

Definition 2.3. Under Condition (i) in Assumption 2.2, for each $i \in \{1, \cdots, E\}$ we define[18]

$$\begin{aligned} e(d, s, \mathbf{x}, z, l) &= P\Big(D_i = d, \sum_{k \in N_i} D_k = s \Big| Z_{k \in N_i} = \mathbf{x}, Z_i = z, |N_i| = l\Big) \\ &= P\Big(D_i = d | Z_i = z\Big) \sum_{u_1, \cdots, u_l : \sum_v u_v = s} \prod_{k=1}^{l} P\Big(D_{N_i^{(k)}} = u_k \Big| Z_{N_i^{(k)}} = \mathbf{x}^{k,\cdot}\Big). \end{aligned} \tag{9}$$

for $d \in \{0, 1\}, s \in \mathbb{Z}, s \leq l$.

Definition 2.3 defines the probability of treatment given individual and neighbors covariates. Condition (i) in Assumption 2.2 guarantees that the propensity score does not depend on the index of unit i or of its neighbors. The definition of the propensity score follows from the literature on multi-valued treatments (Imbens, 2000), while here, the individuals' exposures also depend on the random number of neighbors and

[18] The second equaition decomposes the probability of the event into sums of probabilities of disjoint events, each corresponding to a given combination of treatment assignments $(u_1, \cdots, u_l)$, whose sum equals to s.

neighbors' covariates. Here, the number of treated neighbors $\sum_{k\in N_i}\pi(X_k)$ depends on the array of covariates $X_{k\in N_i}$, via the policy function π. If we were to construct the propensity score, which did not also depend on such covariates, we would not identify welfare effects. We can now state the following identification result.

Lemma 2.2 (Identification)**.** *Let Assumption 2.1, 2.2, 2.4 hold. Let $S_i(\pi)=\sum_{k\in N_i}\pi(X_k)$. Then for any function $\pi\in\Pi$*

$$\begin{aligned}\frac{1}{n}\sum_{i=1}^n\mathbb{E}\Big[r(\pi(X_i),\sum_{k\in N_i}\pi(X_k),Z_i,|N_i|,\varepsilon_i)\Big] &= \frac{1}{n}\sum_{i=1}^n\mathbb{E}\Big[m\Big(\pi(X_i),\sum_{k\in N_i}\pi(X_k),Z_i,|N_i|\Big)\Big]\\ &= \frac{1}{n}\sum_{i=1}^n\mathbb{E}\Big[\frac{1\{S_i(\pi)=\sum_{k\in N_i}D_k,\pi(X_i)=D_i\}}{e\Big(\pi(X_i),S_i(\pi),Z_{k\in N_i},Z_i,|N_i|\Big)}Y_i\Big].\end{aligned}$$

The proof is in Appendix C.2. The above result guarantees the identification of the welfare of the experimental participants. The result does *not* require conditions on the distribution of covariates and on the dependence across unobservables.

3 Network Empirical Welfare Maximization

We now discuss the procedure and its theoretical properties. We defer a discussion on the optimization method to Section 4. We state our results as non-asymptotics, valid in finite sample.

3.1 Empirical welfare

We start introducing some notation. We let $S_i(\pi)=\sum_{k\in N_i}\pi(X_k)$ be the assigned treatment to neighbors of i under policy π.

A nonparametric estimator of the welfare function is denoted as

$$W_n^{ipw}(\pi,e)=\frac{1}{n}\sum_{i=1}^n\frac{1\{S_i(\pi)=\sum_{k\in N_i}D_k,\pi(X_i)=D_i\}}{e\Big(\pi(X_i),S_i(\pi),Z_{k\in N_i},Z_i,|N_i|\Big)}Y_i, \tag{10}$$

which also depends on the propensity score e. One possible disadvantage of the above estimator is the large variance. Therefore, we also consider the following double robust estimator (AIPW) of the welfare function of the following form:

$$W_n^{aipw}(\pi, m^c, e^c) = \frac{1}{n}\sum_{i=1}^n \frac{1\{S_i(\pi) = \sum_{k\in N_i} D_k, \pi(X_i) = D_i\}}{e^c\Big(\pi(X_i), S_i(\pi), Z_{k\in N_i}, Z_i, |N_i|\Big)}\Big(Y_i - m^c\Big(\pi(X_i), S_i(\pi), Z_i, |N_i|\Big)\Big) + \frac{1}{n}\sum_{i=1}^n m^c\Big(\pi(X_i), S_i(\pi), Z_i, |N_i|\Big), \tag{11}$$

for some arbitrary functions (m^c, e^c) which we will refer to as the pseudo-true conditional mean and propensity score. The welfare inherits double robust properties similarly to what discussed in Robins et al. (1994), i.e. if either $m^c = m$ or $e^c = e$, the welfare is equal in expectation to Equation (5) (see Appendix C.2). Throughout our discussion, the estimated conditional mean and propensity score will be denoted as $\hat{m}$ and $\hat{e}$, respectively. The propensity score can be easily estimated as a function of conditional marginal probabilities (see Definition 2.3).

3.2 Known propensity score

We start discussing theoretical guarantees of the policy that maximizes $W_n^{aipw}(\pi, m^c, e^c)$, for some *arbitrary* functions m^c and e^c, which, only throughout this subsection, are assumed not to be data-dependent unless otherwise specified. Formally,

$$\hat{\pi}^{aipw}_{m^c,e^c} \in \arg\max_{\pi\in\Pi} W_n^{aipw}(\pi, m^c, e^c), \tag{12}$$

defines the "oracle" policy, which has access to the pseudo-true functions (m^c, e^c).

Assumption 3.1. Let the following hold:

(LP) $m^c(d, s, z, l)$ is L-Lipschitz a.e. in its second argument, for some $\infty > L \geq 0$;

(OV) For all $d \in \{0, 1\}$, there exist some $\delta \in (0, 1)$ such that $e^c\Big(d, s, \mathbf{x}, z, l\Big) \in (\delta, 1-\delta)$ for all $s \leq l$, for all $z \in \mathcal{Z}, l \in \mathbb{Z}, \mathbf{x} \in \mathcal{Z}^l \times \mathbb{Z}^l$.

(TC) For $\Gamma_1^2, \Gamma_2^2 < \infty$, $\forall i \in \{1, ..., n\}$, $\mathbb{E}\Big[\sup_{d\in\{0,1\}, s\leq |N_i|} r(d, s, Z_i, |N_i|, \varepsilon_i)^3 \Big| A, Z\Big] < \Gamma_1^2$, and $\mathbb{E}\Big[\sup_{d\in\{0,1\}} m^c(d, 0, Z_i, |N_i|)^3 \Big| A, Z\Big] < \Gamma_2^2$, almost surely for each $i \in \{1, \cdots, E\}$.

(VC) π belongs to a function class of point-wise measurable functions[19] Π, where Π has finite VC dimension.[20]

Condition (LP) is satisfied for *any* bounded $m^c(.)$[21], but it also accomodates for unbounded $m^c(.)$. Remarkably, the condition is agnostic on the *true* conditional mean function $m(.)$. Condition (OV) is the usual overlap condition, often imposed in the causal inference literature. We discuss trimming at the end of this subsection. Condition (TC) imposes moment conditions on the outcome. Condition (VC) imposes restrictions on the geometric complexity of the function class of interest. It is commonly assumed in the literature on empirical welfare maximization, and examples include Kitagawa and Tetenov (2018), and Athey and Wager (2021), among others. For instance, if individuals are assigned based on a threshold crossing rule, the VC dimension equals the number of variables used in the assignment. Under condition (OV) in Assumption 3.1 we define δ_0 a constant such that $\max_{d\in\{0,1\}} e^c(d, 0, Z_{k\in N_i}, Z_i, |N_i|) \in (\delta_0, 1-\delta_0)$ almost surely, where $\delta_0 \geq \delta$ by definition.[22]

Covariates can be endogenous to the network and exhibit arbitrary dependence. In the following lines, we impose conditions on the distribution (and dependence) of covariates.

Assumption 3.2 (Covariates' distribution)**.** Assume that we can write $Z_i = h(Q_i, L_i)$, for some unknown function $h(\cdot)$, where $Q_i \in \mathcal{Q}, L_i \in \mathcal{L}$, where

(A) $(Q_j)_{j=1}^E \perp (\varepsilon_j)_{j=1}^E \Big| A, (L_j)_{j=1}^E, \quad Q_j \perp Q_{k\notin N_{j,M}} | A, (L_j)_{j=1}^E, \quad (Q_j, Q_{k\in N_j}) | A, (L_k)_{k=1}^E \sim \mathcal{T}(|N_j|)$;

(B) $|\mathcal{L}| < \infty$,

for some possibly unknown distribution $\mathcal{T}$.

[19]Point-wise measurability can be replaced by measurability of each function $\pi \in \Pi$, but, in this latter case, since the pointwise supremum may not be necessarily measurable the supremum function must be interpreted as the *lattice* supremum. See Appendix F.

[20]The VC dimension denotes the cardinality of the largest set of points that the function π can shatter. The VC dimension is commonly used to measure the complexity of a class. See for example Devroye et al. (2013).

[21]To see why the claim hold, let $|m^c(d, S, Z_i)| < B$ a.e. Since $S \in \mathbb{Z}$, $|m^c(d, S, Z_i) - m^c(d, S', Z_i)| \leq 2B|S - S'|$.

[22]To observe why, note that under (OV) δ defines the lowest propensity score over all possible configurations of number of treated neighbors, also including the case where none of the neighbors is treated.

Assumption 3.2 decomposes covariates into two main components which are possibly *unknown* to the researcher. The first component Q_i has arbitrary support, but it is exogenous to the network structure and locally dependent. The second component L_i can instead exhibit arbitrary dependence, but it has finite support. Intuitively, Assumption 3.2 states that either we have finitely many "types" of individuals, or we have infinitely many, but continuities are not predictive of the network formation. In the presence of continuous variables arbitrarily dependent on the network, the assumption holds after binning.

Example 3.1 (Finitely many types)**.** Suppose that $|\mathcal{Z}| < \infty$, i.e., the support of Z_i is finite dimensional. Then Assumption 3.2 holds. □

Let $\mathcal{N}_E = \max_{i \in \{1, \cdots, E\}} |N_i| + 1$ denote the maximum degree of the network of experimental participants[23], and $\mathcal{N}_n = \max_{i \in \{1, \cdots, n\}} |N_i| + 1$ the maximum degree of the sampled units. We can now state the first theorem.

Theorem 3.1 (Oracle Regret)**.** *Let Assumptions 2.1, 2.2, 2.4, 2.5, 3.1, 3.2 hold. Assume that either (or both) (i) $m^c(.) = m(.)$ and/or (ii) both $e^c(.) = e(.)$. Then,*

$$\mathbb{E}\Big[\sup_{\pi \in \Pi} W(\pi) - W(\hat{\pi}^{aipw}_{m^c, e^c})\Big] \leq \mathbb{E}[\mathcal{P}_{M, |\mathcal{L}|}(\mathcal{N}_E)] \frac{\bar{C}(L+1)(\Gamma_1 + \Gamma_2)}{\delta \delta_0^{\sqrt{3}}} \sqrt{\frac{\text{VC}(\Pi)}{n}} + 2\mathcal{K}_\Pi(n),$$

for a finite constant $\bar{C} < \infty$ independent of $(L, \Gamma_1, \Gamma_2, \delta, \delta_0, E, n)$ and $\mathcal{P}_{M, |\mathcal{L}|}$ being a polynomial function with finite degree.

Corollary (Known propensity score)**.** *Let $\hat{\pi}^{ipw} \in \arg\max_{\pi \in \Pi} W_n^{ipw}(\pi, e)$ for known propensity score and let the conditions in Theorem 3.1 and Lemma 2.1 hold. Suppose that $|N_i| < J < \infty$ almost surely. Then*

$$\mathbb{E}\Big[\sup_{\pi \in \Pi} W(\pi) - W(\hat{\pi}^{ipw})\Big] \leq \bar{C}' \sqrt{\text{VC}(\Pi)/n}$$

for a constant $\bar{C}' < \infty$ indendent of (n, E).

The proof of the theorem is in Appendix D. Theorem 3.1 provides a non-asymptotic upper bound on the regret, and it is the first result of this type under network inter-

[23] In our proofs, we only use information from the maximum degree of the n sampled units and their neighbors up to the Mth degree. Here we instead consider the maximum degree $\mathcal{N}_E$ of all experimental participants, which is larger than what is considered in our proofs, for expositional convenience only.

ference. The theorem is double robust to the misspecification of m^c and $1/e^c$. The corollary provides regret guarantees when researchers maximize the welfare with a known propensity score. The bound in the above corollary is distribution-free (i.e., it holds for all data-generating processes satisfying the conditions in the corollary). For bounded degree, which in our application using data from Cai et al. (2015) equals five, the regret scales at rate $1/\sqrt{n}$, which has been shown to match the maximin lower-bound in the *i.i.d.* with no-interference (Kitagawa and Tetenov, 2018). We observe that assuming that sampled units are representative of the target population (Assumption 2.3) is not invoked in Theorem 3.1 since the regret is expressed as a function of the discrepancy between the target and sample units $\mathcal{K}_\Pi(n)$. However, it is assumed in Lemma 2.1 and the corollary. The proof consists of (i) controlling the Rademacher complexity in the presence of spillover effects, using extensions of contraction inequalities to accommodate lack of sub-exponential tail decay; (ii) dealing with dependence for symmetrization, and bounding the groups of dependent units using the chromatic number.

A direct corollary of Theorem 3.1 is that the bound holds if the propensity score is known and the conditional mean estimated on an independent sample. The polynomial function depends on the degree of dependence and on the number of endogenous types captured by $|\mathcal{L}|$ (see Assumption 3.2). In Section 3.4.2 we illustrate how the rate can improve in terms of the maximum degree under weaker dependence conditions.

3.3 Estimated nuisances and network cross-fitting

Next, we discuss regret guarantees in the presence of unknown propensity score, with the estimated policy being

$$\hat{\pi}^{aipw}_{\hat{m},\hat{e}} \in \arg\max_{\pi \in \Pi} W_n^{aipw}(\pi, \hat{m}, \hat{e}). \tag{13}$$

A key challenge is represented by the dependence structure.

We propose a modification of the *cross-fitting* algorithm (Chernozhukov et al., 2018), to account for the local dependence of observations and a possibly fully connected network. The algorithm assumes that the researcher knows M and observes the network of the experimental participants (but not necessarily of the target units). Extensions, when the dependence structure is unknown, are discussed in Section 3.4.3.

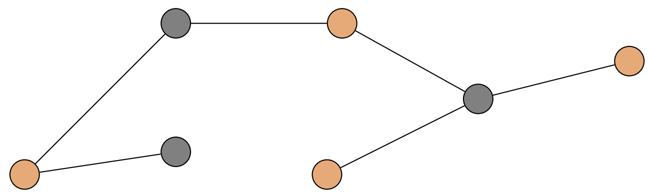

Figure 2: Network cross-fitting: simple illustration with $M = 1$ (note that in practice $M \geq 2$). For each colored units, estimators are constructed using information from the other units assigned to the same color.

Algorithm 1 : Network Cross-Fitting

1. Create K folds and assign to each fold individuals who are not neighbors and do not share common neighbors up to the M^{th} degree conditional on A;

2. For each unit j in fold k, estimate the conditional mean, and the propensity score using all observations *within* the fold k *only*, with the exception of j.

To the best of our knowledge, Algorithm 1 is novel to the literature on network interference.[24] The key intuition of the above algorithm is to construct folds of independent observations and estimate via leave-one-out each conditional mean *within* each fold. The condition of individuals not sharing common neighbors guarantees that also treatment assignments are independent.

In the following assumption, we discuss conditions on the convergence rate of the proposed procedure.

Assumption 3.3. Assume that for each $d \in \{0, 1\}, s \in \mathbb{Z}$, $\mathcal{A}_n \times \mathcal{B}_n = \mathcal{O}(n^{-v})$ for some $v \geq 1/2$, where

$$\begin{aligned} \mathcal{A}_n &= \sqrt{\frac{1}{n}\sum_{i=1}^n \mathbb{E}\Big[\sup_{d,s}\Big(\hat{m}(d, s, Z_i, |N_i|) - m(d, s, Z_i, |N_i|)\Big)^2\Big]} \\ \mathcal{B}_n &= \sqrt{\frac{1}{n}\sum_{i=1}^n \mathbb{E}\Big[\sup_{d,s}\Big(\frac{1}{\hat{e}(d, s, Z_{k\in N_i}, Z_i, |N_i|)} - \frac{1}{e(d, s, Z_{k\in N_i}, Z_i, |N_i|)}\Big)^2\Big]}. \end{aligned} \tag{14}$$

Assume in addition that $\hat{m}$ and $1/\hat{e}$ are uniformly bounded by a finite constant.

[24] Alternative algorithmic procedures have been proposed for multi-way and clustered data (Chiang et al., 2019), while here we consider the different setting of networked observations, which requires using chromatic properties of the network in the construction of the fitting procedure, and introduce novel trade-offs in the rate of convergence that depend on the structure of the network.

Under the above conditions, we can state the following theorem. The proof is contained in Appendix D.

Theorem 3.2. *Let Assumption 2.1, 2.2, 2.4, 2.5, 3.1, 3.2, 3.3 hold, with* $m^c = m, e^c = e$. *Let estimation being performed as in Algorithm 1. Then*

$$\mathbb{E}\Big[\sup_{\pi \in \Pi} W(\pi) - W(\hat{\pi}^{aipw}_{\hat{m},\hat{e}})\Big] \leq \bar{C}' \mathbb{E}[\mathcal{P}_{M,|\mathcal{L}|}(\mathcal{N}_E)] \sqrt{\frac{\mathrm{VC}(\Pi)}{n}} + 2\mathcal{K}_\Pi(n)$$

for a constant $\bar{C}' < \infty$ *independent of* (n, E), *and* $\mathcal{P}_{M,|\mathcal{L}|}$ *being a polynomial function with finite degree.*

Theorem 3.2 shows that the rate of convergence of the estimator does not affect the convergence rate of the regret under Assumption 3.3. Under the conditions in Lemma 2.1, $\mathcal{K}_\Pi(n) = 0$, and hence the regret converges to zero as the degree grows at an appropriate slower rate than the sample size.

Remark 4 (Assumption 3.3 and increasing degree)**.** Assumption 3.3 imposes conditions on the *product* of the convergence rate of the estimator to the *true* conditional mean and propensity score function, in the same spirit of standard conditions in the i.i.d. setting (e.g., Farrell 2015). Observe that the condition in Assumption 3.3 is satisfied for general machine-learning estimators under bounded degree, since, in the presence of bounded degree, by Brooks (1941)'s theorem, we can construct finitely many partitions of independent observations. Assumption 3.3 instead imposes a faster rate of convergence than $n^{-1/4}$ for both the propensity score and the conditional mean function, as the maximum degree is increasing. For instance, letting $M = 2$ (i.e., individuals are dependent with friends and friends of friends), the number of partitions is of order $\mathcal{N}_E^2$. In this cases, each nuisance function is estimated on a sample with $n/\mathcal{N}_E^2$ many observations, and hence Assumption 3.3 is satisfied, if each nuisance convergences at a rate $n^{1/4}\mathcal{N}_E^2$. This equals the parametric rate if $\mathcal{N}_E = E^{1/8}$ and $n = E$ (all participants are sampled). However, if the degree grows at a faster rate than $E^{1/8}$ Assumption 3.3 can be relaxed to hold for an arbitrary rate n^{-v} and the regret also depends on such a rate. We omit this case for the sake of brevity only. □

In Section 3.4.3 we discuss instead the case where the network of the sampled units is not accessible, and the degree of dependence M is unknown.

3.4 Three additional results

Next, we discuss three additional results: targeting with poor overlap; faster rates with clustered networks; estimation with a partially observed network.

3.4.1 Trimming to control overlap

Strict overlap can be violated in the presence of few nodes with an unbounded degree. We address this challenge by proposing a trimming estimator. To guarantee overlap, we introduce the following trimming estimator:

$$W_n^{tr}(\pi;\kappa_n) = \frac{1}{n}\sum_{i=1}^{n}\Big\{\frac{1\{S_i(\pi) = \sum_{k\in N_i} D_k, \pi(X_i) = D_i\}}{e\Big(\pi(X_i), S_i(\pi), Z_{k\in N_i}, Z_i, |N_i|\Big)} Y_i \times 1\Big\{|N_i| \leq \kappa_n\Big\}\Big\}, \quad (15)$$

for a particular choice of κ_n which we discuss below. The above estimator accounts for the spillover effect of treating nodes with a large set of neighbors to their connections, but it excludes from estimation the direct effect on the largely connected nodes.

Theorem 3.3. *Let* $\hat{\pi}_{\kappa_n}^{tr} \in \arg\max_{\pi\in\Pi} W_n^{tr}(\pi;\kappa_n)$ *and suppose that* $e(d,s,\mathbf{x},z,l) \in (\delta_n, 1-\delta_n)$ *for all* $d \in \{0,1\}, \mathbf{x} \in \mathcal{Z}^l, z \in \mathcal{Z}, s \leq \kappa_n, l \leq \kappa_n$*, and assume that* $P(|N_i| \leq \kappa_n) > c > 0$*, for some* $c \in (0,1)$*. Suppose that conditions in Theorem 3.1 hold. Suppose in addition* $Y_i \in [-B, B]$ *for a universal constant* B*. Let the conditions in Lemma 2.1 hold. Then the following holds:*

$$\mathbb{E}\Big[\sup_{\pi\in\Pi} W(\pi) - W(\hat{\pi}_{\kappa_n}^{tr})\Big] \leq \frac{\bar{C}\mathbb{E}[\mathcal{P}_{M,|\mathcal{L}|}(\mathcal{N}_E)]B}{\delta_n}\sqrt{\frac{\mathrm{VC}(\Pi)}{n}} + 2BP(|N_i| > \kappa_n)$$

for a finite constant $\bar{C} < \infty$ *independent of* (E, n)*.*

The proof of the theorem is in Appendix D. Theorem 3.3 shows that the rate of convergence still depends polynomially on the maximum degree, but the overlap constant δ_n is potentially much larger than the worst-case constant δ. The additional price to pay is that the regret also depends on the probability that the degree exceeds a certain threshold. Such probability can be small whenever the *number* of nodes with a large degree grows at a slower rate than the sample size. A simple example is $\sqrt{n}$-many individuals (assuming $n = E$) having a growing degree greater than κ, in which case $P(|N_i| > \kappa) = O(1/\sqrt{n})$.

Remark 5 (Increasing overlap by restricting the number of treatment exposures)**.** Additional modeling restrictions may be imposed to reduce the dimensionality of the problem and guarantee strict overlap (also without trimming). For example, suppose that the researcher imposes the following restriction, for some ordered τ_1, τ_2, τ_3:

$$r(d, s, z, l, e) = \begin{cases} \bar{r}_1(d, z, l, e) \text{ if } s/l \leq \tau_1 \\ \bar{r}_2(d, z, l, e) \text{ if } \tau_1 < s/l \leq \tau_2 \\ \bar{r}_3(d, z, l, e) \text{ if } \tau_2 < s/l \leq \tau_3 \end{cases} \tag{16}$$

for some possibly unknown functions $\bar{r}_1, \bar{r}_2, \bar{r}_3$. Then the number of possible exposures reduces to six different exposures. In this case, the balancing score defines the probability of each of this exposure, having a larger probability compared to the propensity score in Definition 2.3. □

3.4.2 Faster rates with growing degree under weaker dependence

A natural question is whether we can achieve faster rates as a function of the degree. We formalize this in the following theorem, where we show that weaker dependence conditions guarantee a faster rate in the degree.

Theorem 3.4. *Let Assumption 2.1, 2.2, 2.4, 2.5, 3.1 hold. Suppose that* $K_i = \Big(Y_i, D_i, D_{k \in N_i}, Z_i, Z_{k \in N_i}, |N_i|\Big) \sim \mathcal{P}$ *and that we can partition* $(K_i)_{i=1}^n$ *into* K *groups each containing mutually independent observations. Let either* $m^c = m$ *or* $e^c = e$*. Then*

$$\mathbb{E}\Big[\sup_{\pi \in \Pi} W(\pi) - W(\hat{\pi}^{aipw}_{m^c, e^c})\Big] \leq \bar{C} K \sqrt{\frac{\mathbb{E}[\log(\mathcal{N}_n)\mathcal{N}_n]}{n}}$$

for a finite constant $\bar{C}$ *independent of* (n, E, K)*.*

The proof is contained in Appendix D. Such stronger result requires the additional restriction: we can construct finitely many groups of vectors $\Big(Y_i, D_i, D_{k \in N_i}, Z_i, Z_{k \in N_i}, |N_i|\Big)$ with units independent within each group (as in the case of many small clusters). The result shows that the degree affects the function class complexity *sublinearly*, while the remaining polynomial terms in Theorem 3.1 capture the dependence structure of the data. From an inspection of the proof the reader may observe that the result holds regardless of whether Z_i is discrete or continuous.

Given a simpler dependence structure, for Theorem 3.4, the proof follows similarly to the *i.i.d.* case of Kitagawa and Tetenov (2018) in its symmetrization argument, but with an important modification that uses contraction arguments, extended to deal under lack of sub-exponential tails' decay. To gain further intuition on its main motivation, consider first as a "loose" definition of the Rademacher complexity

$$\mathcal{R}_n(\Pi) = \mathbb{E}_\sigma\Big[\sup_{\pi \in \Pi}\Big|\frac{1}{n}\sum_{i=1}^n \sigma_i f_i\Big(\sum_{k \in N_i} \pi(X_i), \pi(X_i)\Big)\Big|\Big] \tag{17}$$

for some data-dependent functions $f_i(\cdot)$ Lipschitz in their second argument, and $(\sigma_i)_{i=1}^n$ denoting *i.i.d.* Rademacher random variables. For example, with a bounded outcome f_i may be defined as the outcome multiplied by its corresponding inverse probability weight.

In standard *i.i.d.* settings Equation (17) bounds (up-to a constant term) the regret for any realization of the data, with $f_i(\cdot, \pi(X_i))$ not depending on the first entry due to lack of interference. However, under interference, the conditional Rademacher complexity as in Equation (17) is not a valid upper bound of the regret due to dependence. It is, however, a valid upper bound under lack of statistical dependence among observations, used in Theorem 3.4.

This result illustrates the trade-off among different assumptions: when individuals exhibit arbitrary network dependence as in Theorem 3.1, and independence only conditional on the network, standard symmetrization arguments cannot directly be used, and the regret also depends on higher moments of the network of all experimental participants. This instead does not occur if stronger independence conditions are imposed.

3.4.3 Regret with partially observed networks

We conclude this section by characterizing the regret rate with M unknown, and information of the neighbors (but not of the entire network of the sampled participants) is accessible. We assume that the researcher estimates nuisance functions using the entire sample, and we characterize the rate as a function of the convergence rate of each estimator. We first impose the following condition.

Assumption 3.4. For some $\xi_1, \xi_2 > 0$,

$$\mathbb{E}\Big[\frac{1}{n}\sum_{i=1}^{n}\sup_{d\in\{0,1\},s\leq|N_i|}\Big|\hat{m}(d,s,Z_i,|N_i|) - m^c(d,s,Z_i,|N_i|)\Big|\Big] = \mathcal{O}(1/n^{\xi_1}).$$
$$\mathbb{E}\Big[\frac{1}{n}\sum_{i=1}^{n}\sup_{d\in\{0,1\},s\leq|N_i|}\Big|\Big(Y_i - m^c(d,s,Z_i,|N_i|)\Big)\Big(e^c(d,s,O_i) - \hat{e}(d,s,O_i)\Big)\Big|\Big] = \mathcal{O}(1/n^{\xi_2}). \tag{18}$$

where $O_i = (Z_{k\in N_i}, Z_i, |N_i|)$. In addition, assume that $\hat{e}(\cdot) \in (\delta, 1-\delta)$ almost surely.

Assumption 3.4 imposes conditions on the convergence rate of $\hat{m}$ to its *pseudo*-true value m^c and the convergence rate of $\hat{e}$ to the pseudo true e^c.[25] Here, we do not require the network cross-fitting algorithm for our results to hold, at the expense of a slower convergence rate for semi-parametric estimators.

Theorem 3.5. *Let Assumption 2.1, 2.2, 2.4, 2.5, 3.1, 3.2, 3.4 hold. Assume that either $m^c(.) = m(.)$ or $e^c(.) = e(.)$. Then,*

$$\mathbb{E}\Big[\sup_{\pi\in\Pi} W(\pi) - W(\hat{\pi}^{aipw}_{\hat{m},\hat{e}})\Big] \leq \mathcal{O}\Big(1/n^{\xi} + \mathbb{E}[\mathcal{P}_{M,|\mathcal{L}|}(\mathcal{N}_E)]\sqrt{\frac{\mathrm{VC}(\Pi)}{n}}\Big) + 2\mathcal{K}_{\Pi}(n),$$

where $\xi = \min\{\xi_1, \xi_2\}$, for $\mathcal{P}_{M,|\mathcal{L}|}$ being a polynomial function with finite degree.

Theorem 3.5 provides a uniform bound on the regret, and it is double robust to correct specification of the conditional mean and the propensity score. The theorem's result depends on the convergence rate of $\hat{e}$ and $\hat{m}$ to their *pseudo*-true value. For parametric estimators of the conditional mean and the propensity score and bounded degree, the regret bounds scale at rate $1/\sqrt{n}$. However, for the general machine-learning estimator, the rate can be slower than the parametric one, reflecting the "cost" of the lack of knowledge of the degree of dependence M.

4 Mixed-integer linear program formulation

In this section, we discuss the optimization procedure. Firstly, we define the estimated effect of assigning to unit i treatment d, after treating s of its neighbors. For the double

[25]Convergence rate for penalized regression on networks is found in He and Song (2018) among others. Estimation via the method of moments that satisfy the high-level conditions in Hansen (1982) also guarantees parametric convergence rates $1/\sqrt{n}$ of the estimators of interest.

robust estimator, this quantity is defined as

$$g_i(d, s) = \frac{1\{\sum_{k \in N_i} D_k = s, D_i = d\}}{e^c\Big(d, s, Z_{k \in N_i}, Z_i, |N_i|\Big)}\Big(Y_i - m^c\Big(d, s, Z_i, |N_i|\Big)\Big) + m^c\Big(d, s, Z_i, |N_i|\Big), \tag{19}$$

where we omit the dependence of $g_i(\cdot)$ with m^c and e^c for sake of brevity. Secondly, we define $I_i(\pi, h) = 1\Big\{\sum_{k \in N_i} \pi(X_k) = h\Big\}$ the indicator of whether h neighbors of individual i have been treated under policy π. We observe that the following holds:

$$\sum_{h=0}^{|N_i|}\Big\{\Big(g_i(1, h) - g_i(0, h)\Big)\pi(X_i)I_i(\pi, h) + I_i(\pi, h)g_i(0, h)\Big\} = g_i\Big(\pi(X_i), \sum_{k \in N_i}\pi(X_k)\Big). \tag{20}$$

Namely, the estimated treatment effect on unit i, obtained after implementing policy π, is the sum of effects obtained by treating zero to all neighbors of i. Each element in the sum is weighted by the indicator $I_i(\pi, h)$, and only one of these indicators is equal to one. We can then define n variables p_i that denote the treatment assignment of each unit after restricting the policy function in the function class of interest. Namely, we let $p_i = \pi(X_i), \pi \in \Pi$. For example, for policy functions of the form

$$\pi(X_i) = 1\{X_i^\top \beta \geq 0\}, \quad \beta \in \mathcal{B},$$

similarly to Kitagawa and Tetenov (2018) we write

$$\frac{X_i^\top \beta}{|C_i|} < p_i \leq \frac{X_i^\top \beta}{|C_i|} + 1, \quad C_i > \sup_{\beta \in \mathcal{B}}|X_i^\top \beta|, \quad p_i \in \{0, 1\},$$

where p_i is equal to one if $X_i^\top \beta$ is positive and zero otherwise. The key intuition is to introduce additional decision variables to represent $I_i(\pi, h)$ through linear constraints. We define the following variables:

$$t_{i,h,1} = 1\Big\{\sum_{k \in N_i} p_k \geq h\Big\}, \quad t_{i,h,2} = 1\Big\{\sum_{k \in N_i} p_k \leq h\Big\}, \quad h \in \{0, \cdots, |N_i|\}.$$

The first variable is one if at least h neighbors are treated, and the second variable is one if at most h neighbors are treated. The goal is to achieve a linear representation of such variables. Observe that

$$t_{i,h,1} + t_{i,h,2} = \begin{cases} 1 \text{ if and only if } \sum_{k \in N_i} p_k \neq h \\ 2 \text{ otherwise} \end{cases} \quad \Rightarrow \quad t_{i,h,1} + t_{i,h,2} - 1 = I_i(\pi, h). \tag{21}$$

Therefore, we can define $t_{i,h,1}, t_{i,h,2}$ using mixed-integer linear constraints. Namely, the variable $t_{i,h,1}$ can equivalently be defined as[26]

$$\frac{(\sum_k A_{i,k} p_k - h)}{|N_i| + 1} < t_{i,h,1} \leq \frac{(\sum_k A_{i,k} p_k - h)}{|N_i| + 1} + 1, \quad t_{i,h,1} \in \{0, 1\}. \tag{22}$$

Similar reasoning follows for $t_{i,h,2}$.

We now can write the objective function as

$$\frac{1}{n} \sum_{i=1}^{n} \sum_{h=0}^{|N_i|} \Big\{ \Big(g_i(1,h) - g_i(0,h) \Big) p_i(t_{i,h,1} + t_{i,h,2} - 1) + (t_{i,h,1} + t_{i,h,2} - 1) g_i(0,h) \Big\}. \tag{23}$$

The above objective function leads to a quadratic mixed integer program. On the other hand, quadratic programs can be computationally expensive to solve. We write the problem as a mixed-integer linear program introducing one additional set variables, that we call $u_{i,h}$ for $h \in \{0, \cdots, |N_i|\}$, with $u_{i,h} = p_i(t_{i,h,1} + t_{i,h,2} - 1)$. We provide the complete formulation below.

$$\max_{\{u_{i,h}\}, \{p_i\}, \{t_{i,1,h}, t_{i,2,h}\}, \beta \in \mathcal{B}} \frac{1}{n} \sum_{i=1}^{n} \sum_{h=0}^{|N_i|} \Big\{ \Big(g_i(1,h) - g_i(0,h) \Big) u_{i,h} + g_i(0,h)(t_{i,h,1} + t_{i,h,2} - 1) \Big\} \tag{24}$$

[26] The equation holds for the following reason. Suppose that $h < \sum_k A_{i,k} p_k$. Since $\frac{(\sum_k A_{i,k} p_k - h)}{|N_i|+1} < 0$, the left-hand side of the inequality is negative and the right hand side is positive and strictly smaller than one. Since $t_{i,h,1}$ is constrained to be either zero or one, in the latter case, it equals zero. Suppose now that $h \geq \sum_k A_{i,k} p_k$. Then the left-hand side is bounded from below by zero, and the right-hand side is bounded from below by one. Therefore $t_{i,h,1}$ is set to be one.

under the constraints:[27]

$$
\begin{aligned}
&(A) \quad p_i = \pi(X_i), \quad \pi \in \Pi \\
&(B) \quad \frac{p_i + t_{i,h,1} + t_{i,h,2}}{3} - 1 < u_{i,h} \leq \frac{p_i + t_{i,h,1} + t_{i,h,2}}{3}, \quad u_{i,h} \in \{0,1\} \quad \forall h \in \{0, \cdots, |N_i|\}, \\
&(C) \quad \frac{(\sum_k A_{i,k} p_k - h)}{|N_i| + 1} < t_{i,h,1} \leq \frac{(\sum_k A_{i,k} p_k - h)}{|N_i| + 1} + 1, \quad t_{i,h,1} \in \{0,1\}, \quad \forall h \in \{0, \cdots, |N_i|\} \\
&(D) \quad \frac{(h - \sum_k A_{i,k} p_k)}{|N_i| + 1} < t_{i,h,2} \leq \frac{(h - \sum_k A_{i,k} p_k)}{|N_i| + 1} + 1, \quad t_{i,h,2} \in \{0,1\}, \quad \forall h \in \{0, \cdots, |N_i|\}.
\end{aligned}
\tag{25}
$$

The first constraint can be replaced by methods discussed in previous literature such as maximum scores (Florios and Skouras, 2008), whereas the additional constraints are justified by the presence of interference.[28] In the presence of capacity constraints, the problem can be formulated as above after adding additional linear constraints on the maximum number of treated units. Whenever units have no-neighbors, the objective function is proportional to the one discussed in Kitagawa and Tetenov (2018) under no interference.[29] Therefore, the formulation provided generalizes the MILP formulation to the case of interference.[30]

Theorem 4.1. *The π^* solves the optimization problem in Equation (24) under the constraints in Equation (25) if and only if $\pi^* \in argmax_{\pi \in \Pi} W_n^{aipw}(\pi, m^c, e^c)$.*

Theorem 4.1 is a direct consequence of the argument in the current section, and it permits to solve the optimization problem over the function class Π using off-the-shelf methods.

---

[27]To motivate the constraint for $u_{i,h}$, notice first that we can write $p_i(t_{i,h,1} + t_{i,h,2} - 1) = p_i \times t_{i,h,1} \times t_{i,h,2}$ since $(t_{i,h,1} + t_{i,h,2} - 1)$ is equal to one if both variables are ones and zero if either of the two variables are ones and the other is zero. The case where both variables are zero never occurs by construction. Therefore we can write $\frac{p_i + t_{i,h,1} + t_{i,h,2}}{3} - 1 < u_{i,h} \leq \frac{p_i + t_{i,h,1} + t_{i,h,2}}{3}, \quad u_{i,h} \in \{0,1\}$.

[28]In practice, we observe that including additional (superfluous) constraints stabilizes the optimization problem. These are $\sum_h (t_{i,h,1} + t_{i,h,2} - 1) = 1$ for each i and $\sum_i \sum_h u_{i,h} = \sum_i p_i$.

[29]This follows from the fact that under no interference the second component in the objective function is constant and the first component only depends on the individual treatment allocation.

[30]Also, observe that the formulation differs from those provided for allocation of an individual into small peer groups (Li et al., 2019) since the latter case does not account for the individualized treatment assignments, encoded in the constraints (A)-(D), and in the variables in the objective function $t_{i,h}$.

5 Extensions

In this section, we discuss two extensions: spillovers on non-compliance and target and sampled units having different distributions.

5.1 Spillovers on non-compliance: identification of welfare effects

In the presence of non-compliance, we interpret the problem from an intention to treat perspective. A challenge is that the treatment assigned to an individual also influences the selection into the treatment of her friends. We use as a working model the one in Equation (7). The model is consistent with an exogenous interference model over the compliance: it states that the selection into the treatment of an individual depends on her and neighbors' treatment assignment. Denote

$$T_i(\pi) = h_\theta\Big(\pi(X_i), \sum_{k \in N_i} \pi(X_k), Z_i, |N_i|, \nu_i\Big)$$

the potential selected treatment, if treatments D_i were assigned under policy $\pi(\cdot)$ (i.e., $D_i = \pi(X_i)$)). In the following theorem, we show that the utilitarian welfare is the expected outcome, once treatments D_i are assigned according to the deterministic rule $\pi(X_i)$. Denote $\mathbb{E}_\pi$ the expectation conditional on the event that $D_i = \pi(X_i)$ for all $i \in \{1, \cdots, E\}$.

Theorem 5.1. *Consider the model in Equation (7). Assume that the following holds:*

$$\varepsilon_i \perp \Big(A, Z, (\nu_j)_{j=1}^E, (\varepsilon_{D_j})_{j=1}^E\Big)\Big| Z_i, |N_i|, \quad \nu_i \perp \Big(A, Z, (\varepsilon_{D_j})_{j=1}^E\Big),$$

with ν_i being i.i.d.. Then, for each $i \in \{1, \cdots, E\}$,

$$\mathbb{E}_\pi[Y_i] = \mathbb{E}\Big[r\Big(T_i(\pi), \sum_{k \in N_i} T_k(\pi), Z_i, |N_i|, \varepsilon_i\Big)\Big] = \mathbb{E}\Big[\sum_{d \in \{0,1\}, s \in \{0, \cdots, |N_i|\}} J_i(d,s) \times H_i(d,s,\pi)\Big], \tag{26}$$

where

$$
\begin{aligned}
J_i(d,s) &= \mathbb{E}\Big[Y_i\Big|Z_i, |N_i|, T_i = d, \sum_{k\in N_i} T_k = s\Big] \\
H_i(d,s,\pi) &= P_\theta\Big(T_i = d\Big|Z_i, |N_i|, V_i(\pi)\Big)\times \\
&\sum_{u_1,\cdots,u_l:\sum_v u_v = s}\prod_{k=1}^{|N_i|} P_\theta\Big(T_{N_i^{(k)}} = u_k\Big|Z_{N_i^{(k)}}, |N_{N_i^{(k)}}|, V_{N_i^{(k)}}(\pi)\Big),
\end{aligned}
\tag{27}
$$

where $V_i(\pi) = \Big\{D_i = \pi(X_i), \sum_{k\in N_i} D_k = \sum_{k\in N_i}\pi(X_k), \pi(X_i), \sum_{k\in N_i}\pi(X_k)\Big\}$, *and* $P_\theta(T_i = 1|\cdot)$ *denotes the conditional probability of selection into treatment indexed by the parameters* θ*.*

The proof is contained in Appendix E. Theorem 5.1 defines a causal estimand in terms of simple to estimate conditional expectations. This is commonly adopted in the literature on mediation analysis and g-estimation in bio-statistics (see, e.g., Naimi et al. 2017), while here we adopt such a framework in the context of interference.

The welfare effect of an incentive π depends on conditional mean functions that can be estimated from observed data. Here $J_i(\cdot)$ denotes the conditional mean of the outcome variable given the *selected* treatment, and $H_i(\cdot)$ denotes the conditional probability of selecting into treatment, conditional on the individual and neighbors' incentives. Observe that $H_i(\cdot)$ takes a *simple* expression which only depends on the individual probability of selected treatments $P_\theta(T_i = 1|\cdot)$, conditional on individual's and neighbors' treatment assignments.

Interestingly, $H_i(\cdot)$ also depends on the treatment assigned to the second-degree neighbors. This is intuitive: an incentive to Elena affects the selection into the treatment of her friend Matteo whose decision affects the outcome of Matteo's friend Riccardo. Therefore identification requires information from first-degree and second-degree neighbors.

Contrasting with existing literature, a strand of literature on interference and non compliance includes Imai et al. (2020), Kang and Imbens (2016), Vazquez-Bare (2020), DiTraglia et al. (2020), and the recent work of Kim (2020). The above references discuss the identification of treatment effects under non-compliance without discussing the problem of welfare maximization. This is an important difference that motivates the exogenous interference model we propose in Equation (7) and the identification strategy that uses information from the second-degree neighbors.[31]

[31] We note in particular that the difference with Kim (2020) is that our model provides a simple

5.2 Different target and sample units

Finally, we study the case where target and sampled units are drawn from different populations. The following condition is imposed.

Assumption 5.1. Assume that target and sample units follow the following laws:

$$
\begin{aligned}
(Z_i, Z_{k\in N_i})\Big| |N_i| &\sim F_{s,|N_i|}, \quad |N_i| \sim G_s \quad \forall i \in \{1, \cdots, E\}, \\
(Z_i, Z_{k\in N_i})\Big| |N_i| &\sim F_{t,|N_i|}, \quad |N_i| \sim G_t \quad \forall i \in \mathcal{I}.
\end{aligned}
$$

Assumption 5.1 states that the distributions of individual covariates, neighbors' covariates, and the number of neighbors differ across the target and sampled units only. The assumption reads as follows: the joint distribution of $(Z_i, Z_{k\in N_i})$, given the degree $|N_i|$ is the same across units having the same degree and being in the same population (either target or sampled population) and similarly the marginal distribution of $|N_i|$, whereas such distributions may be different between target and sampled units.

Under the second condition in Assumption 5.1 the conditional mean function is the same on target and sampled units, whereas the distribution of covariates and network may differ. We define $f_{s,|N_i|}, f_{t,|N_i|}$ the corresponding Radon-Nikodym derivatives of $F_{s,|N_i|}, F_{t,|N_i|}$, of sampled and target units with respect to a common dominating measure on $\mathcal{Z} \times \mathcal{Z}^{|N_i|}$ and similarly g_s, g_t the corresponding Radon-Nikodym derivatives of G_s, G_t.[32] Then, we impose the following condition:

Assumption 5.2. *Assume that*

$$
f_{t,|N_i|}(Z_i, Z_{k\in N_i})g_t(|N_i|) = \rho(Z_i, Z_{k\in N_i}, |N_i|)f_{s,|N_i|}(Z_i, Z_{k\in N_i})g_s(|N_i|),
$$

with $\rho(Z_i, Z_{k\in N_i}, |N_i|) \leq \bar{\rho} < \infty$ *almost surely.*

Assumption 5.2 follows similarly to Kitagawa and Tetenov (2018), where the ratio of the two densities is assumed to be bounded almost surely. The empirical welfare

closed-form expression of potential outcomes as functions of the *realized* treatments, and exploit a local interference model also over the compliance, while Kim (2020) proposes potential outcomes as functions of expected treatments.

[32]Since $|N_i|$ is discrete, $g_s(l), g_t(l)$ corresponds to the probability of the number of neighbors being equal to l.

criterion is constructed as follows:

$$W_n^t(\pi, m^c, e^c) = \frac{1}{n}\sum_{i=1}^n \Big[\frac{1\{S_i(\pi) = \sum_{k\in N_i} D_k, \pi(X_i) = D_i\}}{e^c\Big(\pi(X_i), S_i(\pi), Z_{k\in N_i}, Z_i, |N_i|\Big)}\Big(Y_i - m^c(\pi(X_i), S_i(\pi), Z_i, |N_i|)\Big)\rho(Z_i, Z_{k\in N_i}, |N_i|)\Big] + \frac{1}{n}\sum_{i=1}^n \Big[m^c(\pi(X_i), S_i(\pi), Z_i, |N_i|)\rho(Z_i, Z_{k\in N_i}, |N_i|)\Big]. \tag{28}$$

We can now state the following theorem whose proof is contained in Appendix D.

Theorem 5.2. *Let Assumptions 2.1, 2.2, 2.3, 2.4, 2.5, 3.1, 3.2, 5.1, 5.2, and either $m^c = m$ or $e^c = e$ hold. Then we obtain that for $\hat{\pi}^{aipw}_{m^c,e^c}$ maximizing Equation (28),*

$$\mathbb{E}\Big[\sup_{\pi\in\Pi} W(\pi) - W(\hat{\pi}^{aipw}_{m^c,e^c})\Big] = \mathcal{O}\Big(\mathbb{E}[\mathcal{P}_{M,|\mathcal{L}|}(\mathcal{N}_E)]\sqrt{\frac{\text{VC}(\Pi)}{n}}\Big),$$

for a polynomial function $\mathcal{P}_{M,|\mathcal{L}|}$ with bounded degree.

6 Empirical application

We now illustrate the proposed method using data originated from Cai et al. (2015). The authors study the effect of an information session on insurance adoption in 47 villages in China, documenting (i) positive spillover effects resulting from direct treatments to neighbors; (ii) absence of endogenous spillovers.[33] To evaluate our procedure's performance, we "simulate" the following environment: researchers collect information on the first 25 villages. We estimate the policy to target the individuals in the remaining villages, which in total are 22. On the remaining villages, we assume that the policy-maker does not have access to the network information. Throughout our discussion, we consider the population of individuals having at least one neighbor.[34] Individuals are connected under a "strong" adjacency matrix, whose edges are equal to one if both individuals of a given pair indicated the other as a connection. The training set contains $n = 1315$ observations, and the maximum degree is bounded

[33]Endogenous spillovers define the effect of increasing insurance take-up as a function of the purchase decision of direct neighbors.

[34]This follows from the fact that the optimization problem over individuals having no connections can be treated as a separate problem.

by five. The test set has 1401 units. For simplicity, we assume full compliance with the treatment.[35]

The outcome of interest is insurance adoption, and it is binary. While the experiment has more than two arms, we only focus on the effects of assigning individuals to intensive information sessions for simplicity. The intensive information sessions were randomized at the household level. The experiment of Cai et al. (2015) consists of two rounds: two consecutive sets of information sessions were performed within a few days. The authors assume that spillovers occur only to individuals participating in the second information session. To capture these effects, we estimate the policy function using an asymmetric adjacency matrix, where individuals participating in the first round of information sessions have no incoming edges. We evaluate the performance of the remaining villages using the true population adjacency matrix.

Estimation of the conditional mean function is performed non-parametrically using Random Forest (Breiman, 2001).[36] We maximize welfare using the double robust estimator. We estimate the individual probability of treatment as in Equation (9) using logistic regression.[37] We compare the methods using the estimated doubly-robust estimator on observations from the remaining villages.[38]

Table 1: Percentage of individual with a certain number of neighbors. Two individuals are friends if both indicate the other as a friend.

1	2	3	4	5
53.3 %	30.3 %	12.2 %	2.99 %	0.92 %

We consider a linear policy rule of the following form:

$$\pi(X_i) = 1\Big\{\beta_0 + \text{age}\beta_1 + \text{rice\_area}\beta_2 + \text{risk\_adversion}\beta_3 \geq 0\Big\}. \tag{29}$$

[35] In the experiment, more than 90% of farmers attended the sessions.

[36] The function depends on the percentage of treated neighbors, the individual treatment assignments, and their interaction. We also condition on age, rice area, risk aversion, their interactions with the treatment assignments as well as gender, age, literacy level, the index of risk aversion, the probability of a climate disaster, education, and the number of friends of the participant.

[37] The same covariates used for the conditional mean function are also used for the propensity score.

[38] To guarantee overlap, we trim the propensity score whenever the joint probability of individual and neighbors' treatment is below 5% when evaluating the methods out-of-sample. Results are robust if we choose 2% trimming. For trimmed units, extrapolation using the conditional mean function is performed, which guarantees that we evaluate the performance on all the 1401 units.

We study the effect with four different levels of capacity constraints, namely 20%, 30%, 40% of individuals are treated.[39] We compare the proposed method to four competitors: (i) the EWM rule discussed in Kitagawa and Tetenov (2018), with estimated propensity score of individual treatment (i.e., it ignores network effects), and a policy function class as in Equation (29); (ii) the doubly robust method of Athey and Wager (2021) where, however, the conditional mean function also controls for the network information; (iii) the method that targets at random the same number of individuals as NEWM.[40] In Table 1 we report the percentage of individuals with a certain number of neighbors. The strong network presents a small degree that facilitates computations but may *under*-estimate spillover effects.

We collect results in Table 2. In the table, we report the Feasible NEWM with policy as in Equation (29) and the best competitor among a random allocation, EWM with propensity score and EWM with also a regression adjustment. Table 2 also collects results for the "oracle" NEWM method. This method consists of maximizing the double-robust welfare over each village on the *target* sample, estimating a village-specific linear decision rule which also depends on the *number of neighbors* of each individual. The method is named "oracle" since it has access to the outcomes and the network information from the target villages. Welfare is measured as 12 RMB per mu per season.[41]

We observe that the proposed method uniformly outperforms those methods that do not account for spillover effects. The improvements is up to twelve percentage points. This can be economically relevant once the policy is implemented at scale. The method underperforms relative to the oracle method since the feasible NEWM estimator does not directly use information from the target villages other than the age, rice area, and risk aversion of each individual. In Figure 3 we compare the oracle method to the feasible method under different capacity constraints. Two facts are

[39] We also compute results for capacities 50% whose results are comparable to the results with 40% constraint. We omit those since the optimization of feasible NEWM does not achieve a zero dual gap within the time constraint for 50% constraint. Also, whenever the number of treated unit on the target sample exceeds the capacity constraints, we treat the same number of units as the capacity constraint, treating those with the largest estimated score $X_i^\top \widehat{\beta}$.

[40] We estimate the model using a MILP program as in Kitagawa and Tetenov (2018) for the EWM methods and as in the main text for the NEWM method. The dual gap of each estimated method is zero, with the exception of the oracle method. For the oracle method, since we estimate different policies for different villages we incur a dual gap over a few villages. However, as noted in Table 2, the dual gap of the oracle method does not affect its performance relative to the other methods.

[41] 1 RMB is 0.15$ and one mu is 0.067 hectare

worth noticing. First, (i) the assignments under the oracle method positively correlate with the degree of individuals. However, the method does *not* treat all the units with the largest degree; instead, it balances treatments across different sub-populations to exploit heterogeneity in treatment effects. Second, (ii) the feasible method treats more individuals with the largest degree, although network information is not used by the policy function. The figure shows how the NEWM method exploits information on the dependence between the degree and observable covariates from in-sample units for best targeting individuals when network information is not directly accessible by the policy-maker. In Table 3, we report the estimated policy function's coefficients. We observe a positive dependence of the optimal treatment rule with the risk aversion and the rice area of the individual and a negative dependence with age.

In Appendix A we include a numerical study that shows that our procedure uniformly outperforms those methods that ignore network effects.

Table 2: *Out-of-sample* welfare comparisons. Welfare is measured as the probability of insurance adoption times insurance premium. Feas NEWM stands for feasible NEWM. Best competitor reports the welfare among the random assignment EWM with propensity score and double-robust EWM. Oracle is the NEWM that estimates different policies in the target villages, having access to outcomes and network information in those villages. Third column reports the welfare improvement between the feasible NEWM and the best competitor. Different rows correspond to the case where capacity constraints. 90% lower bound denotes the lower bound constructed using cluster robust standard errors.

	Feas NEWM	Best Comp	Welfare Imp	Oracle NEWM
Welf: Capacity 0.2	2.168	2.118	2%	4.219
90% *Lower bound*	1.18	1.13		
Welf: Capacity 0.3	2.574	2.299	12%	4.842
90% *Lower bound*	2.23	1.92		
Welf: Capacity 0.4	2.781	2.492	12%	5.363
90% *Lower bound*	1.96	1.87		

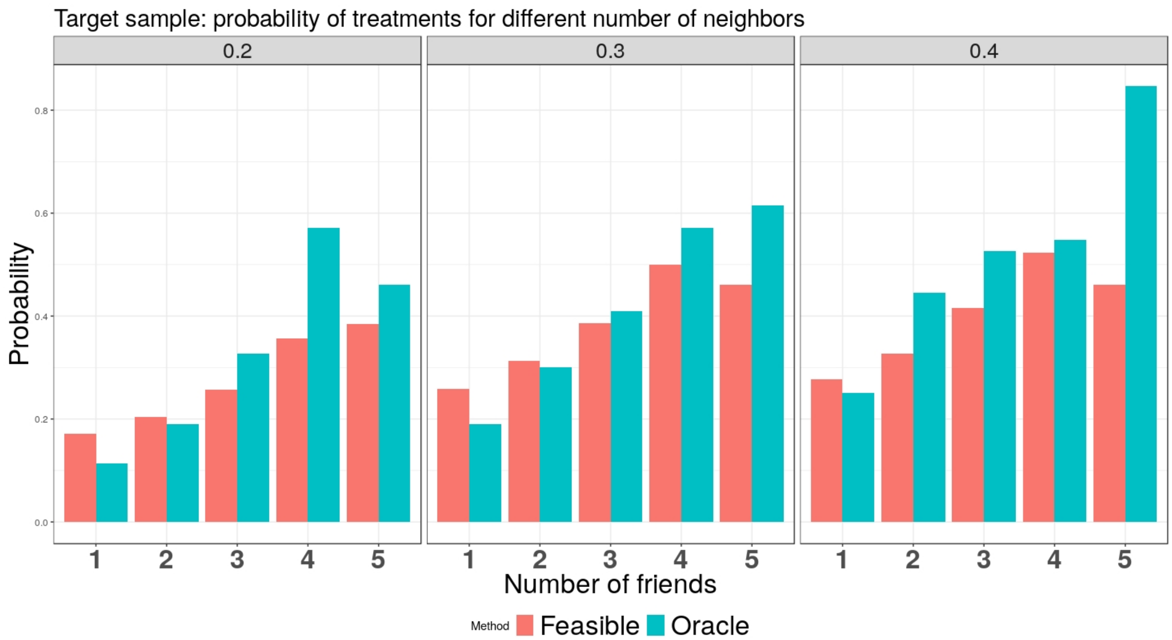

Figure 3: Target sample. The plot reports the probability of being treated as a function of the number of neighbors. In red the feasible NEWM and in blue the oracle NEWM. Different panels correspond to different capacity constraints.

Table 3: Estimated coefficients of the policy function using the feasible NEWM, rescaled by the size of the risk-adversion coefficient.

	Intercept	Age	Rice Area	Risk Adversion
Capacity 0.2	6.944	-1	2.204	1
Capacity 0.3	-0.919	-0.426	1.512	1
Capacity 0.4	0.594	-0.132	0.330	1

7 Conclusions and discussion

In this paper, we have introduced a method for estimating treatment allocation rules under network interference. We consider constrained environments, and we accommodate policy functions that do not necessarily depend on network information. The proposed methodology is valid for a generic class of network formation models under network exogeneity, and it relies on semi-parametric estimators. We cast the optimization problem into a mixed-integer linear program, and we derive guarantees on the regret of the estimated policy.

Our method assumes anonymous and exogenous interactions. Future research can address the case of endogenous interactions by explicitly modeling the endogenous component. Similarly, an avenue for future research is to replace the exogeneity of the network formation with assumptions on the (endogenous) network formation process.

We estimate welfare-maximizing policies when the network information on the target sample is not observed by directly maximizing the empirical welfare. However, it may be interesting to extend our method by incorporating partial information on the network of the target sample. In this regard, combining the high-dimensional estimator of the network on the target sample as in Alidaee et al. (2020) with the empirical welfare maximization procedure is a promising future direction.

Finally, the literature on influence maximization has often relied on structural models, while the literature on statistical treatment choice has focused on semiparametric estimation procedures. This paper opens new questions on the trade-off between structural assumptions and model-robust estimation of policy functions, and exploring this trade-off remains an open research question.

References

Akbarpour, M., S. Malladi, and A. Saberi (2018). Just a few seeds more: value of network information for diffusion. *Available at SSRN 3062830*.

Alidaee, H., E. Auerbach, and M. P. Leung (2020). Recovering network structure from aggregated relational data using penalized regression. *arXiv preprint arXiv:2001.06052*.

Ananth, A. (2021). Optimal treatment assignment rules on networked populations. *Working paper*.

Armstrong, T. and S. Shen (2015). Inference on optimal treatment assignments. *Available at SSRN 2592479*.

Aronow, P. M., C. Samii, et al. (2017). Estimating average causal effects under general interference, with application to a social network experiment. *The Annals of Applied Statistics 11*(4), 1912–1947.

Athey, S., D. Eckles, and G. W. Imbens (2018). Exact p-values for network interference. *Journal of the American Statistical Association 113*(521), 230–240.

Athey, S. and S. Wager (2021). Policy learning with observational data. *Econometrica 89*(1), 133–161.

Auerbach, E. (2019). Identification and estimation of a partially linear regression model using network data. *arXiv preprint arXiv:1903.09679*.

Banerjee, A., A. G. Chandrasekhar, E. Duflo, and M. O. Jackson (2013). The diffusion of microfinance. *Science 341*(6144), 1236498.

Banerjee, A., A. G. Chandrasekhar, E. Duflo, and M. O. Jackson (2014). Gossip: Identifying central individuals in a social network. Technical report, National Bureau of Economic Research.

Banerjee, A., A. G. Chandrasekhar, E. Duflo, and M. O. Jackson (2019). Using gossips to spread information: Theory and evidence from two randomized controlled trials. *The Review of Economic Studies 86*(6), 2453–2490.

Barrera-Osorio, F., M. Bertrand, L. L. Linden, and F. Perez-Calle (2011). Improving the design of conditional transfer programs: Evidence from a randomized education experiment in colombia. *American Economic Journal: Applied Economics 3*(2), 167–95.

Bhattacharya, D. (2009). Inferring optimal peer assignment from experimental data. *Journal of the American Statistical Association 104*(486), 486–500.

Bhattacharya, D. and P. Dupas (2012). Inferring welfare maximizing treatment assignment under budget constraints. *Journal of Econometrics 167*(1), 168–196.

Bhattacharya, D., P. Dupas, and S. Kanaya (2019). Demand and welfare analysis in discrete choice models with social interactions. *Available at SSRN 3116716*.

Bloch, F., M. O. Jackson, and P. Tebaldi (2017). Centrality measures in networks. *Available at SSRN 2749124*.

Bond, R. M., C. J. Fariss, J. J. Jones, A. D. Kramer, C. Marlow, J. E. Settle, and J. H. Fowler (2012). A 61-million-person experiment in social influence and political mobilization. *Nature 489*(7415), 295.

Breiman, L. (2001). Random forests. *Machine learning 45*(1), 5–32.

Breza, E., A. G. Chandrasekhar, T. H. McCormick, and M. Pan (2020). Using aggregated relational data to feasibly identify network structure without network data. *American Economic Review 101*(8), 2454–84.

Brooks, R. L. (1941). On colouring the nodes of a network. In *Mathematical Proceedings of the Cambridge Philosophical Society*, Volume 37, pp. 194–197. Cambridge University Press.

Cai, J., A. De Janvry, and E. Sadoulet (2015). Social networks and the decision to insure. *American Economic Journal: Applied Economics 7*(2), 81–108.

Chernozhukov, V., D. Chetverikov, M. Demirer, E. Duflo, C. Hansen, W. Newey, and J. Robins (2018). Double/debiased machine learning for treatment and structural parameters.

Chernozhukov, V., D. Chetverikov, K. Kato, et al. (2014). Gaussian approximation of suprema of empirical processes. *The Annals of Statistics 42*(4), 1564–1597.

Chiang, H. D., K. Kato, Y. Ma, and Y. Sasaki (2019). Multiway cluster robust double/debiased machine learning. *arXiv preprint arXiv:1909.03489*.

Chin, A., D. Eckles, and J. Ugander (2018). Evaluating stochastic seeding strategies in networks. *arXiv preprint arXiv:1809.09561*.

Christakis, N. A. and J. H. Fowler (2010). Social network sensors for early detection of contagious outbreaks. *PloS one 5*(9).

Devroye, L., L. Györfi, and G. Lugosi (2013). *A probabilistic theory of pattern recognition*, Volume 31. Springer Science & Business Media.

DiTraglia, F. J., C. Garcia-Jimeno, R. O'Keeffe-O'Donovan, and A. Sanchez-Becerra (2020). Identifying causal effects in experiments with social interactions and non-compliance. *arXiv preprint arXiv:2011.07051*.

Duflo, E., P. Dupas, and M. Kremer (2011). Peer effects, teacher incentives, and the impact of tracking: Evidence from a randomized evaluation in kenya. *American Economic Review 101*(5), 1739–74.

Eckles, D., H. Esfandiari, E. Mossel, and M. A. Rahimian (2019). Seeding with costly network information. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pp. 421–422.

Egger, D., J. Haushofer, E. Miguel, P. Niehaus, and M. W. Walker (2019). General equilibrium effects of cash transfers: experimental evidence from kenya. Technical report, National Bureau of Economic Research.

Elliott, G. and R. P. Lieli (2013). Predicting binary outcomes. *Journal of Econometrics 174*(1), 15–26.

Farrell, M. H. (2015). Robust inference on average treatment effects with possibly more covariates than observations. *Journal of Econometrics 189*(1), 1–23.

Florios, K. and S. Skouras (2008). Exact computation of max weighted score estimators. *Journal of Econometrics 146*(1), 86–91.

Galeotti, A., B. Golub, and S. Goyal (2020). Targeting interventions in networks. *Econometrica, Forthcoming*.

Goldsmith-Pinkham, P. and G. W. Imbens (2013). Social networks and the identification of peer effects. *Journal of Business & Economic Statistics 31*(3), 253–264.

Graham, B. S., G. W. Imbens, and G. Ridder (2010). Measuring the effects of segregation in the presence of social spillovers: A nonparametric approach. Technical report, National Bureau of Economic Research.

Hajłasz, P. and J. Malỳ (2002). Approximation in sobolev spaces of nonlinear expressions involving the gradient. *Arkiv för Matematik 40*(2), 245–274.

Hansen, L. P. (1982). Large sample properties of generalized method of moments estimators. *Econometrica: Journal of the Econometric Society*, 1029–1054.

He, X. and K. Song (2018). Measuring diffusion over a large network. *arXiv preprint arXiv:1812.04195*.

Hirano, K. and J. R. Porter (2009). Asymptotics for statistical treatment rules. *Econometrica 77*(5), 1683–1701.

Hudgens, M. G. and M. E. Halloran (2008). Toward causal inference with interference. *Journal of the American Statistical Association 103*(482), 832–842.

Imai, K., Z. Jiang, and A. Malani (2020). Causal inference with interference and noncompliance in two-stage randomized experiments. *Journal of the American Statistical Association*, 1–13.

Imbens, G. W. (2000). The role of the propensity score in estimating dose-response functions. *Biometrika 87*(3), 706–710.

Jackson, M. O. and E. Storms (2018). Behavioral communities and the atomic structure of networks. *Available at SSRN 3049748*.

Kang, H. and G. Imbens (2016). Peer encouragement designs in causal inference with partial interference and identification of local average network effects. *arXiv preprint arXiv:1609.04464*.

Kempe, D., J. Kleinberg, and É. Tardos (2003). Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 137–146. ACM.

Kim, B. (2020). Analysis of randomized experiments with network interference and noncompliance. *arXiv preprint arXiv:2012.13710*.

Kim, D. A., A. R. Hwong, D. Stafford, D. A. Hughes, A. J. O'Malley, J. H. Fowler, and N. A. Christakis (2015). Social network targeting to maximise population behaviour change: a cluster randomised controlled trial. *The Lancet 386*(9989), 145–153.

Kitagawa, T. and A. Tetenov (2018). Who should be treated? Empirical welfare maximization methods for treatment choice. *Econometrica 86*(2), 591–616.

Kitagawa, T. and A. Tetenov (2019). Equality-minded treatment choice. *Journal of Business & Economic Statistics*, 1–14.

Kitagawa, T. and G. Wang (2020). Who should get vaccinated? individualized allocation of vaccines over sir network. *arXiv preprint arXiv:2012.04055*.

Laber, E. B., N. J. Meyer, B. J. Reich, K. Pacifici, J. A. Collazo, and J. M. Drake (2018). Optimal treatment allocations in space and time for on-line control of an emerging infectious disease. *Journal of the Royal Statistical Society: Series C (Applied Statistics) 67*(4), 743–789.

Ledoux, M. and M. Talagrand (2011). Probability in banach spaces. classics in mathematics.

Leung, M. P. (2020). Treatment and spillover effects under network interference. *Review of Economics and Statistics 102*(2), 368–380.

Li, X., P. Ding, Q. Lin, D. Yang, and J. S. Liu (2019). Randomization inference for peer effects. *Journal of the American Statistical Association*, 1–31.

Liu, L., M. G. Hudgens, B. Saul, J. D. Clemens, M. Ali, and M. E. Emch (2019). Doubly robust estimation in observational studies with partial interference. *Stat 8*(1), e214.

Manresa, E. (2013). Estimating the structure of social interactions using panel data. *Unpublished Manuscript. CEMFI, Madrid*.

Manski (2004). Statistical treatment rules for heterogeneous populations. *Econometrica 72*(4), 1221–1246.

Manski, C. F. (2013). Identification of treatment response with social interactions. *The Econometrics Journal 16*(1), S1–S23.

Mbakop, E. and M. Tabord-Meehan (2016). Model selection for treatment choice: Penalized welfare maximization. *arXiv preprint arXiv:1609.03167*.

Muralidharan, K., P. Niehaus, and S. Sukhtankar (2017). General equilibrium effects of (improving) public employment programs: Experimental evidence from india. Technical report, National Bureau of Economic Research.

Naimi, A. I., S. R. Cole, and E. H. Kennedy (2017). An introduction to g methods. *International journal of epidemiology 46*(2), 756–762.

Opper, I. M. (2016). Does helping john help sue? evidence of spillovers in education. *American Economic Review 109*(3), 1080–1115.

Paulin, D. (2012). Concentration inequalities in locally dependent spaces. *arXiv preprint arXiv:1212.2013*.

Pollard, D. (1990). Empirical processes: theory and applications. In *NSF-CBMS regional conference series in probability and statistics*, pp. i–86. JSTOR.

Robins, J. M., A. Rotnitzky, and L. P. Zhao (1994). Estimation of regression coefficients when some regressors are not always observed. *Journal of the American statistical Association 89*(427), 846–866.

Sävje, F., P. M. Aronow, and M. G. Hudgens (2020). Average treatment effects in the presence of unknown interference. *Annals of Statistics, Forthcoming*.

Sinclair, B., M. McConnell, and D. P. Green (2012). Detecting spillover effects: Design and analysis of multilevel experiments. *American Journal of Political Science 56*(4), 1055–1069.

Sobel, M. E. (2006). What do randomized studies of housing mobility demonstrate? causal inference in the face of interference. *Journal of the American Statistical Association 101*(476), 1398–1407.

Stoye, J. (2009). Minimax regret treatment choice with finite samples. *Journal of Econometrics 151*(1), 70–81.

Stoye, J. (2012). Minimax regret treatment choice with covariates or with limited validity of experiments. *Journal of Econometrics 166*(1), 138–156.

Su, L., W. Lu, and R. Song (2019). Modelling and estimation for optimal treatment decision with interference. *Stat 8*(1), e219.

Tetenov, A. (2012). Statistical treatment choice based on asymmetric minimax regret criteria. *Journal of Econometrics 166*(1), 157–165.

Van Der Vaart, A. W. and J. A. Wellner (1996). Weak convergence. In *Weak convergence and empirical processes*, pp. 16–28. Springer.

Vazquez-Bare, G. (2020). Causal spillover effects using instrumental variables. *arXiv preprint arXiv:2003.06023*.

Viviano, D. (2020). Policy choice in experiments with unknown interference. *arXiv preprint arXiv:2011.08174*.

Wager, S. and K. Xu (2020). Experimenting in equilibrium. *Management Science, Forthcoming*.

Wainwright, M. J. (2019). *High-dimensional statistics: A non-asymptotic viewpoint*, Volume 48. Cambridge University Press.

Wenocur, R. S. and R. M. Dudley (1981). Some special vapnik-chervonenkis classes. *Discrete Mathematics 33*(3), 313–318.

Zhou, Z., S. Athey, and S. Wager (2018). Offline multi-action policy learning: Generalization and optimization. *arXiv preprint arXiv:1810.04778*.

Zubcsek, P. P. and M. Sarvary (2011). Advertising to a social network. *Quantitative Marketing and Economics 9*(1), 71–107.

Supplementary Materials

Appendix A A numerical study

In this section, we study the numerical performance of the proposed methodology in a small simulation study. We simulate data according to the following data generating process (DGP):

$$\begin{aligned} Y_i &= |N_i|^{-1}\Big(X_i\beta_1 + X_i\beta_2 D_i + \mu\Big)\sum_{k\in N_i} D_k + X_i\beta_3 D_i + \varepsilon_i \\ \varepsilon_i &= \eta_i/\sqrt{2} + \sum_{k\in N_i}\eta_k/\sqrt{2|N_i|}, \quad \eta_i \sim_{i.i.d.} \mathcal{N}(0,1), \end{aligned} \tag{A.1}$$

where $|N_i|$ is set to be one for the elements with no neighbors. We simulate covariates as $X_{i,u} \sim_{i.i.d.} \mathcal{U}(-1,1)$ for $u \in \{1,2,3,4\}$. We draw $\beta_1, \beta_2, \mu \in \{-1,1\}^3$ independently and with equal probabilities. We draw $\beta_3 \in \{-1.5, 1.5\}$ with equal probabilities. We evaluate the NEWM method with and without balancing score adjustment, under correct specification of the conditional mean function. We impose trimming at 1% on the estimation of the balancing score. We compare the performance of the proposed methodology to two competing methods that ignore network effects. These are the empirical welfare maximization procedures with and without regression adjustment (Kitagawa and Tetenov, 2018; Athey and Wager, 2021).[42] For any of the method under consideration, we consider a policy function of the form

$$\pi(X_i) = 1\Big\{X_{i,1}\phi_1 + X_{i,2}\phi_2 + \phi_3 \geq 0\Big\}. \tag{A.2}$$

Optimization is performed using the MILP formulation.

In the first set of simulations, we consider a geometric network formation of the form

$$A_{i,j} = 1\Big\{|X_{i,2} - X_{j,2}|/2 + |X_{i,4} - X_{j,4}|/2 \leq r_n\Big\} \tag{A.3}$$

where $r_n = \sqrt{4/2.75n}$ similarly to simulations in Leung (2020). In the second set of simulations, we generate Barabasi-Albert networks, where we first draw $n/5$ edges

[42]Namely, first, we consider the empirical welfare maximization method that does not account for network interference discussed in Kitagawa and Tetenov (2018) with a known balancing score of individual treatment. Second, we also compare the double robust formulation of the EWM method discussed in Athey and Wager (2021), with a linear model specification without network information.

uniformly according to Erdős-Rényi graph with probabilities $p = 10/n$, and second, we draw sequentially connections of the new nodes to the existing ones with probability equal to the number of connection of each pre-existing node divided by the overall number of connections in the graph. We estimate the methods over 200 data sets, and we evaluate the performance of each estimate over 1000 networks, drawn from the same DGP. Results are collected in Table 4. The table reports the welfare for different sample sizes. Observe that by construction, the sample size also corresponds to different data-generating processes of the network formation model, which, in turn, may shift by a constant the value of the regret.[43] The table shows that our procedure uniformly outperforms methods that ignore network effects.

Table 4: Out-of-sample median *welfare* over 200 replications. DR is the method in Athey and Wager (2021) with estimated balancing score and EWM PS is the method in Kitagawa and Tetenov (2018) with known balancing score. NEWM1 is the proposed method with a correctly specified outcome model. NEWM2 is the double robust NEWM. G denotes the geometric network, and AB the Albert-Barabasi network. Different sample sizes correspond to different data-generating processes.

Welfare	$n = 50$		$n = 70$		$n = 100$		$n = 150$		$n = 200$	
	G	AB	G	AB	G	AB	G	AB	G	AB
DR	1.51	0.94	1.50	1.08	1.42	1.05	1.53	0.95	1.41	0.95
EWM PS	1.21	0.93	1.23	0.92	1.32	0.93	1.38	0.90	1.29	0.95
NEWM1	1.74	1.31	1.87	1.38	1.93	1.37	1.91	1.40	2.00	1.39
NEWM2	1.78	1.22	1.89	1.33	1.89	1.37	1.94	1.28	1.95	1.33

Appendix B Proofs: Notation and Definitions

Before discussing the main results, we need to introduce the necessary notation. We define $S_i(\pi) = \sum_{k \in N_i} \pi(X_k)$. We denote $\mathcal{A}_n$ the space of symmetric matrices in $\mathbb{R}^{n \times n}$ with entries being either zero or ones.

Definition B.1 (Proper Cover). *Given an adjacency matrix $A_n \in \mathcal{A}_n$, with n rows and columns, a family $\mathcal{C}_n = \{\mathcal{C}_n(j)\}$ of disjoint subsets of $[n]$ is a proper cover of A_n*

[43] Namely $\sup_{\pi \in \Pi} W(\pi)$ may differ across sample sizes because sample sizes correspond to different DGPs.

if $\cup\mathcal{C}_n(j) = [n]$ *and* $\mathcal{C}_n(j)$ *contains units such that for any pair of elements* $\{(i,k) \in \mathcal{C}_n(j), k \neq i\}$*,* $A_n^{(i,k)} = 0$.

The size of the smallest proper cover is the chromatic number, defined as $\chi(A_n)$.

Definition B.2 (Chromatic number)**.** *The chromatic number* $\chi(A_n)$*, denotes the size of the smallest proper cover of* A_n*.*

Definition B.3. *For a given matrix* $A \in \mathcal{A}$*, we define* $A_n^2(A) \in \mathcal{A}_n$ *the adjacency matrix where each row corresponds to a unit* $i \in \{1, ..., n\}$ *and where two of such units are connected if they are neighbor or they share one first or second degree neighbor. Similarly* $A_n^M(A)$ *is the adjacency matrix obtained after connecting such units sharing common neighbors up to* Mth *degree. Here* $N_{i,M}$ *defines the set of neighbors of individual* $i \in \{1, \cdots, n\}$ *with adjacency matrix* A_n^M*.*

The proper cover of A_n^2 is defined as $\mathcal{C}_n^2 = \{\mathcal{C}_n^2(j)\}$ with chromatic number $\chi(A_n^2)$. Similarly $\mathcal{C}_n^M = \{\mathcal{C}_n^M(j)\}$ with chromatic number $\chi(A_n^M)$ is the proper cover of A_n^M.

Next, we discuss definitions on covering numbers.[44] For $z_1^n = (z_1, ..., z_n)$ be arbitrary points in $\mathcal{Z}$, for a function class $\mathcal{F}$, with $f \in \mathcal{F}$, $f : \mathcal{Z} \mapsto \mathbb{R}$, we define,

$$\mathcal{F}(z_1^n) = \{f(z_1), ..., f(z_n) : f \in \mathcal{F}\}. \tag{B.1}$$

Definition B.4. For a class of functions $\mathcal{F}$, with $f : \mathcal{Z} \mapsto \mathbb{R}$, $\forall f \in \mathcal{F}$ and n data points $z_1, ..., z_n \in \mathcal{Z}$ define the l_q-covering number $\mathcal{N}_q\Big(\varepsilon, \mathcal{F}(z_1^n)\Big)$ to be the cardinality of the smallest cover $\mathcal{S} := \{s_1, ..., s_N\}$, with $s_j \in \mathbb{R}^n$, such that for each $f \in \mathcal{F}$, there exist an $s_j \in \mathcal{S}$ such that $(\frac{1}{n}\sum_{i=1}^n |f(z_i) - s_j^{(i)}|^q)^{1/q} < \varepsilon$.

Throughout our analysis, for a random variable $X = (X_1, ..., X_n)$ we denote $\mathbb{E}_X[.]$ the expectation with respect to X. The Rademacher complexity is defined as follows.

Definition B.5. Let $X_1, ..., X_n$ be arbitrary random variables. Let $\sigma = \{\sigma_i\}_{i=1}^n$ be $i.i.d$ Rademacher random variables (i.e., $P(\sigma_i = -1) = P(\sigma_i = 1) = 1/2$), independent of $X_1, ..., X_n$. We define the empirical Rademacher complexity as

$$\mathcal{R}_n(\mathcal{F}) = \mathbb{E}_\sigma\Big[\sup_{f\in\mathcal{F}}|\frac{1}{n}\sum_{i=1}^n \sigma_i f(X_i)|\Big| X_1, ..., X_n\Big]. \tag{B.2}$$

[44] Here we adopt similar notation to Chapter 28 and Chapter 29 of Devroye et al. (2013).

Appendix C Main Lemmas

Lemma C.1. *The following holds:* $\chi(A_n) \leq \chi(A_n^M) \leq M\mathcal{N}_{n,M+1}^{M+1}$.

Proof of Lemma C.1. The first inequality follows by Definition B.3. The second inequality follows by Brook's Theorem (Brooks, 1941), since two individuals $(i,j) \in \{1,\cdots,n\}^2$ are connected if they share a common neighbor up to the Mth degree. The maximum degree is bounded by $\mathcal{N}_{n,1}+\mathcal{N}_{n,1}\times\mathcal{N}_{n,2}+\cdots+\prod_{s=1}^{M+1}\mathcal{N}_{n,s} \leq M\mathcal{N}_{n,M+1}^{M+1}$. $\square$

Lemma C.2. *Let $\mathcal{F}_1,\cdots,\mathcal{F}_k$ be classes of bounded functions with VC-dimension $v<\infty$ and envelope $\bar{F}<\infty$, for $k\geq 2$. Let*

$$\mathcal{J} = \Big\{f_1(f_2+...+f_k), \quad f_j \in \mathcal{F}_j, \quad j=1,...,k\Big\}, \quad \mathcal{J}_n(z_1^n) = \Big\{h(z_1),...,h(z_n); h \in \mathcal{J}\Big\}.$$

For arbitrary fixed points $z_1^n \in \mathbb{R}^d$, $\int_0^{2\bar{F}}\sqrt{\log\Big(\mathcal{N}_1\Big(u,\mathcal{J}(z_1^n)\Big)\Big)}du < c_{\bar{F}}\sqrt{k\log(k)v}$. for a constant $c_{\bar{F}}<\infty$ that only depend on $\bar{F}$.

Proof of Lemma C.2. Let $\mathcal{F}_{-1}(z_1^n) = \{f_2(z_1^n)+...+f_k(z_1^n), f_j \in \mathcal{F}_j, j=2,...,k\}$. By Theorem 29.6 in Devroye et al. (2013),

$$\mathcal{N}_1\Big(\varepsilon,\mathcal{F}_{-1}(z_1^n)\Big) \leq \prod_{j=2}^k \mathcal{N}_1\Big(\varepsilon/(k-1),\mathcal{F}_j(z_1^n)\Big).$$

By Theorem 29.7 in Devroye et al. (2013),

$$\mathcal{N}_1\Big(\varepsilon,\mathcal{J}_n(z_1^n)\Big) \leq \prod_{j=2}^k \mathcal{N}_1\Big(\frac{\varepsilon}{2(k-1)\bar{F}},\mathcal{F}_j(z_1^n)\Big)\mathcal{N}_1\Big(\frac{\varepsilon}{2\bar{F}},\mathcal{F}_1(z_1^n)\Big). \tag{C.1}$$

By standard properties of covering numbers, for a generic set $\mathcal{H}$, $\mathcal{N}_1(\varepsilon,\mathcal{H}) \leq \mathcal{N}_2(\varepsilon,\mathcal{H})$. Therefore

$$(\text{C.1}) \leq \prod_{j=2}^k \mathcal{N}_2\Big(\frac{\varepsilon}{2(k-1)\bar{F}},\mathcal{F}_j(z_1^n)\Big)\mathcal{N}_2\Big(\frac{\varepsilon}{2\bar{F}},\mathcal{F}_1(z_1^n)\Big).$$

We apply now a uniform entropy bound for the covering number. By Theorem 2.6.7 of Van Der Vaart and Wellner (1996), we have that for a universal constant $C<\infty$,

$$\mathcal{N}_2\Big(\frac{\varepsilon}{2(k-1)\bar{F}},\mathcal{F}_j(z_1^n)\Big) \leq C(v+1)(16e)^{(v+1)}\Big(\frac{2\bar{F}^2(k-1)}{\varepsilon}\Big)^{2v}$$

which implies that

$$\log\Big(\mathcal{N}_1\Big(\varepsilon, \mathcal{J}_n(z_1^n)\Big)\Big) \leq \sum_{j=1}^{k-1}\log\Big(\mathcal{N}_2\Big(\frac{\varepsilon}{2\bar{F}(k-1)}, \mathcal{F}_j(z_1^n)\Big)\Big) + \log\Big(\mathcal{N}_2\Big(\frac{\varepsilon}{2\bar{F}}, \mathcal{F}_1(z_1^n)\Big)\Big)$$
$$\leq k\log\Big(C(v+1)(16e)^{v+1}\Big) + k2v\log(2C\bar{F}^2(k-1)/\varepsilon).$$

Since $\int_0^{2\bar{F}}\sqrt{k\log\Big(C(v+1)(16e)^{v+1}\Big) + k2v\log(2C\bar{F}^2(k-1)/\varepsilon)}d\varepsilon \leq c_{\bar{F}}\sqrt{k\log(k)v}$ for a constant $c_{\bar{F}} < \infty$, the proof completes. □

The following lemma controls the Rademacher complexity with individualized treatments.

Lemma C.3. *Let Π be a function class with $\pi : \mathbb{R}^d \mapsto \{0,1\}$ for any $\pi \in \Pi$, with finite VC-dimension, denoted as $\mathrm{VC}(\Pi)$. Take arbitrary random variables $(X_i^{(1)}, \cdots, X_i^{(h_i)}, V_i), X_i, V_i \in \mathbb{R}^d$, functions $f_i : \mathbb{Z} \mapsto \mathbb{R}$, with f_i L-lipschitz for all $i \in \{1, \cdots, n\}$. Let $\bar{h} = \max_i h_i$. Let $\sigma_1, \cdots, \sigma_n$ be independent Rademacher random variables (independent of $(X_i, V_i)_{i=1}^n$). Then for a constant $C < \infty$ independent of L, $\bar{h}$ and n*

$$\mathbb{E}_\sigma\Big[\sup_{\pi\in\Pi}\Big|\frac{1}{n}\sum_{i=1}^n f_i\Big(\sum_{k\in\{1,\dots,h_i\}}\pi(X_i^{(k)})\Big)\pi(V_i)\sigma_i\Big|\Big] \leq CL\sqrt{\frac{(\bar{h}+1)\log(\bar{h}+1)\mathrm{VC}(\Pi)}{n}}. \quad \text{(C.2)}$$

Proof of Lemma C.7. We first discuss the case where $h_i = h$ for all i. We then turn to the case where this condition fails at the end of the proof.

Lipschitz properties First, we add and subtract the value of the function $f_i(0)$ at zero. Namely,

$$\begin{aligned}
&\mathbb{E}_\sigma\Big[\sup_{\pi\in\Pi}\Big|\frac{1}{n}\sum_{i=1}^n \sigma_i\Big(f_i\Big(\sum_{k\in\{1,\dots,h\}}\pi(X_i^{(k)})\Big) - f_i(0) + f_i(0)\Big)\pi(V_i)\Big|\Big] \\
&\leq \underbrace{\mathbb{E}_\sigma\Big[\sup_{\pi\in\Pi}\Big|\frac{1}{n}\sum_{i=1}^n \sigma_i\Big(f_i\Big(\sum_{k\in\{1,\dots,h\}}\pi(X_i^{(k)})\Big) - f_i(0)\Big)\pi(V_i)\Big|\Big]}_{(1)} \\
&+ \underbrace{\mathbb{E}_\sigma\Big[\sup_{\pi\in\Pi}\Big|\frac{1}{n}\sum_{i=1}^n \sigma_i f_i(0)\pi(V_i)\Big|\Big]}_{(2)}
\end{aligned} \quad \text{(C.3)}$$

where the last inequality follows by the triangular inequality. We bound first (1). We decompose the supremum over $\pi \in \Pi$ as follows.

$$\mathbb{E}_\sigma\Big[\sup_{\pi\in\Pi}\Big|\frac{1}{n}\sum_{i=1}^n \sigma_i\Big(f_i\Big(\sum_{k\in\{1,\dots,h\}}\pi(X_i^{(k)})\Big) - f_i(0)\Big)\pi(V_i)\Big|\Big]$$
$$\leq \mathbb{E}_\sigma\Big[\sup_{\pi_1\in\Pi,\pi_2\in\Pi}\Big|\frac{1}{n}\sum_{i=1}^n \sigma_i\Big(f_i\Big(\sum_{k\in\{1,\dots,h\}}\pi_2(X_i^{(k)})\Big) - f_i(0)\Big)\pi_1(V_i)\Big|\Big].$$

Let $\phi_i(s) = \Big(f_i(s) - f_i(0)\Big)$. Conditional on the data, ϕ_i is not random. It is Lipschitz and $\phi_i(0) = 0$ by assumption and by the fact that Y_i is uniformly bounded. By Theorem 4.12 in Ledoux and Talagrand (2011) (see also Lemma C.5)

$$\begin{aligned}&\mathbb{E}_\sigma\Big[\sup_{\pi_1\in\Pi,\pi_2\in\Pi}\Big|\frac{1}{n}\sum_{i=1}^n \sigma_i\Big(f_i\Big(\sum_{k\in\{1,\dots,h\}}\pi_2(X_i^{(k)})\Big) - f_i(0)\Big)\pi_1(V_i)\Big|\Big]\\&\leq L2\mathbb{E}_\sigma\Big[\sup_{\pi_1\in\Pi,\pi_2\in\Pi}\Big|\frac{1}{n}\sum_{i=1}^n \sigma_i\Big(\sum_{k\in\{1,\dots,h\}}\pi_2(X_i^{(k)})\Big)\pi_1(V_i)\Big|\Big].\end{aligned} \tag{C.4}$$

Function class decomposition We decompose the supremum as follows:

$$L\mathbb{E}_\sigma\Big[\sup_{\pi_1\in\Pi,\pi_2\in\Pi}\Big|\frac{1}{n}\sum_{i=1}^n \sigma_i\Big(\sum_{k\in\{1,\dots,h\}}\pi_2(X_i^{(k)})\Big)\pi_1(V_i)\Big|\Big]$$
$$\leq \underbrace{L\,\mathbb{E}_\sigma\Big[\sup_{\pi_1\in\Pi,\pi_2\in\Pi,\cdots,\pi_{h+1}\in\Pi}\Big|\frac{1}{n}\sum_{i=1}^n \sigma_i\Big(\sum_{k\in\{1,\dots,h\}}\pi_{k+1}(X_i^{(k)})\Big)\pi_1(V_i)\Big|\Big]}_{(J)}.$$

We re-parametrize the class of functions as follows.

$$(J) = \mathbb{E}_\sigma\Big[\sup_{\pi_1\in\Pi,\tilde\pi_2\in\Pi_2,\cdots,\tilde\pi_{h+1}\in\Pi_{h+1}}\Big|\frac{1}{n}\sum_{i=1}^n \sigma_i\Big(\sum_{k\in\{1,\dots,h\}}\tilde\pi_{k+1}(X_i^{(1)},\cdots,X_i^{(h)})\Big)\pi_1(V_i)\Big|\Big],$$

where $\Pi_2, ..., \Pi_{h+1}$, are such that $\tilde\pi_j \in \Pi_j : \tilde\pi_j(X_i^{(1)},\cdots,X_i^{(k)},\cdots,X_i^{(h)}) = \pi(X_i^{(j)})$. By Theorem 29.4 in Devroye et al. (2013), $\mathrm{VC}(\Pi_j) = \mathrm{VC}(\Pi)$ for all $j \in \{1, ..., h\}$. Let

$$\tilde\Pi = \Big\{\pi_1\Big(\sum_{j=1}^h \pi_{j+1}\Big), \pi_k \in \Pi_k, k = 1,\cdots,h+1\Big\}.$$

For any fixed point data point z_1^n, by Lemma C.2, the Dudley's integral of the function class $\tilde\Pi(z_1^n)$ is bounded by $C\sqrt{(h+1)\log(h+1)\mathrm{VC}(\Pi)}$, for a finite constant C.

Therefore, by C.2 and the Dudley's entropy integral bound (Lemma F.8) we can bound the conditional Rademacher complexity

$$\mathbb{E}_\sigma\Big[\sup_{\pi_1\in\Pi_1,\pi_2\in\Pi_2,\dots,\pi_{h+1}\in\Pi_{h+1}}\Big|\frac{1}{n}\sum_{i=1}^n\sigma_i\Big(\sum_{k\in\{1,\dots,h\}}\pi_{k+1}(X_i^{(k)})\Big)\pi_1(V_i)\Big|\Big]$$
$$\leq C\sqrt{(h+1)\frac{\mathrm{VC}(\Pi)\log(h+1)}{n}}$$

for a constant $C<\infty$.

Remaining components For (2) in Equation (C.3) we can use the same argument used for (1). In particular, since Π has finite VC-dimension, and it contains functions mapping to $\{0,1\}$, by Theorem 2.6.7 of Van Der Vaart and Wellner (1996) [45], $\int_0^2\sqrt{\mathcal{N}_2(u,\Pi(z_1^n))}du < C\sqrt{\mathrm{VC}(\Pi)}$ for a constant C.

Conclusion We are now left to discuss the case where h_i is different for each individual. In this case, we can assign to each unit covariates $\{X_{k\in N_i},\emptyset,\emptyset,\cdots,\emptyset\}$ with in total $\bar{h}$ and parametrizing $\pi(\emptyset)=0$. The rest of the proof follows similarly as before. □

Lemma C.4. *$e^c(.)\in(\delta,1-\delta)$ for $\delta\in(0,1)$. Then $\frac{1\{N=\sum_{k\in N_i}D_k,d=D_i\}}{e^c\Big(d,N,Z_{k\in N_i},Z_i,|N_i|\Big)}$ is $2/\delta$-Lipschitz in its second argument for all $d\in\{0,1\}$*

Proof of Lemma C.4. Let $O_i=(Z_{k\in N_i},Z_i,|N_i|)$. Then for any $N,N'\in\mathbb{Z}$

$$\Big|\frac{1\{N=\sum_{k\in N_i}D_k,d=D_i\}}{e^c\Big(d,N,O_i\Big)}-\frac{1\{N'=\sum_{k\in N_i}D_k,d=D_i\}}{e^c\Big(d,N',O_i\Big)}\Big|\leq\frac{2}{\delta}.$$

for $N\neq N'$. The last inequality follows from the fact that by the overlap condition and the triangular inequality. Since N is discrete, the right-hand side is at least $2/\delta|N-N'|$ which completes the proof. □

C.1 Auxiliary lemmas for an unbounded outcome

The next lemma provides a bound on the Rademacher complexity in the presence of the composition of functions. We discuss the Ledoux-Talagrand contraction inequality (Ledoux and Talagrand, 2011; Chernozhukov et al., 2014) to the case of interest in this paper.

[45] The argument is the same used in the proof of Lemma C.2.

Lemma C.5. *Let $\phi_i : \mathbb{R} \mapsto \mathbb{R}$ be Lipschitz functions with parameter L, $\forall i \in \{1, ..., n\}$, i.e., $|\phi_i(a) - \phi_i(b)| \leq L|a - b|$ for all $a, b \in \mathbb{R}$, with $\phi_i(0) = 0$. Then, for any $\mathcal{T} \subseteq \mathbb{R}^n$, with $t = (t_1, ..., t_n) \in \mathbb{R}$, $\alpha = (\alpha_1, ..., \alpha_n) \in \mathcal{A} \subseteq \{0, 1\}^n$,*

$$\frac{1}{2}\mathbb{E}_\sigma\Big[sup_{t\in\mathcal{T},\alpha\in\mathcal{A}}\Big|\frac{1}{n}\sum_{i=1}^n \sigma_i\phi_i(t_i)\alpha_i\Big|\Big] \leq L\mathbb{E}_\sigma\Big[sup_{t\in\mathcal{T},\alpha\in\mathcal{A}}\Big|\frac{1}{n}\sum_{i=1}^n \alpha_i\sigma_i t_i\Big|\Big].$$

Proof of Lemma C.5. The proof follows closely the one in Theorem 4.12 of Ledoux and Talagrand (2011) while dealing with the additional α vector. First note that if $\mathcal{T}$ is unbounded, there will be some setting so that the right hand side is infinity and the result trivially holds. Therefore, we can focus to the case where $\mathcal{T}$ is bounded. First we aim to show that for $\mathcal{T} \subseteq \mathbb{R}^2$, for $\alpha \in \{0, 1\}^2$,

$$\mathbb{E}\Big[\sup_{t\in\mathcal{T},\alpha\in\mathcal{A}}\alpha_1 t_1 + \sigma_2\phi(t_2)\alpha_2\Big] \leq \mathbb{E}\Big[\sup_{t\in\mathcal{T},\alpha\in\mathcal{A}}\alpha_1 t_1 + L\sigma_2 t_2\alpha_2\Big]. \tag{C.5}$$

If the claim above is true, than it follows that

$$\mathbb{E}\Big[\sup_{t\in\mathcal{T},\alpha\in\mathcal{A}}\alpha_1\phi_1(t_1)\sigma_1 + \sigma_2\phi(t_2)\alpha_2|\sigma_1\Big] \leq \mathbb{E}\Big[\sup_{t\in\mathcal{T},\alpha\in\mathcal{A}}\alpha_1\phi_1(t_1)\sigma_1 + L\sigma_2 t_2\alpha_2\Big|\sigma_1\Big].$$

as $\sigma_1\phi(t_1)$ simply transforms $\mathcal{T}$ (and it is still bounded, if not the claim would trivially holds), and we can iteretively apply this result. Hence, we first prove Equation (C.5). Define for $a, b \in \mathcal{A}$, $I(t, s, a, b) := \frac{1}{2}\Big(t_1 a_1 + a_2\phi(t_2)\Big) + \frac{1}{2}\Big(s_1 b_1 - b_2\phi(s_2)\Big)$. We want to show that the right hand side in Equation (C.5) is larger thant $I(t, s, a, b)$ for all $t, s \in \mathcal{T}$ and $a, b \in \mathcal{A}$. Since we are taking the supremum over t, s, a, b, we can assume without loss of generality that

$$t_1 a_1 + a_2\phi(t_2) \geq s_1 b_1 + b_2\phi(s_2), \quad s_1 b_1 - b_2\phi(s_2) \geq t_1 a_1 - a_2\phi(t_2). \tag{C.6}$$

We can now define four quantities of interest, being

$$m = b_1 s_1 - b_2\phi(s_2), \quad n = b_1 s_1 - L s_2 b_2, \quad m' = a_1 t_1 + L a_2 t_2, \quad n' = a_1 t_1 + a_2\phi(t_2).$$

We would like to show that $2I(t, s, a, b) = m + n' \leq m' + n$. We consider four different cases, similarly to the proof of Ledoux and Talagrand (2011) and argue that for any value of $(a_1, a_2, b_1, b_2) \in \{0, 1\}^4$ the claim holds.

Case 1 Start from the case $a_2 t_2, s_2 b_2 \geq 0$. We know that $\phi(0) = 0$, so that $|b_2\phi(s_2)| \leq L b_2 s_2$. Now assume that $a_2 t_2 \geq b_2 s_2$. In this case

$$m - n = L b_2 s_2 - b_2\phi(s_2) \leq L a_2 t_2 - a_2\phi(t_2) = m' - n' \tag{C.7}$$

since $|a_2\phi(t_2) - b_2\phi(s_2)| \leq L|a_2t_2 - b_2s_2| = L(a_2t_2 - b_2s_2)$. To see why this last claim holds, note that for $a_2, b_2 = 1$, then the results hold by the condition $a_2t_2 \geq b_2s_2$ and Lipschitz continuity. If instead $a_2 = 1, b_2 = 0$, the claim trivially holds. While the case $a_2 = 0, b_2 = 1$, then it must be that $s_2 = 0$ since we assumed that $a_2t_2 \geq 0, b_2s_2 \geq 0$ and $a_2t_2 \geq b_2s_2$. Thus $m - n \leq m' - n'$. If instead $b_2s_2 \geq a_2t_2$, then use $-\phi$ instead of ϕ and switch the roles of s, t giving a similar proof.

Case 2 Let $a_2t_2 \leq 0, b_2s_2 \leq 0$. Then the proof is the same as Case 1, switching the signs where necessary.

Case 3 Let $a_2t_2 \geq 0, b_2s_2 \leq 0$. Then we have $a_2\phi(t_2) \leq La_2t_2$, since $a_2 \in \{0, 1\}$ and by Lipschitz properties of ϕ, $-b_2\phi(s_2) \leq -b_2Ls_2$ so that $a_2\phi(t_2) - b_2\phi(s_2) \leq a_2Lt_2 - b_2Ls_2$ proving the claim.

Case 4 Let $a_2t_2 \leq 0, b_2s_2 \geq 0$. Then the claim follows simmetrically to Case 3.

We now conclude the proof. Denote $[x]_+ = \max\{0, x\}$ and $[x]_- = \max\{-x, 0\}$. Then we have

$$\mathbb{E}\Big[\frac{1}{2}\sup_{t \in \mathcal{T}, \alpha \in \mathcal{A}}\Big|\sum_{i=1}^n \sigma_i \phi_i(t_i)\alpha_i\Big|\Big] \leq \mathbb{E}\Big[\frac{1}{2}\sup_{t \in \mathcal{T}, \alpha \in \mathcal{A}}\Big(\sum_{i=1}^n \sigma_i \phi_i(t_i)\alpha_i\Big)_+\Big] + \mathbb{E}\Big[\frac{1}{2}\sup_{t \in \mathcal{T}, \alpha \in \mathcal{A}}\Big(\sum_{i=1}^n \sigma_i \phi_i(t_i)\alpha_i\Big)_-\Big]$$
$$\leq \mathbb{E}\Big[\sup_{t \in \mathcal{T}, \alpha \in \mathcal{A}}\Big(\sum_{i=1}^n \sigma_i \phi_i(t_i)\alpha_i\Big)_+\Big]$$

where the last inequality follows by symmetry of σ_i and the fact that $(-x)_- = (x)_+$. note that $sup_x(x)_+ = (sup_x x)_+$. Therefore, using Equation (C.5)

$$\mathbb{E}\Big[\sup_{t \in \mathcal{T}, \alpha \in \mathcal{A}}\Big(\sum_{i=1}^n \sigma_i \phi_i(t_i)\alpha_i\Big)_+\Big] = \mathbb{E}\Big[\Big(\sup_{t \in \mathcal{T}, \alpha \in \mathcal{A}}\sum_{i=1}^n \sigma_i \phi_i(t_i)\alpha_i\Big)_+\Big] \leq \mathbb{E}\Big[\Big(\sup_{t \in \mathcal{T}, \alpha \in \mathcal{A}}\sum_{i=1}^n L\sigma_i t_i\alpha_i\Big)_+\Big]$$
$$\leq \mathbb{E}\Big[\Big|\sup_{t \in \mathcal{T}, \alpha \in \mathcal{A}}\sum_{i=1}^n \sigma_i \phi_i(t_i)\alpha_i\Big|\Big] \leq \mathbb{E}\Big[\sup_{t \in \mathcal{T}, \alpha \in \mathcal{A}}\Big|\sum_{i=1}^n \sigma_i \phi_i(t_i)\alpha_i\Big|\Big]$$

which completes the proof. □

Lemma C.6. *Let Z be an arbitrary random variable and $\mathcal{F}$ a class of uniformly bounded functions, i.e., there exist $\bar{F} < \infty$, such that $||f||_\infty \leq \bar{F}$ for all $f \in \mathcal{F}$. Let $Y_i \sim \mathcal{P}_i|Z$, where $Y \geq 0$ is a scalar. Let $(Y_i)_{i=1}^n|Z$ be pairwise independent across individuals i. Assume that for some $u > 0$, $\mathbb{E}[Y_i^{2+u}|Z] < B, \quad \forall i \in \{1, ..., n\}$. In addition assume that for any fixed points z_1^n, for some $V < \infty$, $\int_0^{2\bar{F}} \sqrt{\log\Big(\mathcal{N}_1\Big(u, \mathcal{F}(z_1^n)\Big)\Big)}du < \sqrt{V}$. Let σ_i be i.i.d Rademacher random variables independent of Y, Z. Then there* *exist a constant* $0 < C_{\bar{F}} < \infty$ *that only depend on* $\bar{F}$ *and* u*, such that*

$$\int_0^\infty \mathbb{E}\Big[sup_{f\in\mathcal{F}}\Big|\frac{1}{n}\sigma_i\sum_{i=1}^n f(Z_i)1\{Y_i > y\}\Big|Z\Big]dy \leq C_{\bar{F}}\sqrt{\frac{BV}{n}}$$

for all $n \geq 1$.

Proof of Lemma C.6. The proof follows closely the one in Kitagawa and Tetenov (2019) (Lemma A.5) to express the bound as a function of the covering number (instead of the VC-dimension). First, define

$$\xi_n(y) = \sup_{f\in\mathcal{F}}\Big|\frac{1}{n}\sum_{i=1}^n f(Z_i)1\{Y_i > y\}\sigma_i\Big|.$$

Denote $\bar{p}(y) = \frac{1}{n}\sum_{i=1}^n P(Y_i > y|Z)$.

Case 1 Consider first values of y for which $n\bar{p}(y) = \sum_{i=1}^n P(Y_i > y|Z) \leq 1$. Due to the envelope condition, and the definition of Rademacher random variables, we have

$$\Big|\frac{1}{n}\sum_{i=1}^n f(Z_i)1\{Y_i > y\}\sigma_i\Big| \leq \bar{F}\frac{1}{n}\sum_{i=1}^n 1\{Y_i > y\}, \forall f \in \mathcal{F}.$$

Taking expectations we have $\mathbb{E}[\xi_n(y)] \leq \bar{F}\mathbb{E}\Big[\frac{1}{n}\sum_{i=1}^n 1\{Y_i > y\}\Big|Z\Big] = \bar{F}\bar{p}(y)$ and the right hand side is bounded by $\bar{F}\frac{1}{n}$ for this particular case.

Case 2 Consider now values of y such that $n\bar{p}(y) > 1$. Define the random variable $N_y = \sum_{i=1}^n 1\{Y_i > y\}$. Then we can write

$$\frac{1}{n}\sum_{i=1}^n f(Z_i)1\{Y_i > y\}\sigma_i = \begin{cases} 0 \text{ if } N_y = 0 \\ \frac{N_y}{n}\frac{1}{N_y}\sum_{i=1}^n f(Z_i)1\{Y_i > y\}\sigma_i \text{ if } N_y \geq 1. \end{cases}$$

If $N_y \geq 1$, then

$$\begin{aligned}
\xi_n(y) &= \sup_{f\in\mathcal{F}}\Big|\frac{N_y}{n}\frac{1}{N_y}\sum_{i=1}^n f(Z_i)\sigma_i 1\{Y_i > y\}\Big| \\
&= \sup_{f\in\mathcal{F}}\Big|\frac{N_y}{n}\frac{1}{N_y}\sum_{i=1}^n f(Z_i)\sigma_i 1\{Y_i > y\} - \bar{p}(y)\frac{1}{N_y}\sum_{i=1}^n f(Z_i)1\{Y_i > y\}\sigma_i + \\
&+ \bar{p}(y)\frac{1}{N_y}\sum_{i=1}^n f(Z_i)1\{Y_i > y\}\sigma_i\Big| \\
&\leq \Big|\frac{N_y}{n} - \bar{p}(y)\Big|\sup_{f\in\mathcal{F}}\Big|\frac{1}{N_y}\sum_{i=1}^n f(Z_i)1\{Y_i > y\}\sigma_i\Big| + \bar{p}(y)\sup_{f\in\mathcal{F}}\Big|\frac{1}{N_y}\sum_{i=1} f(Z_i)\sigma_i 1\{Y_i > y\}\Big|.
\end{aligned}$$

Denote $\mathbb{E}_\sigma$ the expectation only with respect to the Rademacher random variables σ. Conditional on N_y, $\xi_n(y)$ sums over N_y terms. Therefore, for a constant $0 < C_1 < \infty$ that only depend on $\bar{F}$,

$$\mathbb{E}_\sigma\Big[\bar{p}(y)\sup_{f\in\mathcal{F}}\Big|\frac{1}{N_y}\sum_{i=1} f(Z_i)\sigma_i 1\{Y_i > y\}\Big|\Big] \leq C_1\bar{p}(y)\sqrt{V}g(N_y)$$

by Theorem 5.22 in Wainwright (2019) (Lemma F.8 in Appendix F) where $g(.)$ is defined in Lemma F.4. Similarly,

$$\mathbb{E}_\sigma\Big[\Big|\frac{N_y}{n} - \bar{p}(y)\Big|\sup_{f\in\mathcal{F}}\Big|\frac{1}{N_y}\sum_{i=1}^n f(Z_i)1\{Y_i > y\}\sigma_i\Big||Z\Big] \leq \Big|\frac{N_y}{n} - \bar{p}(y)\Big|C_1\sqrt{V}g(N_y).$$

For $N_y \geq 1$, it follows by the law of iterated expectations,

$$\mathbb{E}[\xi_n(y)|N_y, Z] \leq |\frac{N_y}{n} - \bar{p}(y)|C_1\sqrt{V}g(N_y) + C_1\bar{p}(y)\sqrt{V}g(N_y) \leq |\frac{N_y}{n} - \bar{p}(y)|C_1\sqrt{V} + C_1\bar{p}(y)\sqrt{V}g(N_y) \tag{C.8}$$

where the last inequality follows by the definition of the $g(.)$ function and the fact that $N_y \in \{1, 2, \cdots\}$. For $N_y = 0$ instead, we have

$$\mathbb{E}[\xi_n(y)|N_y, Z] \leq \Big|\frac{N_y}{n} - \bar{p}(y)\Big|C_1\sqrt{V}g(N_y) + C_1\bar{p}(y)\sqrt{V}g(N_y) = 0$$

by the definition of $g(.)$. Hence, the bound in Equation (C.8) always holds.

Unconditional expectation We are left to bound the unconditional expectation with respect to N_y. Notice first that

$$\mathbb{E}\Big[\Big|\frac{N_y}{n} - \bar{p}(y)\Big||Z\Big] \leq C_1\sqrt{V}\sqrt{\mathbb{E}\Big[\Big|\frac{N_y}{n} - \bar{p}(y)\Big|^2\Big|Z\Big]} = C_1\sqrt{V}\sqrt{\mathrm{Var}\Big(\frac{1}{n}\sum_{i=1}^n 1\{Y_i > y\}\Big|Z\Big)}.$$

By independence assumption,

$$C_1\sqrt{V}\sqrt{\mathrm{Var}\Big(\frac{1}{n}\sum_{i=1}^n 1\{Y_i > y\}\Big|Z\Big)} \leq C_1\sqrt{V}\sqrt{\frac{1}{n^2}\sum_{i=1}^n P(Y_i > y|Z)}.$$

In addition, since $n\bar{p}(y) > 1$, by Lemma F.9 , $\mathbb{E}[g(N_y)|Z] \leq \frac{2}{\sqrt{n}\sqrt{\bar{p}(y)}}$. Combining the inequalities, it follows

$$\mathbb{E}[\xi_n(y)|Z] \leq C_1\sqrt{V}\sqrt{\frac{1}{n^2}\sum_{i=1}^n P(Y_i > y|Z)} + \bar{p}(y)C_1\sqrt{V}\frac{2}{\sqrt{n}\sqrt{\bar{p}(y)}} \leq 2(1 + C_1)\sqrt{\frac{V}{n}}\sqrt{\frac{1}{n}\sum_{i=1}^n P(Y_i > y|Z)}.$$

This bound is larger than the bound derived for $n\bar{p}(y) < 1$, up to a constant factor $C_{\bar{F}} < \infty$. Therefore,

$$\mathbb{E}[\xi_n(y)|Z] \leq C_{\bar{F}}\sqrt{\frac{V}{n}}\sqrt{\frac{1}{n}\sum_{i=1}^{n} P(Y_i > y|Z)}.$$

Integral bound We can now write

$$\int_0^\infty \mathbb{E}[\xi_n(y)|Z]dy \leq \int_0^\infty C_{\bar{F}}\sqrt{\frac{V}{n}}\sqrt{\sum_{i=1}^{n}\frac{P(Y_i > y|Z)}{n}}dy = \int_0^1 C_{\bar{F}}\sqrt{\frac{V}{n}}\sqrt{\sum_{i=1}^{n}\frac{P(Y_i > y|Z)}{n}}dy + \int_1^\infty C_{\bar{F}}\sqrt{\frac{V}{n}}\sqrt{\sum_{i=1}^{n}\frac{P(Y_i > y|Z)}{n}}dy.$$

The first term is bounded as follows. $\int_0^1 C_{\bar{F}}\sqrt{\frac{V}{n}}\sqrt{\sum_{i=1}^{n}\frac{P(Y_i > y|Z)}{n}}dy \leq C_{\bar{F}}\sqrt{\frac{V}{n}}$. The second term is bounded instead as follows.

$$\int_1^\infty C_{\bar{F}}\sqrt{\frac{V}{n}}\sqrt{\sum_{i=1}^{n}\frac{P(Y_i > y|Z)}{n}}dy \leq C_{\bar{F}}\sqrt{\frac{V}{n}}\int_1^\infty \sqrt{\sum_{i=1}^{n}\frac{\mathbb{E}[Y_i^{2+u}|Z]}{ny^{2+u}}}dy \leq C_{\bar{F}}\sqrt{\frac{V}{n}}\int_1^\infty \sqrt{\frac{B}{y^{2+u}}}dy \leq C'_{\bar{F}}\sqrt{\frac{VB}{n}}$$

for a constant $C'_{\bar{F}} < \infty$ that only depend on $\bar{F}$ and u. □

Lemma C.7. *Let Π be a function class with $\pi : \mathbb{R}^d \mapsto \{0,1\}$ for any $\pi \in \Pi$, with finite VC-dimension, denoted as $\mathrm{VC}(\Pi)$. Let $X^1, ..., X^{h_i}, V \in \mathbb{R}^d, O \in \mathbb{R}^k, Y \in \mathbb{R}$. Let $K = (V_i, X_i^1, X_i^2, ..., X_i^{h_i}, O_i)_{i=1}^n$. Suppose that $\mathbb{E}[Y_i^{2+u}|K] < B_1 < \infty$ for all $i \in \{1, \cdots, n\}$ for some $u > 0$ and $(Y_i)_{i=1}^n$ are pairwise independent conditional on K. Let $\sigma_1, \cdots, \sigma_n$ be independent Rademacher random variables, independent of $(K_i, Y_i)_{i=1}^n$. Let $f : \mathbb{Z}\times\mathbb{R}^k \mapsto \mathbb{R}$ be L-Lipschitz in its first argument. Assume that $\mathbb{E}[f(0,O_i)^{2+u}Y_i^{2+u}|K] < B_2 < \infty$ for all $i \in \{1, \cdots, n\}$ for some $u > 0$. Let $\bar{h} = \max_i h_i$. Then for a constant $C < \infty$ that only depend on u,*

$$\mathbb{E}_{Y,\sigma}\Big[\sup_{\pi\in\Pi}\Big|\frac{1}{n}\sum_{i=1}^{n} f\Big(\sum_{k\in\{1,\dots,h\}}\pi(X_i^{(k)}), O_i\Big)\pi(V_i)\sigma_i Y_i\Big|\Big] \leq C(L+1)\sqrt{\frac{(\bar{h}+1)\log(\bar{h}+1)\mathrm{VC}(\Pi)(B_1+B_2)}{n}}, \tag{C.9}$$

where $\mathbb{E}_{Y,\sigma}$ denotes the conditional expectation taken only with respect to the variables $(Y_i, \sigma_i)_{i=1}^n$ conditional on K.

Proof of Lemma C.7. The proof follows similarly to Lemma C.3. We can consider the case where all units correspond to sums over h components, and discuss the case where this is violated at the end of the proof.

First decomposition First, we add and subtract the value of the function $f(0, O_i)$ at zero. Namely, the left hand side in Equation (C.9) equals

$$\begin{aligned}
&\mathbb{E}_{Y,\sigma}\Big[\sup_{\pi\in\Pi}\Big|\frac{1}{n}\sum_{i=1}^n \sigma_i\Big(f\Big(\sum_{k\in\{1,\dots,h\}}\pi(X_i^{(k)}), O_i\Big) - f\Big(0, O_i\Big) + f\Big(0, O_i\Big)\Big)Y_i\pi(V_i)\Big|\Big] \\
&\leq \underbrace{\mathbb{E}_{Y,\sigma}\Big[\sup_{\pi\in\Pi}\Big|\frac{1}{n}\sum_{i=1}^n \sigma_i\Big(f\Big(\sum_{k\in\{1,\dots,h\}}\pi(X_i^{(k)}), O_i\Big) - f\Big(0, O_i\Big)\Big)Y_i\pi(V_i)\Big|\Big]}_{(1)} \qquad \text{(C.10)} \\
&+ \underbrace{\mathbb{E}_{Y,\sigma}\Big[\sup_{\pi\in\Pi}\Big|\frac{1}{n}\sum_{i=1}^n \sigma_i f\Big(0, O_i\Big)Y_i\pi(V_i)\Big|\Big]}_{(2)}
\end{aligned}$$

where the last inequality follows by the triangular inequality.

Decomposition with integral We bound first (1). We write

$$\begin{aligned}
&\mathbb{E}_{Y,\sigma}\Big[\sup_{\pi\in\Pi}\Big|\frac{1}{n}\sum_{i=1}^n \sigma_i\Big(f\Big(\sum_{k\in\{1,\dots,h\}}\pi(X_i^{(k)}), O_i\Big) - f\Big(0, O_i\Big)\Big)Y_i\pi(V_i)\Big|\Big] \\
&= \mathbb{E}_{Y,\sigma}\Big[\sup_{\pi\in\Pi}\Big|\frac{1}{n}\sum_{i=1}^n \sigma_i\Big(f\Big(\sum_{k\in\{1,\dots,h\}}\pi(X_i^{(k)}), O_i\Big) - f\Big(0, O_i\Big)\Big)|Y_i|\text{sign}(Y_i)\pi(V_i)\Big|\Big] \qquad \text{(C.11)} \\
&\leq 2\mathbb{E}_{Y,\tilde{\sigma}}\Big[\sup_{\pi\in\Pi}\Big|\frac{1}{n}\sum_{i=1}^n \tilde{\sigma}_i\Big(f\Big(\sum_{k\in\{1,\dots,h\}}\pi(X_i^{(k)}), O_i\Big) - f\Big(0, O_i\Big)\Big)|Y_i|\pi(V_i)\Big|\Big]
\end{aligned}$$

where $\tilde{\sigma}_i$ are Rademacher random variables independent of $Y_1, ..., Y_n$.[46] Since $|Y_i| > 0$, we have

$$\begin{aligned}
\text{(C.11)} &= 2\mathbb{E}_{Y,\tilde{\sigma}}\Big[\sup_{\pi\in\Pi}\Big|\frac{1}{n}\sum_{i=1}^n \tilde{\sigma}_i\Big(f\Big(\sum_{k\in\{1,\dots,h\}}\pi(X_i^{(k)}), O_i\Big) - f\Big(0, O_i\Big)\Big)\int_0^\infty 1\{|Y_i| > y\}dy\pi(V_i)\Big|\Big] \\
&\leq 2\mathbb{E}_{Y,\tilde{\sigma}}\Big[\sup_{\pi\in\Pi}\int_0^\infty\Big|\frac{1}{n}\sum_{i=1}^n \tilde{\sigma}_i\Big(f\Big(\sum_{k\in\{1,\dots,h\}}\pi(X_i^{(k)}), O_i\Big) - f\Big(0, O_i\Big)\Big)1\{|Y_i| > y\}\pi(V_i)\Big|dy\Big] \\
&\leq 2\int_0^\infty \mathbb{E}_{Y,\tilde{\sigma}}\Big[\sup_{\pi\in\Pi}\Big|\frac{1}{n}\sum_{i=1}^n \tilde{\sigma}_i\Big(f\Big(\sum_{k\in\{1,\dots,h\}}\pi(X_i^{(k)}), O_i\Big) - f\Big(0, O_i\Big)\Big)1\{|Y_i| > y\}\pi(V_i)\Big|\Big]dy
\end{aligned}$$

[46] To check the last claim the reader might consider that $P(\tilde{\sigma}_i = 1|Y_i) = P(\sigma_i\text{sign}(Y_i) = 1|Y_i) = 1/2$.

where the last inequality follows by the properties of the supremum function and Fubini theorem.

Lipschitz property We decompose the supremum over $\pi \in \Pi$ as follows.

$$\int_0^\infty \mathbb{E}_{Y,\tilde{\sigma}}\Big[\sup_{\pi\in\Pi}\Big|\frac{1}{n}\sum_{i=1}^n \tilde{\sigma}_i\Big(f\Big(\sum_{k\in\{1,\dots,h\}}\pi(X_i^{(k)}), O_i\Big) - f\Big(0, O_i\Big)\Big)1\{|Y_i| > y\}\pi(V_i)\Big|\Big]dy$$
$$\leq \int_0^\infty \mathbb{E}_{Y,\tilde{\sigma}}\Big[\sup_{\pi_1\in\Pi,\pi_2\in\Pi}\Big|\frac{1}{n}\sum_{i=1}^n \tilde{\sigma}_i\Big(f\Big(\sum_{k\in\{1,\dots,h\}}\pi_2(X_i^{(k)}), O_i\Big) - f\Big(0, O_i\Big)\Big)1\{|Y_i| > y\}\pi_1(V_i)\Big|\Big]dy.$$

Let $\phi_i(N) = f\Big(N, O_i\Big) - f\Big(0, O_i\Big)$. Conditional on the data, ϕ_i is not random. By assumption it is Lipschitz and $\phi_i(0) = 0$. By Lemma C.5,

$$\int_0^\infty \mathbb{E}_{Y,\tilde{\sigma}}\Big[\sup_{\pi_1\in\Pi,\pi_2\in\Pi}\Big|\frac{1}{n}\sum_{i=1}^n \tilde{\sigma}_i\Big(f\Big(\sum_{k\in\{1,\dots,h\}}\pi_2(X_i^{(k)}), O_i\Big) - f\Big(0, O_i\Big)\Big)1\{|Y_i| > y\}\pi_1(V_i)\Big|\Big]dy$$
$$\leq 2L\int_0^\infty \mathbb{E}_{Y,\tilde{\sigma}}\Big[\sup_{\pi_1\in\Pi,\pi_2\in\Pi}\Big|\frac{1}{n}\sum_{i=1}^n \tilde{\sigma}_i\Big(\sum_{k\in\{1,\dots,h\}}\pi_2(X_i^{(k)})\Big)1\{|Y_i| > y\}\pi_1(V_i)\Big|\Big]dy. \tag{C.12}$$

Supremum over Π We decompose the supremum as follows:

$$L\int_0^\infty \mathbb{E}_{Y,\tilde{\sigma}}\Big[\sup_{\pi_1\in\Pi,\pi_2\in\Pi}\Big|\frac{1}{n}\sum_{i=1}^n \tilde{\sigma}_i\Big(\sum_{k\in\{1,\cdots,h\}}\pi_2(X_i^{(k)})\Big)1\{|Y_i| > y\}\pi_1(V_i)\Big|\Big]dy$$
$$\leq \underbrace{L\int_0^\infty \mathbb{E}_{Y,\tilde{\sigma}}\Big[\sup_{\pi_1\in\Pi,\pi_2\in\Pi,\cdots,\pi_{h+1}\in\Pi}\Big|\frac{1}{n}\sum_{i=1}^n \tilde{\sigma}_i\Big(\sum_{k\in\{1,\dots,h\}}\pi_{k+1}(X_i^{(k)})\Big)1\{|Y_i| > y\}\pi_1(V_i)\Big|\Big]dy}_{(J)}.$$

We re-parametrize the class of functions as follows.

$$(J) = \int_0^\infty \mathbb{E}_{Y,\tilde{\sigma}}\Big[\sup_{\pi_1\in\Pi,\tilde{\pi}_2\in\Pi_2,\cdots,\tilde{\pi}_{h+1}\in\Pi_{h+1}}\Big|\frac{1}{n}\sum_{i=1}^n \tilde{\sigma}_i\Big(\sum_{k\in\{1,\dots,h\}}\tilde{\pi}_{k+1}(X_i^{(1)},\cdots,X_i^{(h)})\Big)1\{|Y_i| > y\}\pi_1(V_i)\Big|\Big]dy,$$

where $\Pi_2, ..., \Pi_{h+1}$, are such that $\tilde{\pi}_j \in \Pi_j : \tilde{\pi}_j(X_i^{(1)},\cdots,X_i^{(k)},\cdots,X_i^{(h)}) = \pi(X_i^{(j)})$. By Theorem 29.4 in Devroye et al. (2013), $\text{VC}(\Pi_j) = \text{VC}(\Pi)$ for all $j \in \{1, ..., h\}$.[47]
Let

$$\tilde{\Pi} = \Big\{\pi_1\Big(\sum_{j=1}^h \pi_{j+1}\Big), \pi_k \in \Pi_k, k = 1,\cdots,h+1\Big\}.$$

[47]The reader might recognize that each $\pi_j \in \Pi_j$ can be written as the sum of π and functions constant at zero.

For any fixed point data point z_1^n, by Lemma C.2, the Dudley's integral of the function class $\tilde{\Pi}(z_1^n)$ is bounded by $C\sqrt{(h+1)\log(h+1)\text{VC}(\Pi)}$, for a finite constant C. Therefore, by Lemma C.6 and C.2

$$
\begin{aligned}
&\int_0^\infty \mathbb{E}_{Y,\tilde{\sigma}}\Big[\sup_{\pi_1\in\Pi_1,\pi_2\in\Pi_2,\dots,\pi_{h+1}\in\Pi_{h+1}}\Big|\frac{1}{n}\sum_{i=1}^n \tilde{\sigma}_i\Big(\sum_{k\in\{1,\dots,h\}}\pi_{k+1}(X_i^{(k)})\Big)1\{|Y_i|>y\}\pi_1(V_i)\Big|\Big]dy \\
&\leq C\sqrt{B_1(h+1)\frac{\text{VC}(\Pi)\log(h+1)}{n}}
\end{aligned}
$$

for a constant $C<\infty$.

Remaining component Next, we bound the term (2) in Equation (C.10). Following the same argument used for term (1), we have

$$
\begin{aligned}
\mathbb{E}_{Y,\tilde{\sigma}}\Big[\sup_{\pi\in\Pi}\Big|\frac{1}{n}\sum_{i=1}^n\tilde{\sigma}_i f\Big(0,O_i\Big)Y_i\pi(V_i)\Big|\Big] &\leq 2\mathbb{E}_{Y,\tilde{\sigma}}\Big[\sup_{\pi\in\Pi}\Big|\frac{1}{n}\sum_{i=1}^n\tilde{\sigma}_i|f\Big(0,O_i\Big)Y_i|\pi(V_i)\Big|\Big] \\
&\leq \int_0^\infty 2\mathbb{E}_{Y,\tilde{\sigma}}\Big[\sup_{\pi\in\Pi}\Big|\frac{1}{n}\sum_{i=1}^n\tilde{\sigma}_i 1\{|f\Big(0,O_i\Big)Y_i|>y\}\pi(V_i)\Big|\Big]dy.
\end{aligned}
$$

Since Π has finite VC-dimension, and it contains functions mapping to $\{0,1\}$, by Theorem 2.6.7 of Van Der Vaart and Wellner (1996) [48], $\int_0^2\sqrt{\mathcal{N}_2(u,\Pi(z_1^n))}du < C\sqrt{\text{VC}(\Pi)}$ for a constant C. Hence, by Lemma C.6 and C.2,

$$
\int_0^\infty 2\mathbb{E}_{Y,\tilde{\sigma}}\Big[\sup_{\pi\in\Pi}\Big|\frac{1}{n}\sum_{i=1}^n\tilde{\sigma}_i 1\{|f\Big(0,O_i\Big)Y_i|>y\}\pi(V_i)\Big|\Big]dy \leq c_0'\sqrt{\frac{B_2\text{VC}(\Pi)}{n}}
$$

for a constant $c_0'<\infty$.

Conclusion We are now left to discuss the case where h_i is different for each individual. In this case, we can assign to each unit covariates $\{X_i^{(1)},\cdots,X_i^{(h_i)},\emptyset,\emptyset,\cdots,\emptyset\}$ with in total $\bar{h}$ and parametrizing $\pi(\emptyset)=0$. The rest of the proof follows similarly as before. □

C.2 Lemmas for Identification

We conclude this section discussing the lemmas useful for identification.

[48]The argument is the same used in the proof of Lemma C.2.

Lemma C.8. *Let Assumption 2.1, 2.2, 2.4 hold. Then*

$$\mathbb{E}\Big[\frac{1\{S_i(\pi) = \sum_{k\in N_i} D_k, \pi(X_i) = D_i\}}{e\Big(\pi(X_i), S_i(\pi), Z_{k\in N_i}, Z_i, |N_i|\Big)} Y_i \Big| A, Z\Big] = m\Big(\pi(X_i), S_i(\pi), Z_i, |N_i|\Big).$$

Proof. Under Assumption 2.1, we can write

$$\begin{aligned}&\mathbb{E}\Big[\frac{1\{S_i(\pi) = \sum_{k\in N_i} D_k, \pi(X_i) = D_i\}}{e\Big(\pi(X_i), S_i(\pi), Z_{k\in N_i}, Z_i, |N_i|\Big)} Y_i \Big| A, Z\Big]\\ &= \mathbb{E}\Big[\frac{1\{S_i(\pi) = \sum_{k\in N_i} D_k, \pi(X_i) = D_i\}}{e\Big(\pi(X_i), S_i(\pi), Z_{k\in N_i}, Z_i, |N_i|\Big)} r\Big(\pi(X_i), S_i(\pi), |N_i|, Z_i, \varepsilon_i\Big) \Big| A, Z\Big].\end{aligned} \tag{C.13}$$

Under Assumption 2.2, we can then write

$$\begin{aligned}(\text{C.13}) = \mathbb{E}\Big[&\frac{1\{S_i(\pi) = \sum_{k\in N_i} D_k, \pi(X_i) = D_i\}}{e\Big(\pi(X_i), S_i(\pi), Z_{k\in N_i}, Z_i, |N_i|\Big)} |A, Z\Big] \times\\ &\times \mathbb{E}\Big[r\Big(\pi(X_i), S_i(\pi), |N_i|, Z_i, \varepsilon_i\Big) \Big| A, Z\Big].\end{aligned}$$

By the first condition in Assumption 2.2 $\mathbb{E}\Big[\frac{1\{S_i(\pi) = \sum_{k\in N_i} D_k, \pi(X_i) = D_i\}}{e\Big(\pi(X_i), S_i(\pi), Z_{k\in N_i}, Z_i, |N_i|\Big)} |A, Z\Big] = 1$. Finally, by Assumption 2.2 and Assumption 2.4

$$\mathbb{E}\Big[r\Big(\pi(X_i), S_i(\pi), |N_i|, Z_i, \varepsilon_i\Big) \Big| A, Z\Big] = m\Big(\pi(X_i), S_i(\pi), |N_i|, Z_i, \Big),$$

where $m(\cdot)$ not being individual specific follows by Assumption 2.4. □

Lemma C.9. *Let $S_i(\pi) = \sum_{k\in N_i} \pi(X_k)$. Let Assumption 2.1, 2.2, 2.4 hold. Then*

$$\begin{aligned}&\mathbb{E}\Big[\frac{1\{S_i(\pi) = \sum_{k\in N_i} D_k, d = D_i\}}{e^c\Big(\pi(X_i), S_i(\pi), Z_{k\in N_i}, Z_i, |N_i|\Big)} \Big(Y_i - m^c\Big(\pi(X_i), S_i(\pi), Z_i, |N_i|\Big)\Big) + m^c\Big(\pi(X_i), S_i(\pi), Z_i, |N_i|\Big) \Big| A, Z\Big]\\ &= m\Big(\pi(X_i), S_i(\pi), Z_i, |N_i|\Big),\end{aligned}$$

if either $e^c = e$ or (and) $m^c = m$.

Proof. Whenever $e^c = e$, similarly to what discussed in Lemma C.8, we obtain that

$$\mathbb{E}\Big[\frac{1\{S_i(\pi) = \sum_{k\in N_i} D_k, d = D_i\}}{e^c\Big(\pi(X_i), S_i(\pi), Z_{k\in N_i}, Z_i, |N_i|\Big)} Y_i \Big| A, Z\Big] = m\Big(\pi(X_i), S_i(\pi), Z_i, |N_i|\Big).$$

The remaining component depending on m^c cancels out similarly as above. Let now $m^c = m$. Then

$$
\begin{aligned}
&\mathbb{E}\Big[\frac{1\{S_i(\pi) = \sum_{k\in N_i} D_k, d = D_i\}}{e^c\Big(\pi(X_i), S_i(\pi), Z_{k\in N_i}, Z_i, |N_i|\Big)}\Big(Y_i - m\Big(\pi(X_i), S_i(\pi), Z_i, |N_i|\Big)\Big)\Big|A, Z\Big] \\
&= \mathbb{E}\Big[\frac{1\{S_i(\pi) = \sum_{k\in N_i} D_k, d = D_i\}}{e^c\Big(\pi(X_i), S_i(\pi), Z_{k\in N_i}, Z_i, |N_i|\Big)}\Big(r\Big(\pi(X_i), S_i(\pi), Z_i, |N_i|, \varepsilon_i\Big) - m\Big(\pi(X_i), S_i(\pi), Z_i, |N_i|\Big)\Big)\Big|A, Z\Big] \\
&= \mathbb{E}\Big[\frac{1\{S_i(\pi) = \sum_{k\in N_i} D_k, d = D_i\}}{e^c\Big(\pi(X_i), S_i(\pi), Z_{k\in N_i}, Z_i, |N_i|\Big)}\Big|A, Z\Big]\times \\
&\times \mathbb{E}\Big[\Big(r\Big(\pi(X_i), S_i(\pi), Z_i, |N_i|, \varepsilon_i\Big) - m\Big(\pi(X_i), S_i(\pi), Z_i, |N_i|\Big)\Big)\Big|A, Z\Big] = 0.
\end{aligned}
$$

□

C.2.1 Proofs of lemmas in the main text

In this section we prove the lemmas in the main text.

Proof of Lemma 2.1. Under Assumption 2.3, the welfare on the target sample reads as follows

$$
W(\pi) = \frac{1}{|\mathcal{I}|}\sum_{j\in\mathcal{I}}\mathbb{E}\Big[r\Big(\pi(X_j), \sum_{k\in N_j}\pi(X_k), Z_j, |N_j|, \varepsilon_j\Big)\Big].
$$

Under the distributional condition in the statement of the lemma we have

$$
W(\pi) = \mathbb{E}\Big[r\Big(\pi(X_j), \sum_{k\in N_j}\pi(X_k), Z_j, |N_j|, \varepsilon_j\Big)\Big]
$$

is constant in j. In addition, by Assumption 2.1 and 2.2, welfare on the sampled units takes the expression

$$
\frac{1}{n}\sum_{i=1}^{n}\mathbb{E}\Big[r\Big(\pi(X_i), \sum_{k\in N_i}\pi(X_k), Z_i, |N_i|, \varepsilon_i\Big)\Big].
$$

The proof completes again by the distributional assumption in the statement of the lemma. □

Proof of Lemma 2.2. By the law of iterated expectations, we write

$$
\mathbb{E}\Big[r(\pi(X_i), \sum_{k\in N_i}\pi(X_k), Z_i, |N_i|, \varepsilon_i)\Big] = \mathbb{E}\Big[\mathbb{E}\Big[r(\pi(X_i), \sum_{k\in N_i}\pi(X_k), Z_i, |N_i|, \varepsilon_i)\Big|Z_i, N_i, X_{k\in N_i}\Big]\Big]
$$

Under Assumption 2.2, 2.4 we obtain

$$\mathbb{E}\Big[\mathbb{E}\Big[r(\pi(X_i), \sum_{k\in N_i}\pi(X_k), Z_i, |N_i|, \varepsilon_i)\Big|Z_i, N_i, X_{k\in N_i}\Big]\Big] = \mathbb{E}\Big[m(\pi(X_i), \sum_{k\in N_i}\pi(X_k), Z_i, |N_i|)\Big].$$

The rest of the proof follows from Lemma C.8 and the law of iterated expectations. □

Appendix D Regret Bounds

Theorem D.1. *Suppose the conditions in Theorem 3.1 holds. Then,*

$$\mathbb{E}\Big[\sup_{\pi\in\Pi}|W(\pi) - W_n(\hat{\pi}^{aipw}_{m^c,e^c})|\Big]$$
$$\leq |\mathcal{L}|^2(\Gamma_1+\Gamma_2)\frac{\bar{C}}{\delta}\Big(L + \frac{L}{\delta_0^{\sqrt{3}}} + \frac{1}{\delta_0^{\sqrt{3}}}\Big)\mathbb{E}[\mathcal{N}_{n,M}^{M}\mathcal{N}_n^{|\mathcal{L}|}\mathcal{N}_n^{3/2}\sqrt{\log(\mathcal{N}_{n,1})}]^2\sqrt{\frac{\mathrm{VC}(\Pi)}{n}} + \mathcal{K}_\Pi(n),$$

for a finite constant $\bar{C} < \infty$.

Proof of Theorem D.1. Let $S_i(\pi) = \sum_{k\in N_i}\pi(X_k)$. We denote $L = (L_j)_{j=1}^E$ under Assumption 3.2. Recall under Assumption 3.2 that we can decompose covariates Z_i into two components, one L_i having arbitrary dependence and endogenous with respect to the network but with finite support, and the other, Q_i having local dependence but arbitrary support.

Preliminaries Notice first that by Definition 2.1,

$$\mathbb{E}\Big[\sup_{\pi\in\Pi}|W(\pi) - W_n(\pi)|\Big] \leq \mathcal{K}_\Pi(n) + \mathbb{E}\Big[\sup_{\pi\in\Pi}|W^*(\pi) - W_n(\pi)|\Big]$$

where

$$W^*(\pi) = \frac{1}{n}\sum_{i=1}^n\mathbb{E}\Big[r\Big(\pi(X_i), S_i(\pi), Z_i, |N_i|, \varepsilon_i\Big)\Big] = \frac{1}{n}\sum_{i=1}^n\mathbb{E}\Big[m\Big(\pi(X_i), S_i(\pi), Z_i, |N_i|\Big)\Big],$$

and where the second equality follows from Assumption 2.1, 2.2 and 2.4 similarly to what discussed in the proof of Lemma C.9. By Lemma C.9, we have that $W^*(\pi) = \mathbb{E}[W_n(\pi, m^c, e^c)]$ under correct specification of either m^c or e^c. Therefore for $i \in \{1,\cdots,n\}$

$$W^*(\pi) = \frac{1}{n}\sum_{i=1}^n\mathbb{E}\Big[\frac{1\{\pi(X_i) = D_i, \sum_{k\in N_i}\pi(X_k) = \sum_{k\in N_i}D_k\}}{e^c\Big(\pi(X_i), S_i(\pi), O_i\Big)}\Big(Y_i - m^c\Big(\pi(X_i), S_i(\pi), Z_i, |N_i|\Big)\Big)\Big] +$$
$$+ \mathbb{E}\Big[m^c\Big(\pi(X_i), S_i(\pi), Z_i, |N_i|\Big)\Big],$$

where $O_i = \Big(Z_{k\in N_i}, Z_i, |N_i|\Big)$.

Decomposition Using the triangular inequality, we decompose the supremum of the empirical process into three terms as follows.

$$
\begin{aligned}
\mathbb{E}\Big[\sup_{\pi\in\Pi}|W^*(\pi) - W_n(\pi, m^c, e^c)|\Big] \leq & \underbrace{\mathbb{E}\Big[\sup_{\pi\in\Pi}\Big|T_1(\pi, m^c, e^c) - \mathbb{E}[T_1(\pi, m^c, e^c)]\Big|\Big]}_{(A)} \\
& + \underbrace{\mathbb{E}\Big[\sup_{\pi\in\Pi}\Big|T_2(\pi, m^c, e^c) - \mathbb{E}[T_2(\pi, m^c, e^c)]\Big|\Big]}_{(B)} \\
& + \underbrace{\mathbb{E}\Big[\sup_{\pi\in\Pi}\Big|T_3(\pi, m^c) - \mathbb{E}[T_3(\pi, m^c)]\Big|\Big]}_{(C)},
\end{aligned}
$$

where

$$
\begin{aligned}
T_1(\pi, m^c, e^c) &= \frac{1}{n}\sum_{i=1}^{n}\frac{1\{\pi(X_i) = D_i, \sum_{k\in N_i}\pi(X_k) = \sum_{k\in N_i} D_k\}}{e^c\Big(\pi(X_i), S_i(\pi), O_i\Big)}Y_i \\
T_2(\pi, m^c, e^c) &= \frac{1}{n}\sum_{i=1}^{n}\frac{1\{\pi(X_i) = D_i, \sum_{k\in N_i}\pi(X_k) = \sum_{k\in N_i} D_k\}}{e^c\Big(\pi(X_i), S_i(\pi), O_i\Big)}m^c\Big(\pi(X_i), \sum_{k\in N_i}\pi(X_k), Z_i, |N_i|\Big) \\
T_3(\pi, m^c) &= \frac{1}{n}\sum_{i=1}^{n} m^c\Big(\pi(X_i), S_i(\pi), Z_i, |N_i|\Big).
\end{aligned} \tag{D.1}
$$

The rest of the proof consists in providing upper bounds for (A), (B), and (C). We discuss (A) first.

Preparation for (A) We create an independent copy of $(T_1(\cdot), A, L)$, defined as $(T_1'(\cdot), A', L')$ with the same distribution of $(T_1(\cdot), A, L)$. Using Jensen's inequality, we can write

$$
\begin{aligned}
\mathbb{E}\Big[\sup_{\pi\in\Pi}\Big|T_1(\pi, m^c, e^c) - \mathbb{E}[T_1(\pi, m^c, e^c)]\Big|\Big] &= \mathbb{E}\Big[\sup_{\pi\in\Pi}\Big|T_1(\pi, m^c, e^c) - \mathbb{E}[T_1'(\pi, m^c, e^c)]\Big|\Big] \\
&\leq \mathbb{E}\Big[\sup_{\pi\in\Pi}\Big|T_1(\pi, m^c, e^c) - T_1'(\pi, m^c, e^c)\Big|\Big].
\end{aligned}
$$

We now use the law of iterated expectations to write

$$
\mathbb{E}\Big[\sup_{\pi\in\Pi}\Big|T_1(\pi, m^c, e^c) - \mathbb{E}[T_1(\pi, m^c, e^c)]\Big|\Big] \leq \mathbb{E}\Big[\mathbb{E}\Big[\sup_{\pi\in\Pi}\Big|T_1(\pi, m^c, e^c) - T_1'(\pi, m^c, e^c)\Big|\Big|A, L, A', L'\Big]\Big].
$$

Denote each summand corresponding to T_1 as U_i and each summand corresponding to T_1' as U_i'. That is,

$$U_i = \frac{1\{\pi(X_i) = D_i, \sum_{k \in N_i} \pi(X_k) = \sum_{k \in N_i} D_k\}}{e^c\Big(\pi(X_i), S_i(\pi), O_i\Big)} Y_i.$$

Observe that $\Big[(U_i)_{i=1}^n, A, L\Big]$ and $\Big[(U_i')_{i'=1}^n, A', L'\Big]$ are independent by construction. Since we condition on (A, L), (A', L'), we can write

$$U_i\Big(L_i, |N_i|, L_{k \in N_i}, \pi\Big)$$

as an explicit function of $L_i, |N_i|, L_{k \in N_i}$ (and π) conditional on (A, L) and similarly $U_i'(L_i', |N_i|', L_{k \in N_i}', \pi)$. Conditional on $(A, L), (A', L')$, the marginal distribution of U_i, U_i' depends on (A, L, A', L') only through the individual, neighbors' covariates and number of neighbors under Assumption 2.2, Assumption 2.4.[49] Also, observe $U_i(L_i, |N_i|, L_{k \in N_i}, \pi)$ is exchangeable in $L_{k \in N_i}$.[50] Define $\mathcal{N}_n'$ the maximum degree under A' and write $R_n = \max\{\mathcal{N}_n, \mathcal{N}_n'\}$. Define $\tilde{\mathcal{L}}_m \subset \mathcal{L}^m$ the space of vectors of dimension m with each component taking value in $\mathcal{L}$ after removing exchangable duplicates, where two vectors are duplicates if they are equal up-to exchanging the entries of the vectors. We can then write

$$\begin{aligned}
&\mathbb{E}\Big[\sup_{\pi \in \Pi} \Big| T_1(\pi, m^c, e^c) - T_1'(\pi, m^c, e^c)\Big| \Big| A, Z, A', Z'\Big] \\
&= \mathbb{E}\Big[\sup_{\pi \in \Pi} \Big| \frac{1}{n} \sum_{i=1}^n U_i(L_i, |N_i|, L_{k \in N_i}, \pi) - \frac{1}{n} \sum_{i=1}^n U_i'(L_i', |N_i|', L_{k \in N_i}', \pi)\Big| \Big| A, L, A', L'\Big] \\
&\leq \sum_{m, m' \leq R_n} \sum_{l \in \mathcal{L} \times \tilde{\mathcal{L}}_m, l' \in \mathcal{L} \times \tilde{\mathcal{L}}_{m'}} \\
&\mathbb{E}\Big[\sup_{\pi \in \Pi} \Big| \frac{1}{n} \sum_{i=1}^n \Big[U_i(L_i, |N_i|, L_{k \in N_i}, \pi) - U_i'(L_i', |N_i|', L_{k \in N_i}', \pi)\Big] I_i(l, l', m, m')\Big| \Big| A, L, A', L'\Big]
\end{aligned} \tag{D.2}$$

where $I_i(l, l', m, m') = 1\{(L_i, L_{k \in N_i}) = l, (L_i', L_{k \in N_i}') = l', m = |N_i|, m' = |N_i|'\}$. Finally, define $\mathcal{C}_n^M$ the proper cover of the adjacency matrix W_n where two elements

[49] To observe why, note that conditional on (L, A), U_i is just a function of $(L_i, Q_i, L_{k \in N_i}, Q_{k \in N_i}, |N_i|, \varepsilon_i, \varepsilon_{D_i}, \varepsilon_{D_{k \in N_i}})$, where the distribution of $(\varepsilon_i, \varepsilon_{D_i}, \varepsilon_{D_{k \in N_i}})$ given (A, L) only depends on $(L_i, L_{k \in N_i}, |N_i|)$ since ε_{D_i} are *i.i.d.* and exogenous.

[50] This can be observed from Definition 2.3 and Assumption 2.2 for the treatment assignments.

are connected if they are neighbors of M degree either (or both) under A or A'. We can then write

$$\mathbb{E}\Big[\sup_{\pi\in\Pi}\Big|\frac{1}{n}\sum_{i=1}^{n}\Big[U_i(L_i,|N_i|,L_{k\in N_i},\pi)-U_i'(L_i',|N_i|',L_{k\in N_i}',\pi)\Big]I_i(l,l',m,m')\Big|\Big|A,L,A',L'\Big]$$

$$\leq \sum_{\mathcal{C}_n^M(j)\in\mathcal{C}_n^M} \underbrace{\mathbb{E}\Big[\sup_{\pi\in\Pi}\Big|\frac{1}{n}\sum_{i\in\mathcal{C}_n^M(j)}\Big[U_i(L_i,|N_i|,L_{k\in N_i},\pi)-U_i'(L_i',|N_i|',L_{k\in N_i}',\pi)\Big]I_i(l,l',m,m')\Big|\Big|A,L,A',L'\Big]}_{(i)}.$$

The goal is to bound (i).

Analysis of (i) Observe first that each summand $U_i(L_i,|N_i|,L_{k\in N_i},\pi)$ is a function of $\Big(L_i,|N_i|,L_{k\in N_i},\varepsilon_i,Q_i,\varepsilon_{D_i},Q_{k\in N_i},\varepsilon_{D_{k\in N_i}},\pi\Big)$ only. In addition, by construction of $\mathcal{C}_n^M$, each summand is mutually independent with the others conditional on (A,L,A',L') by Assumption 2.5 and the first condition in Assumption 2.2 and Assumption 3.2. Finally, each non-zero summand is evaluated at the same values of individual, neighbors' covariates L_i and number of neighbors for both U_i, U_i'.

Symmetrization Under the above observation, we note that each summand whose indicator I_i is non-zero in (i) is $i.i.d.$ conditional on A,L,A',L' under 2.5, 3.2 the fact that ε_{D_i} are $i.i.d.$ and exogenous (Assumption 2.2), since each summand is evaluated at the same values of $(L_i,L_{k\in N_i},|N_i|)$. The remaining summands are instead equal to zero almost surely conditional on (A,L,A',L'). Hence, using a standard symmetrization argument, we bound (i) as follows

$$
\begin{aligned}
(i) &= \mathbb{E}\Big[\sup_{\pi\in\Pi}\Big|\frac{1}{n}\sum_{i\in\mathcal{C}_n^M(j)}\sigma_i\Big[U_i(L_i,|N_i|,L_{k\in N_i},\pi)-U_i'(L_i',|N_i|',L_{k\in N_i}',\pi)\Big]I_i(l,l',m,m')\Big|\Big|A,L,A',L'\Big] \\
&\leq \underbrace{\mathbb{E}\Big[\sup_{\pi\in\Pi}\Big|\frac{1}{n}\sum_{i\in\mathcal{C}_n^M(j)}\sigma_iU_i(L_i,|N_i|,L_{k\in N_i},\pi)I_i(l,l',m,m')\Big|\Big|A,L,A',L'\Big]}_{(j)} \\
&+ \underbrace{\mathbb{E}\Big[\sup_{\pi\in\Pi}\Big|\frac{1}{n}\sum_{i\in\mathcal{C}_n^M(j)}\sigma_iU_i'(L_i',|N_i|',L_{k\in N_i}',\pi)I_i(l,l',m,m')\Big|\Big|A,L,A',L'\Big]}_{(jj)}
\end{aligned}
\tag{D.3}
$$

where σ_i are $i.i.d.$ Rademacher random variables. Note here that a key step is to have $i.i.d.$ of all non-zero components. For those components that are zero, the value of σ_i is not relevant conditional on (A,L,A',L'), and hence we can just use σ_i to be all $i.i.d.$. We then bound (j) (and symmetrically (jj)) with Lemma C.7. To invoke Lemma C.7 we observe that under Assumption 3.1, by Lemma C.4

$$
\phi_i(d,s) := \frac{1\{d=D_i, s=\sum_{k\in N_i}D_k\}}{e^c\Big(d,N,O_i\Big)}
$$

is $\frac{2}{\delta}$-Lipschitz in s conditional on the data. In addition, the function is bounded by Assumption 3.1, therefore

$$
\mathbb{E}\Big[Y_i^3\frac{1\{\pi(X_i)=D_i, 0=\sum_{k\in N_i}D_k\}}{e^c\Big(\pi(X_i),0,O_i\Big)^3}\Big|A,L\Big] < \frac{\Gamma_1^2}{\delta_0^3} < \infty,
$$

since $\sigma(L)\subseteq\sigma(Z)$ Since Y_i is independent of (A',L') the moment bound also holds conditionally on (A',L'). Hence we can invoke Lemma C.7 for each term (j) and (jj) symmetrically.

Collecting the outer sum for (A) Using a symmetric argument for each of the two terms in Equation (D.3), we can write by Lemma C.7 and collecting the outer sums in Equation (D.2)

$$
(A) \leq 2|\mathcal{L}|^2\bar{C}\mathbb{E}\Big[|\tilde{\mathcal{L}}_{\mathcal{N}_n}|^2R_n^2|\mathcal{C}_n^M|\frac{\Gamma_1 L}{\delta\delta_0^{3/2}}\sqrt{\frac{\mathcal{N}_n\log(\mathcal{N}_n)}{n}}\Big],
$$

for a universal constant $\bar{C}$.

Bounding the constants to obtain a polynomial bound in the degree We can bound $R_n \leq \mathcal{N}_n + \mathcal{N}_n'$, and $|\mathcal{C}_n^M| \leq \mathcal{N}_E^M \times \mathcal{N}_E^{M'}$ (see Lemma C.1). Note that $\mathcal{N}_E^M, \mathcal{N}_E^{M'}$ are independent. To characterize $|\tilde{\mathcal{L}}_{\mathcal{N}_n}|$, we can assume, since $|\mathcal{L}|$ is fixed with n that $\mathcal{N}_n > |\mathcal{L}|$. Observe that we can find at most $\mathcal{N}_n^{|\mathcal{L}|}$ many exchangeable combinations. To observe why, note that the number of exchangeable combinations is less or equal than the the number of distinct values that the vector $[V_1, \cdots, V_{|\mathcal{L}|}]$ can take with $V_1 \in \{1, \cdots, \mathcal{N}_n\}$.[51] The number of entries is of order $\mathcal{N}_n^{|\mathcal{L}|}$.

Conclusion Finally, observe that the same reasoning directly applies also to (B), (C), completing the proof.

□

D.1 Theorem 3.1 and Theorem 3.5

We state these two theorems as corollaries of Theorem D.1.

Corollary. *Theorem 3.1 holds.*

Proof. Following Kitagawa and Tetenov (2018),

$$\begin{aligned}\mathbb{E}\Big[\sup_{\pi\in\Pi} W(\pi) - W(\hat{\pi}_{m^c,e^c})\Big] &= \mathbb{E}\Big[\sup_{\pi\in\Pi} W(\pi) - W_n(\hat{\pi}_{m^c,e^c}, m^c, e^c) + W_n(\hat{\pi}_{m^c,e^c}, m^c, e^c) - W(\hat{\pi}_{m^c,e^c})\Big] \\ &\leq \mathbb{E}\Big[\sup_{\pi\in\Pi} W(\pi) - W_n(\pi, m^c, e^c) + W_n(\hat{\pi}_{m^c,e^c}, m^c, e^c) - W(\hat{\pi}_{m^c,e^c})\Big] \\ &\leq \mathbb{E}\Big[2\sup_{\pi\in\Pi} |W(\pi) - W_n(\pi, m^c, e^c)|\Big]\end{aligned}$$

where the last inequality follows by the triangular inequality. The bound on the right hand side is given by Theorem D.1. □

Corollary. *Theorem 3.5 hold.*

Proof. Let $O_i = \Big(Z_{k\in N_i}, Z_i, |N_i|\Big)$ and $I_i(\pi) = 1\{D_i = \pi(X_i), \sum_{k\in N_i} D_k = S_i(\pi)\}$. We

[51] Namely, we can represent the vector L_i as a vector of binary dummies. Due to exchangeability we are interested in the number of combinations that the sum of $\mathcal{N}_n$ of such vectors can take.

have

$$\mathbb{E}\Big[\sup_{\pi\in\Pi} W(\pi) - W(\hat{\pi}^{aipw}_{\hat{m},\hat{e}})\Big] \leq 2\,\underbrace{\mathbb{E}\Big[\sup_{\pi\in\Pi}|W_n^{aipw}(\pi, m^c, e^c) - W(\pi)|\Big]}_{(I)}$$
$$+\, 2\,\underbrace{\mathbb{E}\Big[\sup_{\pi\in\Pi}|W_n^{aipw}(\pi, \hat{m}, \hat{e}) - W_n(\pi, m^c, e^c)|\Big]}_{(II)}.$$

Term (I) is bounded by Theorem D.1.

Bounding (II) We now study (II). We separate two components of (II).

$$\begin{aligned}(II) = \mathbb{E}\Big[\sup_{\pi\in\Pi}\Big|\frac{1}{n}\sum_{i=1}^n \frac{I_i(\pi)}{\hat{e}(\pi(X_i), S_i(\pi), O_i)}\Big(Y_i - \hat{m}(\pi(X_i), S_i(\pi), Z_i, |N_i|\Big) \\ + \frac{1}{n}\sum_{i=1}^n \Big(\hat{m}(\pi(X_i), S_i(\pi), Z_i, |N_i|\Big) - m^c(\pi(X_i), S_i(\pi), Z_i, |N_i|\Big)\Big) \\ - \frac{1}{n}\sum_{i=1}^n \frac{I_i(\pi)}{e^c(\pi(X_i), S_i(\pi), O_i)}\Big(Y_i - m^c(\pi(X_i), S_i(\pi), Z_i, |N_i|\Big)\Big|\Big].\end{aligned} \tag{D.4}$$

We can now use the triangular inequality.

$$\begin{aligned}(D.4) \leq \mathbb{E}\Big[\sup_{\pi\in\Pi}\Big|\frac{1}{n}\sum_{i=1}^n \frac{I_i(\pi)}{e^c(\pi(X_i), S_i(\pi), O_i)\hat{e}(\pi(X_i), S_i(\pi), O_i)}(Y_i - m^c(\pi(X_i), S_i(\pi), Z_i, |N_i|))\times \\ \times \Big(e^c(\pi(X_i), S_i(\pi), O_i) - \hat{e}(\pi(X_i), S_i(\pi), O_i)\Big)\Big|\Big] \\ + \mathbb{E}\Big[\sup_{\pi\in\Pi}\Big|\frac{1}{n}\sum_{i=1}^n \Big(\hat{m}(\pi(X_i), S_i(\pi), Z_i, |N_i|\Big) - m^c(\pi(X_i), S_i(\pi), Z_i, |N_i|\Big)\Big)\Big|\Big] \\ + \mathbb{E}\Big[\sup_{\pi\in\Pi}\Big|\frac{1}{n}\sum_{i=1}^n \frac{I_i(\pi)}{\hat{e}(\pi(X_i), S_i(\pi), O_i)}(\hat{m}(\pi(X_i), S_i(\pi), Z_i, |N_i|) - m^c(\pi(X_i), S_i(\pi), Z_i, |N_i|)\Big|\Big].\end{aligned}$$

By Assumption 3.1, and Holder's inequality we have that

$$\begin{aligned}\mathbb{E}\Big[\sup_{\pi\in\Pi}\Big|\frac{1}{n}\sum_{i=1}^n \frac{I_i(\pi)}{e^c(\pi(X_i), S_i(\pi), O_i)\hat{e}(\pi(X_i), S_i(\pi), O_i)}(Y_i - m^c(\pi(X_i), S_i(\pi), Z_i, |N_i|)) \\ \times \Big(e^c(\pi(X_i), S_i(\pi), O_i) - \hat{e}(\pi(X_i), S_i(\pi), O_i)\Big)\Big|\Big] \\ \leq \frac{1}{\delta^2}\mathbb{E}\Big[\sup_{\pi\in\Pi}\frac{1}{n}\sum_{i=1}^n \Big|(Y_i - m^c(\pi(X_i), S_i(\pi), Z_i, |N_i|)) \times \Big(e^c(\pi(X_i), S_i(\pi), O_i) - \hat{e}(\pi(X_i), S_i(\pi), O_i)\Big|\Big)\Big].\end{aligned} \tag{D.5}$$

Similarly,

$$\mathbb{E}\Big[\sup_{\pi\in\Pi}\Big|\frac{1}{n}\sum_{i=1}^{n}\frac{I_i(\pi)}{\hat{e}(\pi(X_i),S_i(\pi),O_i)}\Big(\hat{m}(\pi(X_i),S_i(\pi),Z_i,|N_i|)-m^c(\pi(X_i),S_i(\pi),Z_i,|N_i|\Big)\Big|\Big]$$
$$\leq \frac{1}{\delta}\mathbb{E}\Big[\sup_{\pi\in\Pi}\frac{1}{n}\sum_{i=1}^{n}\Big|\Big(\hat{m}(\pi(X_i),S_i(\pi),Z_i,|N_i|)-m^c(\pi(X_i),S_i(\pi),Z_i,|N_i|\Big)\Big|\Big].$$

Now note that for the conditional mean we have that the following expression

$$\sup_{\pi\in\Pi}\Big|\hat{m}(\pi(X_i),\sum_{k\in N_i}\pi(X_k),Z_i,|N_i|)-m^c(\pi(X_i),\sum_{k\in N_i}\pi(X_k),Z_i,|N_i|)\Big|.$$

is bounded by Holder's inequality by

$$\Big|\sum_{h=1}^{|N_i|}(m^c(\pi(X_i),h,Z_i,|N_i|)-\hat{m}(\pi(X_i),h,Z_i,|N_i|))1\{\sum_{k\in N_i}\pi(X_k)=h\}\Big|$$
$$\leq \max_{s\leq|N_i|}|m^c(\pi(X_i),s,Z_i,|N_i|)-\hat{m}(\pi(X_i),s,Z_i,|N_i|)|,$$

since $\sum_{h=1}^{|N_i|}1\{\sum_{k\in N_i}\pi(X_k)=h\}=1$. A similar bound applies to Equation (D.5). By Assumption 3.4 the proof completes. □

D.2 Proof of Theorem 3.3

The proof follows a similar strategy of Theorem D.1.

Studying the expectation Define $I_i(\pi)=1\Big\{D_i=\pi(X_i),\sum_{k\in N_i}D_k=\sum_{k\in N_i}\pi(X_k)\Big\}$. First, by Lemma 2.1 $\mathcal{K}_\Pi(n)=0$, and by the distributional assumption in Lemma 2.1 we have

$$W(\pi)=\frac{1}{|\mathcal{I}|}\sum_{j\in\mathcal{I}}\mathbb{E}\Big[r\Big(\pi(X_j),\sum_{k\in N_j}\pi(X_k),Z_j,|N_j|,\varepsilon_j\Big)\Big]=\mathbb{E}\Big[m(\pi(X_i),\sum_{k\in N_i}\pi(X_k),Z_i,|N_i|)\Big]$$

for all $i\in\{1,\cdots,E\}$. Following the same proof as in Lemma C.8 and by the definition of conditional expectation

$$\mathbb{E}\Big[m\Big(\pi(X_i),\sum_{k\in N_i}\pi(X_k),Z_i,|N_i|\Big)\Big||N_i|\leq\kappa_n\Big]P(|N_i|\leq\kappa_n)$$
$$=\mathbb{E}\Big[Y_i\frac{1\{D_i=\pi(X_i),\sum_{k\in N_i}D_k=\sum_{k\in N_i}\pi(X_k)\}}{e(\pi(X_i),S_i(\pi),Z_i,|N_i|)}1\{|N_i|\leq\kappa_n\}\Big]. \tag{D.6}$$

Bounding the regret with two main terms We can then write

$$\sup_{\pi\in\Pi} W(\pi) - W(\hat{\pi}^{tr}) \leq 2\sup_{\pi\in\Pi}\Big|W(\pi) - W_n^{tr}(\pi)\Big|$$
$$= 2\sup_{\pi\in\Pi}\Big|W(\pi) - \frac{1}{n}\sum_{i=1}^n Y_i \frac{I_i(\pi)}{e(\pi(X_i), S_i(\pi), Z_i, |N_i|)} 1\{|N_i| \leq \kappa_n\}\Big|$$
$$\leq \underbrace{2\sup_{\pi\in\Pi}\Big|\mathbb{E}\Big[m(\pi(X_i), \sum_{k\in N_i}\pi(X_k), Z_i, |N_i|)\Big||N_i| \leq \kappa_n\Big] P(|N_i| \leq \kappa_n) - \frac{1}{n}\sum_{i=1}^n Y_i \frac{I_i(\pi)1\{|N_i| \leq \kappa_n\}}{e(\pi(X_i), S_i(\pi), Z_i, |N_i|)}\Big|}_{(A)}$$
$$+ 2BP(|N_i| \geq \kappa_n),$$

where we used the fact that the outcome is uniformly bounded by B.

Bounding (A) and symmetrization The rest of the proof follows similarly to the proof of Theorem D.1. Namely, we bound the above term using the symmetrization argument as in Equation (D.2) and below. Therefore, we can show using the same argument in Theorem D.1 that

$$\mathbb{E}[(A)] \leq 2\bar{C}\mathbb{E}\Big[\sum_{m,m'\leq R_n}\sum_{l\in\mathcal{L}\times\tilde{\mathcal{L}}_m, l'\in\mathcal{L}\times\tilde{\mathcal{L}}_{m'}}\sum_{\mathcal{C}_n^M(j)\in\mathcal{C}_n^M} \underbrace{\mathbb{E}\Big[\sup_{\pi\in\Pi}\Big|\frac{1}{n}\sum_{i\in\mathcal{C}_n^M(j)}\sigma_i U_i(L_i, |N_i|, L_{k\in N_i}, \pi) I_i(l, l', m, m')\Big| |A, L, A', L'\Big]}_{(A)}\Big]$$

with the above quantities $\Big(I_i(l, l', m, m'), R_n, \mathcal{C}_n^M, A', L', (\sigma_i)_{i=1}^n, \tilde{\mathcal{L}}_m\Big)$ defined above and below Equation (D.2) and $U_i(L_i, |N_i|, L_{k\in N_i}, \pi)$ denoting random variables equal to the summands $Y_i\frac{I_i(\pi)1\{|N_i|\leq\kappa_n\}}{e(\pi(X_i), S_i(\pi), Z_i, |N_i|)}$. Such random variables are measurable with respect to $(\varepsilon_i, Q_i, Q_{k\in N_i}, \varepsilon_{D_i}, \varepsilon_{D_{k\in N_i}}, L_i, L_{k\in N_i}, |N_i|)$.

Rademacher complexity bound and conclusion We then bound the Rademacher complexity obtained from the argument as in Theorem D.1 with Lemma C.3. In particular we observe that since Y_i is bounded we have that

$$Y_i \frac{I_i(\pi)1\{|N_i| \leq \kappa_n\}}{e(\pi(X_i), S_i(\pi), Z_i, |N_i|)}$$

is a Lipschitz function of $S_i(\pi)$ with constant $2B/\delta$, since $S_i(\pi)$ takes discretely many values. The bound on (A) follows by Lemma C.3. The outer sums are bounded as discussed at the end of the proof of Theorem D.1.

D.3 Proof of Theorem 3.2

Notation Throughout the proof we define

$$\begin{aligned}
I_i(\pi) &= 1\Big\{\pi(X_i) = D_i, \sum_{k\in N_i} \pi(X_k) = \sum_{k\in N_i} D_k\Big\}, \quad \tilde{I}_i(d,s) = 1\Big\{d = D_i, s = \sum_{k\in N_i} D_k\Big\}, \\
e_i(\pi) &= e\Big(\pi(X_i), \sum_{k\in N_i} \pi(X_k), Z_{k\in N_i}, Z_i, |N_i|\Big), \quad \tilde{e}_i(d,s) = e\Big(d, s, Z_{k\in N_i}, Z_i, |N_i|\Big) \\
m_i(\pi) &= m\Big(\pi(X_i), \sum_{k\in N_i} \pi(X_k), Z_i, |N_i|\Big), \quad \tilde{m}_i(d,s) = m\Big(d, s, Z_i, |N_i|\Big).
\end{aligned} \tag{D.7}$$

We also denote $\tilde{\varepsilon}_i = Y_i - m(\pi(X_i), \sum_{k\in N_i} \pi(X_k), Z_i, |N_i|)$. Whenever we use $\hat{e}, \hat{m}$ we refer to the estimated nuisances. Recall that Z denotes the matrix of covariates for all individuals $i \in \{1, \cdots, E\}$ in the network of participants.

Preliminaries By Lemma F.3, we have

$$\begin{aligned}
\mathbb{E}\Big[\sup_{\pi\in\Pi} W(\pi) - W(\hat{\pi}^{aipw}_{\hat{m},\hat{e}})\Big] \le\ & 2\,\underbrace{\mathbb{E}\Big[\sup_{\pi\in\Pi} |W_n^{aipw}(\pi, m^c, e^c) - W(\pi)|\Big]}_{(I)} \\
& + 2\,\underbrace{\mathbb{E}\Big[\sup_{\pi\in\Pi} |W_n^{aipw}(\pi, \hat{m}, \hat{e}) - W_n(\pi, m^c, e^c)|\Big]}_{(II)}.
\end{aligned}$$

Term (I) is bounded by Theorem D.1. We now study (II).

$$\begin{aligned}
&\mathbb{E}\Big[\sup_{\pi\in\Pi} |W_n^{aipw}(\pi, \hat{m}, \hat{e}) - W_n(\pi, m^c, e^c)|\Big] \\
&= \mathbb{E}\Big[\sup_{\pi\in\Pi} \Big|\frac{1}{n}\sum_{i=1}^n \frac{I_i(\pi)}{\hat{e}_i(\pi)}(m_i(\pi) - \hat{m}_i(\pi)) + \varepsilon_i \frac{I_i(\pi)}{\hat{e}_i(\pi)} + \hat{m}_i(\pi) - m_i(\pi)\Big|\Big] \\
&= \mathbb{E}\Big[\sup_{\pi\in\Pi} \Big|\frac{1}{n}\sum_{i=1}^n \Big(\frac{I_i(\pi)}{\hat{e}_i(\pi)} - \frac{I_i(\pi)}{e_i(\pi)}\Big)(m_i(\pi) - \hat{m}_i(\pi)) + \tilde{\varepsilon}_i \frac{I_i(\pi)}{\hat{e}_i(\pi)} + \Big(\frac{I_i(\pi)}{e_i(\pi)} - 1\Big)\hat{m}_i(\pi) - m_i(\pi)\Big|\Big].
\end{aligned} \tag{D.8}$$

The last equality follows after adding and subctracting the component $\frac{I_i(\pi)}{e_i(\pi)}(m_i(\pi)-$

$\hat{m}_i(\pi)$). Observe now that we can bound the above component as follows.

$$\mathbb{E}\Big[\sup_{\pi\in\Pi}|\frac{1}{n}\sum_{i=1}^{n}\Big(\frac{I_i(\pi)}{\hat{e}_i(\pi)}-\frac{I_i(\pi)}{e_i(\pi)}\Big)(m_i(\pi)-\hat{m}_i(\pi))+\tilde{\varepsilon}_i\frac{I_i(\pi)}{\hat{e}_i(\pi)}+\Big(\frac{I_i(\pi)}{e_i(\pi)}-1\Big)\hat{m}_i(\pi)-m_i(\pi)|\Big]$$
$$\leq \underbrace{\mathbb{E}\Big[\sup_{\pi\in\Pi}|\frac{1}{n}\sum_{i=1}^{n}\Big(\frac{I_i(\pi)}{\hat{e}_i(\pi)}-\frac{I_i(\pi)}{e_i(\pi)}\Big)(m_i(\pi)-\hat{m}_i(\pi))\Big|\Big]}_{(i)}+\mathbb{E}\Big[\sup_{\pi\in\Pi}\Big|\frac{1}{n}\sum_{i=1}^{n}\underbrace{\tilde{\varepsilon}_i\Big(\frac{I_i(\pi)}{\hat{e}_i(\pi)}-\frac{I_i(\pi)}{e_i(\pi)}\Big)}_{(ii)}\Big|\Big]$$
$$+\mathbb{E}\Big[\sup_{\pi\in\Pi}\Big|\frac{1}{n}\sum_{i=1}^{n}\underbrace{\tilde{\varepsilon}_i\frac{I_i(\pi)}{e_i(\pi)}}_{(iii)}\Big|\Big]+\mathbb{E}\Big[\sup_{\pi\in\Pi}\Big|\frac{1}{n}\sum_{i=1}^{n}\underbrace{\Big(\frac{I_i(\pi)}{e_i(\pi)}-1\Big)(\hat{m}_i(\pi)-m_i(\pi))}_{(iv)}\Big|\Big]. \tag{D.9}$$

The first step consists in showing that the terms (ii), (iii), (iv) are centered around zero. (iii) directly follows since none of the nuisance functions is estimated. We therefore discuss the expectation of (ii) and (iv).

Expectation of (ii) We start from (ii). By Algorithm 1, we claim that

$$\mathbb{E}\Big[\tilde{\varepsilon}_i\Big(\frac{I_i(\pi)}{\hat{e}_i(\pi)}-\frac{I_i(\pi)}{e_i(\pi)}\Big)\Big]=0. \tag{D.10}$$

This follows from the following argument: using the law of iterated expectations

$$E[(ii)]=\mathbb{E}\Big[\mathbb{E}\Big[\Big(r(\pi(X_i),\sum_{k\in N_i}\pi(X_k),Z_i,|N_i|,\varepsilon_i)-m_i(\pi)\Big)\Big(\frac{I_i(\pi)}{\hat{e}_i(\pi)}-\frac{I_i(\pi)}{e_i(\pi)}\Big)\Big|\hat{e}_i,Z,A\Big]\Big]. \tag{D.11}$$

Observe now that $\hat{e}_i(\pi)$ is estimated on an independent sample from ε_i, conditional on (A,Z). Under Assumption 2.5 and the first condition in Assumption 2.2[52],

$$\begin{aligned}&\mathbb{E}\Big[r(\pi(X_i),\sum_{k\in N_i}\pi(X_k),Z_i,|N_i|,\varepsilon_i)\Big(\frac{I_i(\pi)}{\hat{e}_i(\pi)}-\frac{I_i(\pi)}{e_i(\pi)}\Big)\Big|\hat{e}_i,Z,A\Big]=\\&m(\pi(X_i),\sum_{k\in N_i}\pi(X_k),Z_i,|N_i|)\mathbb{E}\Big[\frac{I_i(\pi)}{\hat{e}_i(\pi)}-\frac{I_i(\pi)}{e_i(\pi)}\Big|\hat{e}_i,Z,A\Big].\end{aligned} \tag{D.12}$$

Expectation of (iii) Notice now that the same reasoning for (ii) directly applies to (iii) for the fact that

$$\mathbb{E}\Big[\frac{I_i(\pi)}{e_i(\pi)}-1|\hat{m}_i,A,Z\Big]=\mathbb{E}\Big[\frac{I_i(\pi)}{e_i(\pi)}-1|A,Z\Big]=0$$

[52]Namely ε_{D_i} are *i.i.d.* and exogenous.

by the cross fitting argument.

Concentration bound for $(ii), (iii), (iv)$ We can therefore apply the same argument used for Theorem D.1 also to these three components directly (see the argument on Equation (D.2) and below) and obtain a bound of order $\mathcal{O}\Big(\mathbb{E}[\mathcal{P}_{M,|\mathcal{L}|}(\mathcal{N}_E)]\sqrt{\mathrm{VC}(\Pi)/n}\Big)$. Here, Lipschitz properties are guaranteed by the second condition of Assumption 3.3 for which the estimated $\hat{m}$ and $1/\hat{e}$ are bounded .

Bounding (i) We are now left to bound (i). Observe that we have

$$\begin{aligned}
(i) &\leq \mathbb{E}\Big[\sqrt{\frac{1}{n}\sum_{i=1}^n \sup_{d,s}\Big(\frac{1}{\tilde{e}_i(d,s)} - \frac{1}{\hat{\tilde{e}}_i(d,s)}\Big)^2}\sqrt{\frac{1}{n}\sum_{i=1}^n \sup_{d,s}\Big(\tilde{m}_i(d,s) - \hat{\tilde{m}}_i(d,s)\Big)^2}\Big] \\
&\leq \sqrt{\mathbb{E}\Big[\frac{1}{n}\sum_{i=1}^n \sup_{d,s}\Big(\frac{1}{\tilde{e}_i(d,s)} - \frac{1}{\hat{\tilde{e}}_i(d,s)}\Big)^2\Big]}\sqrt{\mathbb{E}\Big[\frac{1}{n}\sum_{i=1}^n \sup_{d,s}\Big(\tilde{m}_i(d,s) - \hat{\tilde{m}}_i(d,s)\Big)^2\Big]} \leq \bar{C}n^{1/2},
\end{aligned} \tag{D.13}$$

for a finite constant $\bar{C} < \infty$ by Assumption 3.3.

D.4 Proof of Theorem 3.4

As in the above theorem, it suffices to bound $\mathbb{E}\Big[\sup_{\pi\in\Pi}|W(\pi) - W_n(\pi)|\Big]$.

Expectation We can write

$$\mathbb{E}\Big[\sup_{\pi\in\Pi}|W(\pi) - W_n(\pi)|\Big] \leq \mathcal{K}_\Pi(n) + \mathbb{E}\Big[\sup_{\pi\in\Pi}|W^*(\pi) - W_n(\pi)|\Big]$$

where

$$W^*(\pi) = \frac{1}{n}\sum_{i=1}^n \mathbb{E}\Big[r\Big(\pi(X_i), S_i(\pi), Z_i, |N_i|, \varepsilon_i\Big)\Big] = \frac{1}{n}\sum_{i=1}^n \mathbb{E}\Big[m\Big(\pi(X_i), S_i(\pi), Z_i, |N_i|\Big)\Big],$$

and where the second equality follows from Assumption 2.1, 2.2 and 2.4 similarly to what discussed in the proof of Lemma C.9. By Lemma C.9, we have that $W^*(\pi) = \mathbb{E}[W_n(\pi, m^c, e^c)]$ under correct specification of either m^c or e^c. Therefore for $i \in \{1, \cdots, n\}$

$$\begin{aligned}
W^*(\pi) =& \frac{1}{n}\sum_{i=1}^n \mathbb{E}\Big[\frac{1\{\pi(X_i) = D_i, \sum_{k\in N_i}\pi(X_k) = \sum_{k\in N_i} D_k\}}{e^c\Big(\pi(X_i), S_i(\pi), O_i\Big)}\Big(Y_i - m^c\Big(\pi(X_i), S_i(\pi), Z_i, |N_i|\Big)\Big)\Big] + \\
&+ \mathbb{E}\Big[m^c\Big(\pi(X_i), S_i(\pi), Z_i, |N_i|\Big)\Big],
\end{aligned}$$

where $O_i = \Big(Z_{k\in N_i}, Z_i, |N_i|\Big)$.

Bound in terms of K We can then decompose

$$\mathbb{E}\Big[\sup_{\pi\in\Pi}|W^*(\pi) - W_n(\pi)|\Big] \leq \sum_{k=1}^K \mathbb{E}\Big[\sup_{\pi\in\Pi}|W^*(\pi) - W_n^k(\pi)|\Big]$$

where $W_n^k(\pi)$ is the empirical welfare in a group k of independent observations.

Symmetrization and conclusion Using Lemma F.1 we can use the standard symmetrization argument to $\mathbb{E}\Big[\sup_{\pi\in\Pi}|W^*(\pi) - W_n^k(\pi)|\Big]$, since each summand in $W_n^k(\pi)$ is independent and identically distributed. The proof concludes invoking Lemma C.7 and by Jensen's inequality to bound $\mathbb{E}[\sqrt{\mathcal{N}_n\log(\mathcal{N}_n)}] \leq \sqrt{\mathbb{E}[\mathcal{N}_n\log(\mathcal{N}_n)]}$.

D.5 Proof of Theorem 5.2

By the law of iterated expectations we obtain that if either m^c or e^c are correctly specified, and under Assumption 2.3

$$\mathbb{E}[W_n^t(\pi, m^c, e^c)] = W(\pi) \tag{D.14}$$

similarly to what discussed in Lemma C.9. By trivial rearrengement, we obtain that

$$\sup_{\pi\in\Pi} W(\pi) - W(\hat{\pi}^{aipw}) \leq 2\sup_{\pi\in\Pi}\Big|W_n^t(\pi, m^c, e^c) - W(\pi)\Big|. \tag{D.15}$$

Notice now that the exact same argument for bounding the above term follows from the proof of Theorem D.1 holds, whereas in such a case the moment conditions also depend on $\bar{\rho}$. In particular the moment bounds are multiplied by $\bar{\rho}^3$. Similarly, the Lipschitz constant of the conditional mean function also must be multiplied by the term $\bar{\rho}$.

Appendix E Proof of Theorem 5.1

First equality First we want to show that (note that $\Big\{D_i = \pi(X_i)\forall i\in\{1,\cdots,E\}\Big\}$ denotes the event that $D_i = \pi(X_i)$ but do not condition on the values of D_i)

$$\mathbb{E}\Big[r\Big(T_i, \sum_{k\in N_i} T_k, Z_i, |N_i|, \varepsilon_i\Big)\Big|\Big\{D_i = \pi(X_i)\forall i\in\{1,\cdots,E\}\Big\}\Big] = \mathbb{E}\Big[r\Big(T_i(\pi), \sum_{k\in N_i} T_k(\pi), Z_i, |N_i|, \varepsilon_i\Big)\Big].$$

The proof of the claim follows from the assumption that $(\varepsilon_j, \nu_j)_{j=1}^E \perp \Big(Z, A, (\varepsilon_{D_j})_{j=1}^E\Big)$. In particular, we can write

$$
\begin{aligned}
&\mathbb{E}\Big[r\Big(T_i, \sum_{k \in N_i} T_k, Z_i, |N_i|, \varepsilon_i\Big)\Big|\Big\{D_i = \pi(X_i) \forall i \in \{1, \cdots, E\}\Big\}\Big] \\
&= \mathbb{E}\Big[r\Big(T_i(\pi), \sum_{k \in N_i} T_k(\pi), Z_i, |N_i|, \varepsilon_i\Big)\Big|\Big\{D_i = \pi(X_i) \forall i \in \{1, \cdots, E\}\Big\}\Big] \\
&= \mathbb{E}_{A,Z}\Big[\mathbb{E}\Big[r\Big(T_i(\pi), \sum_{k \in N_i} T_k(\pi), Z_i, |N_i|, \varepsilon_i\Big)\Big|\Big\{D_i = \pi(X_i) \forall i \in \{1, \cdots, E\}\Big\}, A, Z\Big]\Big] \\
&= \mathbb{E}_{A,Z}\Big[\mathbb{E}\Big[r\Big(T_i(\pi), \sum_{k \in N_i} T_k(\pi), Z_i, |N_i|, \varepsilon_i\Big)\Big|A, Z\Big]\Big] \\
&= \mathbb{E}\Big[r\Big(T_i(\pi), \sum_{k \in N_i} T_k(\pi), Z_i, |N_i|, \varepsilon_i\Big)\Big].
\end{aligned}
$$

Second equality Second, we want to derive the main decomposition in the second equality. Let denote $\mathbb{E}_\pi$ the expectation conditional on the event $\Big\{D_i = \pi(X_i) \forall i \in \{1, \cdots, E\}\Big\}$. Using the law of iterated expectations and the equality above, we have

$$
\begin{aligned}
&\mathbb{E}_\pi\Big[r(T_i, \sum_{k \in N_i} T_k, Z_i, |N_i|, \varepsilon_i)\Big] \\
&= \mathbb{E}\Big[\mathbb{E}\Big[r(T_i(\pi), \sum_{k \in N_i} T_k(\pi), Z_i, |N_i|, \varepsilon_i)\Big|Z, A\Big]\Big].
\end{aligned}
$$

Using the definition of the expectation, we have

$$
(I) = \underbrace{\sum_{s \in \{0, \cdots, |N_i|\}} \mathbb{E}\Big[r(d, s, Z_i, |N_i|, \varepsilon_i)\Big|T_i(\pi) = d, \sum_{k \in N_i} T_k(\pi) = s, Z, A\Big]}_{(i)} \times \underbrace{P\Big(T_i(\pi) = d, \sum_{k \in N_i} T_k(\pi) = s\Big|A, Z\Big)}_{(ii)}.
$$

We now discuss (i) and (ii) separately. First, we observe that since $(\varepsilon_j)_{j=1}^E \perp \Big(Z, A, (\varepsilon_{D_j}, \nu_j)_{j=1}^E\Big)$

$$
\begin{aligned}
(i) &= \mathbb{E}\Big[r(d, s, Z_i, |N_i|, \varepsilon_i)\Big|Z_i, |N_i|, T_i(\pi) = d, \sum_{k \in N_i} T_k(\pi) = s, Z, A\Big] \\
&= \mathbb{E}_\pi\Big[r(d, s, Z_i, |N_i|, \varepsilon_i)\Big|Z_i, |N_i|, T_i = d, \sum_{k \in N_i} T_k = s, Z, A\Big] \\
&= \mathbb{E}\Big[r(d, s, Z_i, |N_i|, \varepsilon_i)\Big|Z_i, |N_i|, T_i = d, \sum_{k \in N_i} T_k = s, Z, A\Big] \\
&= \mathbb{E}\Big[r(d, s, Z_i, |N_i|, \varepsilon_i)\Big|Z_i, |N_i|, T_i = d, \sum_{k \in N_i} T_k = s\Big].
\end{aligned}
$$

Consider now (ii). Observe that by indepedence and exogeneity of $(\nu_j)_{j=1}^E$, we can write

$$(ii) = P\Big(T_i(\pi) = d\Big|A, Z\Big) \times \sum_{u_1, \cdots, u_l : \sum_v u_v = s} \prod_{k=1}^{|N_i|} P\Big(T_{N_i^{(k)}}(\pi) = u_k\Big|A, Z\Big)$$

where the expression sums over all possible combinations of selected treatments such that the sum of the selected treatments of the neighbors is s. Using again exogeneity of ν_i, we have

$$P\Big(T_i(\pi) = d\Big|A, Z\Big) = P\Big(T_i = d\Big|Z_i, |N_i|, D_i = \pi(X_i), \sum_{k \in N_i} D_k = \sum_{k \in N_i} \pi(X_k), \sum_{k \in N_i} \pi(X_k), \pi(X_i)\Big).$$

Similar reasoning also applies to neighbors' selected treatments, which depend on the assigned treatments to the second-degree neighbors. The proof completes.

Appendix F Preliminary lemmas

This section collects a first set of lemmas from past literature that we invoke in our proofs.

Lemma F.1. *(Van Der Vaart and Wellner (1996)) Let $\sigma_1, ..., \sigma_n$ be Rademacher sequence independent of $X_1, ..., X_n$. Suppose that $X_1, \cdots, X_n$ are i.i.d.. Then*

$$\mathbb{E}\Big[sup_{f \in \mathcal{F}}\Big|\sum_{i=1}^n f(X_i) - \mathbb{E}[f(X_i)]\Big|\Big] \leq 2\mathbb{E}\Big[sup_{f \in \mathcal{F}}\Big|\sum_{i=1}^n \sigma_i f(X_i)\Big|\Big].$$

Lemma F.2. *(Brook's Theorem, Brooks (1941)) For any connected undirected graph G with maximum degree Δ, the chromatic number of G is at most Δ unless G is a complete graph or an odd cycle case the chromatic number is $\Delta + 1$.*

The next lemma discuss the case of estimated conditional mean and balancing score function. The lemma follows from Kitagawa and Tetenov (2018).

Lemma F.3. *(From Kitagawa and Tetenov (2018)) The following holds.*

$$\sup_{\pi \in \Pi} W(\pi) - W(\hat{\pi}^{aipw}_{\hat{m}, \hat{e}}) \leq 2 \sup_{\pi \in \Pi} |W_n^{aipw}(\pi, m^c, e^c) - W(\pi)| + 2 \sup_{\pi \in \Pi} |\hat{W}_n^{aipw}(\pi, \hat{m}, \hat{e}) - W_n(\pi, m^c, e^c)|.$$

Lemma F.4. *(Lemma A.1, Kitagawa and Tetenov (2019)) Let $t_0 > 1$, define*

$$g(t) := \begin{cases} 0, & \text{for } t = 0 \\ t^{-1/2}, t \geq 1 \end{cases}, \quad h(t) = t_0^{-1/2} - \frac{1}{2}t_0^{-3/2}(t - t_0) + t_0^{-2}(t - t_0).$$

Then $g(t) \leq h(t)$, for $t = 0$ and all $t \geq 1$.

In the next two lemmas we formalize two key properties of covering numbers.

Lemma F.5. *(Theorem 29.6, Devroye et al. (2013)) Let $\mathcal{F}_1, ..., \mathcal{F}_k$ be classes of real functions on $\mathbb{R}^d$. For n arbitrary points $z_1^n = (z_1, ..., z_n)$ in $\mathbb{R}^d$, define the sets $\mathcal{F}_1(z_1^n), ..., \mathcal{F}_k(z_1^n)$ in $\mathbb{R}^n$ by*

$$\mathcal{F}_j(z_1^n) = \{f_j(z_1), ..., f_j(z_n) : f_j \in \mathcal{F}_j\}, \quad j = 1, ..., k.$$

Also, introduce

$$\mathcal{F} = \{f_1 + ... + f_k; f_j \in \mathcal{F}_j, \quad j = 1, ..., k\}.$$

Then for every $\varepsilon > 0$ and z_1^n,

$$\mathcal{N}_1\Big(\varepsilon, \mathcal{F}(z_1^n)\Big) \leq \prod_{j=1}^{k} \mathcal{N}_1\Big(\varepsilon/k, \mathcal{F}_j(z_1^n)\Big)$$

Lemma F.6. *(Pollard, 1990) Let $\mathcal{F}$ and $\mathcal{G}$ be classes of real valued functions on $\mathbb{R}^d$ bounded by M_1 and M_2 respectively. For arbitrary fixed points z_1^n in $\mathbb{R}^d$, let*

$$\mathcal{J}(z_1^n) = \{(h(z_1), ..., h(z_n); h \in \mathcal{J}\}, \quad \mathcal{J} = \{fg; f \in \mathcal{F}, g \in \mathcal{G}\}.$$

Then for every $\varepsilon > 0$ and z_1^n,

$$\mathcal{N}_1\Big(\varepsilon, \mathcal{J}(z_1^n)\Big) \leq \mathcal{N}_1\Big(\frac{\varepsilon}{2M_2}, \mathcal{F}(z_1^n)\Big)\mathcal{N}_1\Big(\frac{\varepsilon}{2M_1}, \mathcal{G}(z_1^n)\Big).$$

Lemma F.7. *(Wenocur and Dudley, 1981) Let $g : \mathbb{R}^d \mapsto \mathbb{R}$ be an arbitrary function and consider the class of functions $\mathcal{G} = \{g + f, f \in \mathcal{F}\}$. Then*

$$VC(\mathcal{G}) = VC(\mathcal{F})$$

where $VC(\mathcal{F})$, $VC(\mathcal{G})$ denotes the VC dimension respectively of $\mathcal{F}$ and $\mathcal{G}$.

An important relation between covering numbers and Rademacher complexity is given by Dudley's entropy integral bound.

Lemma F.8. *(From Theorem 5.22 in Wainwright (2019)) For a function class $\mathcal{F}$ of uniformly bounded functions and arbitrary fixed points $z_1^n \in \mathbb{R}^d$, and i.i.d. Rademacher random variables $\sigma_1, ..., \sigma_n$ in $\mathbb{R}$,*

$$\mathbb{E}_\sigma\Big[sup_{f\in\mathcal{F}}\Big|\frac{1}{n}\sum_{i=1}^n \sigma_i f(z_i)\Big|\Big] \leq \frac{32}{\sqrt{n}}\int_0^{D_q}\sqrt{\log\Big(\mathcal{N}_q\Big(u, \mathcal{F}(z_1^n)\Big)\Big)}du,$$

where D_q denotes the maximum diameter of $\mathcal{F}(z_1^n)$ according to the metric $d_q(f,g) = \Big(\frac{1}{n}\sum_{i=1}^n (f(z_i) - g(z_i))^q\Big)^{1/q}$.

Theorem 5.22 in Wainwright (2019) provides a general version of Lemma F.8. Equivalent versions of Lemma F.8 can be found also in Van Der Vaart and Wellner (1996).

Lemma F.9. *Let $X_1, ..., X_n$ be independent Bernoulli random variables with $b_i \sim Bern(p_i)$. Let $\bar{p} = \frac{1}{n}\sum_{i=1}^n p_i$ with $n\bar{p} > 1$ and $g(.)$ as defined in the Lemma F.4. Then*

$$\mathbb{E}\Big[g\Big(\sum_{i=1}^n X_i\Big)\Big] < 2(n\bar{p})^{-1/2}$$

Proof. The proof follows similarly as in Kitagawa and Tetenov (2019). Let $h(.)$ be the function defined in Lemma F.4 with $t_0 = n\bar{p}$. By Lemma F.4,

$$\begin{aligned}
\mathbb{E}\Big[g\Big(\sum_{i=1}^n X_i\Big)\Big] &\le \mathbb{E}\Big[h\Big(\sum_{i=1}^n X_i\Big)\Big] \\
&= \mathbb{E}\Big[(n\bar{p})^{-1/2} - \frac{1}{2}\sqrt{(n\bar{p})^3}\Big(\sum_{i=1}^n X_i - \sum_{i=1}^n \mathbb{E}[X_i]\Big) + (n\bar{p})^{-2}\Big(\sum_{i=1}^n X_i - \sum_{i=1}^n \mathbb{E}[X_i]\Big)^2\Big] \\
&= (n\bar{p})^{-1/2} - \frac{1}{2}(n\bar{p})^{-3/2}0 + (n\bar{p})^{-2}Var\Big(\sum_{i=1}^n X_i\Big) \\
&= (n\bar{p})^{-1/2} + (n\bar{p})^{-2}\sum_{i=1}^n p_i(1-p_i) \\
&\le (n\bar{p})^{-1/2} + (n\bar{p})^{-2}\sum_{i=1}^n p_i \\
&= (n\bar{p})^{-1/2} + (n\bar{p})^{-2} n\bar{p} \\
&= (n\bar{p})^{-1/2} + (n\bar{p})^{-1} < 2(n\bar{p})^{-1/2}
\end{aligned}$$

since $n\bar{p} > 1$ and $(n\bar{p})^{-1} < (n\bar{p})^{-1/2}$.

□

We conclude our discussion with an extension under lack of point-wise measurability.

Lemma F.10. *(Lemma 2.6, Hajłasz and Malỳ (2002)) Let $\mathcal{U}$ be a class of measurable functions defined on a measurable space $E \subset \mathbb{R}^n$. Then the lattice supremum defined as $\bigvee \mathcal{U}$ exists and there is a countable sub-family $\mathcal{V} \subset \mathcal{U}$ such that*

$$\bigvee \mathcal{U} = \bigvee \mathcal{V} = \sup \mathcal{V}. \tag{F.1}$$

The above lemma has one important implication: whenever the function Π is not point-wise measurable, by taking the lattice supremum over a function class, we are guaranteed that such supremum corresponds to a supremum over a countable subset, for which the pointwise supremum is defined. By construction, such class $\mathcal{V}$ has VC-dimension bounded by the VC-dimension of the original function class $\mathcal{U}$.

Appendix G Additional Extensions

G.1 Higher Order Interference

In this section, we discuss the case of two-degree interference. The generalization to K-degree interference follows similarly to what is discussed in this section. For a given A, we denote the set of second degree neighbors of individual i as $N_{i,2}(A)$. The set $N_{i,2}(A)$ contains each element $j \neq i$ such that there exists a connecting path of length two between j and i, repeated by the number of such paths.

Formally, for a given adjacency matrix A we let $N_{i,2}(A) = \{j^{|k_{j,i}(A)|} : k_{j,i}(A) = \{k \notin \{j\} \cup \{i\} : A^{(i,k)}A^{(k,j)} \neq 0\}\}$, the set containing second degree neighbors, where each element j has multiplicity $|k_{j,i}(A)|$. We will refer to the set of $N_i, N_{i,2}$ by omitting their dependence on the adjacency matrix A whenever possible. We will assume that the first and second-degree neighbors are observable. We also assume that researchers observe the treatment assignment D_i of each unit, covariates of individuals, Z_i, that may also contain sufficient statistics of first and second-degree neighbors' statistics. We let the policy function depend on characteristics $X_i \subseteq (Z_i, |N_i|, |N_{i,2}|), X_i \in \mathcal{X}$.

Assumption G.1. (Two-Degrees Interference Model) For all unit i, there exist a function $r(.)$ such that

$$Y_i = r\Big(D_i, \sum_{k \in N_i} D_k, \sum_{k \in N_{i,2}} D_k, Z_i, |N_i|, |N_{i,2}|, \varepsilon_i\Big)$$

almost surely, where $N_{i,2} = \{j^{|k_{j,i}(A)|} : k_{j,i}(A) = \{k \notin \{j\} \cup \{i\} : A^{(i,k)}A^{(k,j)} \neq 0\}\}$.

Assumption G.1 generalizes Assumption 2.1 to two-degrees interference. It assumes each unit's potential outcome only depends on its treatment status, the number of first degree neighbors, the number of second-degree neighbors, the number of treated first-degree neighbors, and the number of treated second-degree neighbors.

Assumption G.2. Let the following holds.

(ID) For all $(i, j) \in \{1, ..., n\}^2$,

$$P(\varepsilon_i \leq x | Z_i = z, |N_i| = l, |N_{i,2}| = l_2) = P(\varepsilon_j \leq x | Z_j = z, |N_j| = l, |N_{j,2}| = l_2)$$

for all $z \in \mathcal{Z}, l, l_2 \in \mathbb{Z}, x \in \mathbb{R}$.

(UN1) For all $i \in \{1, ..., n\}$, $D_i = f_D(Z_i, \varepsilon_{D_i})$ where ε_{D_i} are *i.i.d.* and independent of observables and unobservables.

(UN2) For all $i \in \{1, ..., n\}$, $\varepsilon_i \perp \Big(N_i, Z_{j \in N_i}, N_{i,2}, Z_{j \in N_{i,2}}, N_{i \in N_i}, N_{i \in N_{i,2}}\Big)\Big| Z_i, |N_i|, |N_{i,2}|$.

Based on the previous assumption, we can now formalize the definition of the conditional mean function and the balancing score.

Definition G.1. (Conditional Mean) Under Assumption G.1 and Assumption G.2 the conditional mean function under two degree interference is

$$m_2(d, s_1, s_2, z, l, l_2) = \mathbb{E}\Big[r\Big(d, s_1, s_2, Z_i, |N_i| = l, |N_{i,2}| = l_2, \varepsilon_i\Big)\Big| Z_i = z\Big]$$

as a function of (d, s_1, s_2, z, l, l_2) only.

The second causal estimand of interest is the balancing score.

Definition G.2. (balancing score) Under Assumption G.2 the balancing score is defined as

$$e_2(d, s_1, s_2, \mathbf{x}_1, \mathbf{x}_2, z, l, l_2) = P\Big(D_i = d, \sum_{k \in N_i} D_k = s_1, \sum_{k \in N_{i,2}} D_k = s_2 \Big| \mathcal{V}_i(\mathbf{x}_1, \mathbf{x}_2, z, l, l_2)\Big)$$

$$\mathcal{V}_i(\mathbf{x}_1, \mathbf{x}_2, z, l, l_2) = \{Z_{k \in N_i} = \mathbf{x}_1, Z_{k \in N_{i,2}} = \mathbf{x}_2, Z_i = z, |N_i| = l, |N_{i,2}| = l_2\}$$

Assumption G.2 guarantees that the balancing score is not unit-specific. The expression above can be decomposed as a sum of products of marginal conditional probabilities, similarly to what discussed in Equation (9) under one degree interference.[53]

[53] Namely,

$$\begin{aligned} e_2(d, s_1, s_2, \mathbf{x}_1, \mathbf{x}_2, z, l, l_2) = P\Big(D_i = d | Z_i = z\Big) \times \\ \Big(\sum_{u_1, \ldots, u_{|N_i|} : \sum_v u_v = s_1} \prod_{k=1}^{l} P\Big(D_{N_i^{(k)}} = u_k \Big| Z_{N_i^{(k)}} = \mathbf{x}_1^{k,\cdot}\Big)\Big) \\ \times \Big(\sum_{u_1, \ldots, u_{|N_{i,2}|} : \sum_v u_v = s_2} \prod_{k=1}^{l_2} P\Big(D_{N_{i,2}^{(k)}} = u_k \Big| Z_{N_{i,2}^{(k)}} = \mathbf{x}_2^{k,\cdot}\Big)\Big), \end{aligned} \quad \text{(G.1)}$$

which can be estimated by plugging-in the conditional marginal probabilities obtained via maximum likelihood.

Estimation can be performed similarly to what discussed in Section 4 under one degree interference. For instance the estimated welfare obtained by using the conditional mean function only is defined as

$$\hat{W}_{n,2}(\pi) = \frac{1}{n}\sum_{i=1}^{n} \hat{m}_2\Big(\pi(X_i), \sum_{k\in N_i} \pi(X_k), \sum_{k\in N_{i,2}} \pi(X_k), Z_i, |N_i|, |N_{i,2}|\Big). \tag{G.2}$$

G.2 Capacity Constraints

In this section, we discuss capacity constraints using concentration arguments, and we quantify such an error in the presence of locally dependent random variables. Further discussion on capacity constraints for *i.i.d.* data can be found in Kitagawa and Tetenov (2018), Bhattacharya and Dupas (2012) whereas here we differently account for the network structure. To achieve this goal, we impose conditions on the dependence of the random variables used for the targeting exercise.

Assumption G.3. (Local Dependence 2) *Let $\{W_i, i \in \{1, ..., \infty\}\}$ denote an arbitrary set of possibly unobserved independent random variables. For each $i \in \{1, ..., n\}$ let $S_i \subset \{1, ..., \infty\}$ denote an arbitrary set of such random variables. Suppose that we can write $X_i : W_{S_i} \mapsto \mathcal{X}$ depending on some arbitrary W_{S_i}. Suppose in addition that $\max_{i\in\{1,...,n\}} |j : S_i \cap S_j \neq \emptyset| \leq m < \infty$, where $|j : S_i \cap S_j \neq \emptyset|$ denotes the cardinality of the set of dependent X_i.*

Under Assumption G.3, we show that the constraint equation concentrates uniformly over all possible policies around its expectation at an exponential rate. Such a result guarantees control on the number of treated units, also once the policy is scaled at the population level.

Proposition G.1. *Suppose that Assumption G.3 holds and $X_i \sim F_X$. Assume that Π has finite VC-dimension. Then there exists a universal constant $\bar{C}$ such that with probability at least $1 - \gamma$,*

$$\sup_{\pi\in\Pi} \Big|\frac{1}{n}\sum_{i=1}^{n} \pi(X_i) - \int_{x\in\mathcal{X}} \pi(x) dF_X(x)\Big| \leq m^2\sqrt{\mathrm{VC}(\Pi)/n} + 2m\sqrt{\log(2/\gamma)/n} \tag{G.3}$$

Proposition G.1 guarantees non-asymptotic probability bound on the in-sample capacity constraint. Intuitively, it shows that the sample analog constraint converges to the target constraint at the optimal rate, for the bounded degree, uniformly on the function class, also under local dependence.

Proof. First, it is easy to show that $(X_i)_{i=1}^n$ form a dependency graph of degree m^2. By Lemma F.1 we can bound the expectation as

$$\mathbb{E}\Big[\sup_{\pi \in \Pi}\Big|\frac{1}{n}\sum_{i=1}^n \pi(X_i) - \mathbb{E}[\pi(X_i)]\Big|\Big] \leq 2 \sum_{\mathcal{C}_n^2(j) \in \mathcal{C}} \mathbb{E}\Big[\sup_{\pi \in \Pi}\Big|\frac{1}{n}\sum_{i \in \mathcal{C}_n^2(j)} \sigma_i \pi(X_i)\Big|\Big]. \tag{G.4}$$

By finiteness of the VC dimension, and by the fact that the maximum diameter of Π is two, we have as discussed in Chapter 5 in Wainwright (2019), that for a universal constant $\bar{C}$

$$\int_0^2 \sqrt{\log(\mathcal{N}_2(u, \Pi))} du \leq \bar{C}\sqrt{VC(\Pi)}. \tag{G.5}$$

Therefore by Lemma F.8

$$\mathbb{E}\Big[\sup_{\pi \in \Pi}\Big|\frac{1}{n}\sum_{i=1}^n \sigma_i \pi(X_i)\Big|\Big] \leq \bar{C}'\sqrt{VC(\Pi)/n}. \tag{G.6}$$

for a universal constant $\bar{C}'$. We are left to show concentration of the left-hand side in Equation (G.3) around its expectation. Notice now that

$$\frac{1}{n}\sum_{i=1}^n \pi(X_i) - \int \pi(x) dF_x(x) \tag{G.7}$$

satisfies the bounded difference inequality with constant $2/n$. In addition, by Assumption G.3, units X_i satisfy the definition of (HD,m) discussed in Definition 3 in Paulin (2012) for finite m. In addition by Lemma 1.1 in Paulin (2012), the random variables are also (HD,m,m). By Corollary 2.1 in Paulin (2012), we have that

$$P\Big(\Big|\sup_{\pi \in \Pi}\Big|\frac{1}{n}\sum_{i=1}^n \pi(X_i) - \int \pi(x) dF(x)\Big| - \mathbb{E}\Big[\sup_{\pi \in \Pi}\Big|\frac{1}{n}\sum_{i=1}^n \pi(X_i) - \int \pi(x) dF(x)\Big|\Big]\Big| > t\Big) \\ \leq 2\exp\Big(-\frac{nt^2}{m^2 2}\Big). \tag{G.8}$$

by setting $\gamma = 2\exp\Big(-\frac{nt^2}{m^2 2}\Big)$ and $t = m\sqrt{\log(2/\gamma)/n}$, we obtain that with probability $1-\gamma$,

$$\Big|\sup_{\pi\in\Pi}\Big|\frac{1}{n}\sum_{i=1}^{n}\pi(X_i)-\int\pi(x)dF(x)\Big|-\mathbb{E}\Big[\sup_{\pi\in\Pi}\Big|\frac{1}{n}\sum_{i=1}^{n}\pi(X_i)-\int\pi(x)dF(x)\Big|\Big]\Big|\leq m\sqrt{\log(2/\gamma)/n}. \tag{G.9}$$

Therefore, using the bound on the expectation derived above, we have

$$\sup_{\pi\in\Pi}\Big|\frac{1}{n}\sum_{i=1}^{n}\pi(X_i)-\int\pi(x)dF(x)\Big|\leq\bar{C}m^2\sqrt{VC(\Pi)/n}+m\sqrt{\log(2/\gamma)/n}. \tag{G.10}$$

□

References

Akbarpour, M., S. Malladi, and A. Saberi (2018). Just a few seeds more: value of network information for diffusion. *Available at SSRN 3062830*.

Alidaee, H., E. Auerbach, and M. P. Leung (2020). Recovering network structure from aggregated relational data using penalized regression. *arXiv preprint arXiv:2001.06052*.

Ananth, A. (2021). Optimal treatment assignment rules on networked populations. *Working paper*.

Armstrong, T. and S. Shen (2015). Inference on optimal treatment assignments. *Available at SSRN 2592479*.

Aronow, P. M., C. Samii, et al. (2017). Estimating average causal effects under general interference, with application to a social network experiment. *The Annals of Applied Statistics 11*(4), 1912–1947.

Athey, S., D. Eckles, and G. W. Imbens (2018). Exact p-values for network interference. *Journal of the American Statistical Association 113*(521), 230–240.

Athey, S. and S. Wager (2021). Policy learning with observational data. *Econometrica 89*(1), 133–161.

Auerbach, E. (2019). Identification and estimation of a partially linear regression model using network data. *arXiv preprint arXiv:1903.09679*.

Banerjee, A., A. G. Chandrasekhar, E. Duflo, and M. O. Jackson (2013). The diffusion of microfinance. *Science 341*(6144), 1236498.

Banerjee, A., A. G. Chandrasekhar, E. Duflo, and M. O. Jackson (2014). Gossip: Identifying central individuals in a social network. Technical report, National Bureau of Economic Research.

Banerjee, A., A. G. Chandrasekhar, E. Duflo, and M. O. Jackson (2019). Using gossips to spread information: Theory and evidence from two randomized controlled trials. *The Review of Economic Studies 86*(6), 2453–2490.

Barrera-Osorio, F., M. Bertrand, L. L. Linden, and F. Perez-Calle (2011). Improving the design of conditional transfer programs: Evidence from a randomized education experiment in colombia. *American Economic Journal: Applied Economics 3*(2), 167–95.

Bhattacharya, D. (2009). Inferring optimal peer assignment from experimental data. *Journal of the American Statistical Association 104*(486), 486–500.

Bhattacharya, D. and P. Dupas (2012). Inferring welfare maximizing treatment assignment under budget constraints. *Journal of Econometrics 167*(1), 168–196.

Bhattacharya, D., P. Dupas, and S. Kanaya (2019). Demand and welfare analysis in discrete choice models with social interactions. *Available at SSRN 3116716*.

Bloch, F., M. O. Jackson, and P. Tebaldi (2017). Centrality measures in networks. *Available at SSRN 2749124*.

Bond, R. M., C. J. Fariss, J. J. Jones, A. D. Kramer, C. Marlow, J. E. Settle, and J. H. Fowler (2012). A 61-million-person experiment in social influence and political mobilization. *Nature 489*(7415), 295.

Breiman, L. (2001). Random forests. *Machine learning 45*(1), 5–32.

Breza, E., A. G. Chandrasekhar, T. H. McCormick, and M. Pan (2020). Using aggregated relational data to feasibly identify network structure without network data. *American Economic Review 101*(8), 2454–84.

Brooks, R. L. (1941). On colouring the nodes of a network. In *Mathematical Proceedings of the Cambridge Philosophical Society*, Volume 37, pp. 194–197. Cambridge University Press.

Cai, J., A. De Janvry, and E. Sadoulet (2015). Social networks and the decision to insure. *American Economic Journal: Applied Economics 7*(2), 81–108.

Chernozhukov, V., D. Chetverikov, M. Demirer, E. Duflo, C. Hansen, W. Newey, and J. Robins (2018). Double/debiased machine learning for treatment and structural parameters.

Chernozhukov, V., D. Chetverikov, K. Kato, et al. (2014). Gaussian approximation of suprema of empirical processes. *The Annals of Statistics 42*(4), 1564–1597.

Chiang, H. D., K. Kato, Y. Ma, and Y. Sasaki (2019). Multiway cluster robust double/debiased machine learning. *arXiv preprint arXiv:1909.03489*.

Chin, A., D. Eckles, and J. Ugander (2018). Evaluating stochastic seeding strategies in networks. *arXiv preprint arXiv:1809.09561*.

Christakis, N. A. and J. H. Fowler (2010). Social network sensors for early detection of contagious outbreaks. *PloS one 5*(9).

Devroye, L., L. Györfi, and G. Lugosi (2013). *A probabilistic theory of pattern recognition*, Volume 31. Springer Science & Business Media.

DiTraglia, F. J., C. Garcia-Jimeno, R. O'Keeffe-O'Donovan, and A. Sanchez-Becerra (2020). Identifying causal effects in experiments with social interactions and non-compliance. *arXiv preprint arXiv:2011.07051*.

Duflo, E., P. Dupas, and M. Kremer (2011). Peer effects, teacher incentives, and the impact of tracking: Evidence from a randomized evaluation in kenya. *American Economic Review 101*(5), 1739–74.

Eckles, D., H. Esfandiari, E. Mossel, and M. A. Rahimian (2019). Seeding with costly network information. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pp. 421–422.

Egger, D., J. Haushofer, E. Miguel, P. Niehaus, and M. W. Walker (2019). General equilibrium effects of cash transfers: experimental evidence from kenya. Technical report, National Bureau of Economic Research.

Elliott, G. and R. P. Lieli (2013). Predicting binary outcomes. *Journal of Econometrics 174*(1), 15–26.

Farrell, M. H. (2015). Robust inference on average treatment effects with possibly more covariates than observations. *Journal of Econometrics 189*(1), 1–23.

Florios, K. and S. Skouras (2008). Exact computation of max weighted score estimators. *Journal of Econometrics 146*(1), 86–91.

Galeotti, A., B. Golub, and S. Goyal (2020). Targeting interventions in networks. *Econometrica, Forthcoming*.

Goldsmith-Pinkham, P. and G. W. Imbens (2013). Social networks and the identification of peer effects. *Journal of Business & Economic Statistics 31*(3), 253–264.

Graham, B. S., G. W. Imbens, and G. Ridder (2010). Measuring the effects of segregation in the presence of social spillovers: A nonparametric approach. Technical report, National Bureau of Economic Research.

Hajłasz, P. and J. Malỳ (2002). Approximation in sobolev spaces of nonlinear expressions involving the gradient. *Arkiv för Matematik 40*(2), 245–274.

Hansen, L. P. (1982). Large sample properties of generalized method of moments estimators. *Econometrica: Journal of the Econometric Society*, 1029–1054.

He, X. and K. Song (2018). Measuring diffusion over a large network. *arXiv preprint arXiv:1812.04195*.

Hirano, K. and J. R. Porter (2009). Asymptotics for statistical treatment rules. *Econometrica 77*(5), 1683–1701.

Hudgens, M. G. and M. E. Halloran (2008). Toward causal inference with interference. *Journal of the American Statistical Association 103*(482), 832–842.

Imai, K., Z. Jiang, and A. Malani (2020). Causal inference with interference and noncompliance in two-stage randomized experiments. *Journal of the American Statistical Association*, 1–13.

Imbens, G. W. (2000). The role of the propensity score in estimating dose-response functions. *Biometrika 87*(3), 706–710.

Jackson, M. O. and E. Storms (2018). Behavioral communities and the atomic structure of networks. *Available at SSRN 3049748*.

Kang, H. and G. Imbens (2016). Peer encouragement designs in causal inference with partial interference and identification of local average network effects. *arXiv preprint arXiv:1609.04464*.

Kempe, D., J. Kleinberg, and É. Tardos (2003). Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 137–146. ACM.

Kim, B. (2020). Analysis of randomized experiments with network interference and noncompliance. *arXiv preprint arXiv:2012.13710*.

Kim, D. A., A. R. Hwong, D. Stafford, D. A. Hughes, A. J. O'Malley, J. H. Fowler, and N. A. Christakis (2015). Social network targeting to maximise population behaviour change: a cluster randomised controlled trial. *The Lancet 386*(9989), 145–153.

Kitagawa, T. and A. Tetenov (2018). Who should be treated? Empirical welfare maximization methods for treatment choice. *Econometrica 86*(2), 591–616.

Kitagawa, T. and A. Tetenov (2019). Equality-minded treatment choice. *Journal of Business & Economic Statistics*, 1–14.

Kitagawa, T. and G. Wang (2020). Who should get vaccinated? individualized allocation of vaccines over sir network. *arXiv preprint arXiv:2012.04055*.

Laber, E. B., N. J. Meyer, B. J. Reich, K. Pacifici, J. A. Collazo, and J. M. Drake (2018). Optimal treatment allocations in space and time for on-line control of an emerging infectious disease. *Journal of the Royal Statistical Society: Series C (Applied Statistics) 67*(4), 743–789.

Ledoux, M. and M. Talagrand (2011). Probability in banach spaces. classics in mathematics.

Leung, M. P. (2020). Treatment and spillover effects under network interference. *Review of Economics and Statistics 102*(2), 368–380.

Li, X., P. Ding, Q. Lin, D. Yang, and J. S. Liu (2019). Randomization inference for peer effects. *Journal of the American Statistical Association*, 1–31.

Liu, L., M. G. Hudgens, B. Saul, J. D. Clemens, M. Ali, and M. E. Emch (2019). Doubly robust estimation in observational studies with partial interference. *Stat 8*(1), e214.

Manresa, E. (2013). Estimating the structure of social interactions using panel data. *Unpublished Manuscript. CEMFI, Madrid*.

Manski (2004). Statistical treatment rules for heterogeneous populations. *Econometrica 72*(4), 1221–1246.

Manski, C. F. (2013). Identification of treatment response with social interactions. *The Econometrics Journal 16*(1), S1–S23.

Mbakop, E. and M. Tabord-Meehan (2016). Model selection for treatment choice: Penalized welfare maximization. *arXiv preprint arXiv:1609.03167*.

Muralidharan, K., P. Niehaus, and S. Sukhtankar (2017). General equilibrium effects of (improving) public employment programs: Experimental evidence from india. Technical report, National Bureau of Economic Research.

Naimi, A. I., S. R. Cole, and E. H. Kennedy (2017). An introduction to g methods. *International journal of epidemiology 46*(2), 756–762.

Opper, I. M. (2016). Does helping john help sue? evidence of spillovers in education. *American Economic Review 109*(3), 1080–1115.

Paulin, D. (2012). Concentration inequalities in locally dependent spaces. *arXiv preprint arXiv:1212.2013*.

Pollard, D. (1990). Empirical processes: theory and applications. In *NSF-CBMS regional conference series in probability and statistics*, pp. i–86. JSTOR.

Robins, J. M., A. Rotnitzky, and L. P. Zhao (1994). Estimation of regression coefficients when some regressors are not always observed. *Journal of the American statistical Association 89*(427), 846–866.

Sävje, F., P. M. Aronow, and M. G. Hudgens (2020). Average treatment effects in the presence of unknown interference. *Annals of Statistics, Forthcoming*.

Sinclair, B., M. McConnell, and D. P. Green (2012). Detecting spillover effects: Design and analysis of multilevel experiments. *American Journal of Political Science 56*(4), 1055–1069.

Sobel, M. E. (2006). What do randomized studies of housing mobility demonstrate? causal inference in the face of interference. *Journal of the American Statistical Association 101*(476), 1398–1407.

Stoye, J. (2009). Minimax regret treatment choice with finite samples. *Journal of Econometrics 151*(1), 70–81.

Stoye, J. (2012). Minimax regret treatment choice with covariates or with limited validity of experiments. *Journal of Econometrics 166*(1), 138–156.

Su, L., W. Lu, and R. Song (2019). Modelling and estimation for optimal treatment decision with interference. *Stat 8*(1), e219.

Tetenov, A. (2012). Statistical treatment choice based on asymmetric minimax regret criteria. *Journal of Econometrics 166*(1), 157–165.

Van Der Vaart, A. W. and J. A. Wellner (1996). Weak convergence. In *Weak convergence and empirical processes*, pp. 16–28. Springer.

Vazquez-Bare, G. (2020). Causal spillover effects using instrumental variables. *arXiv preprint arXiv:2003.06023*.

Viviano, D. (2020). Policy choice in experiments with unknown interference. *arXiv preprint arXiv:2011.08174*.

Wager, S. and K. Xu (2020). Experimenting in equilibrium. *Management Science, Forthcoming*.

Wainwright, M. J. (2019). *High-dimensional statistics: A non-asymptotic viewpoint*, Volume 48. Cambridge University Press.

Wenocur, R. S. and R. M. Dudley (1981). Some special vapnik-chervonenkis classes. *Discrete Mathematics 33*(3), 313–318.

Zhou, Z., S. Athey, and S. Wager (2018). Offline multi-action policy learning: Generalization and optimization. *arXiv preprint arXiv:1810.04778*.

Zubcsek, P. P. and M. Sarvary (2011). Advertising to a social network. *Quantitative Marketing and Economics 9*(1), 71–107.