

A Service of

ZBW

Leibniz-Informationszentrum Wirtschaft Leibniz Information Centre for Economics

Kitagawa, Toru; Wang, Guanyi

Working Paper Who should get vaccinated? Individualized allocation of vaccines over SIR network

cemmap working paper, No. CWP28/21

Provided in Cooperation with: Institute for Fiscal Studies (IFS), London

Suggested Citation: Kitagawa, Toru; Wang, Guanyi (2021) : Who should get vaccinated? Individualized allocation of vaccines over SIR network, cemmap working paper, No. CWP28/21, Centre for Microdata Methods and Practice (cemmap), London, https://doi.org/10.47004/wp.cem.2021.2821

This Version is available at: https://hdl.handle.net/10419/246796

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



WWW.ECONSTOR.EU



Who should get vaccinated? Individualized allocation of vaccines over SIR network

Toru Kitagawa Guanyi Wang

The Institute for Fiscal Studies Department of Economics, UCL

cemmap working paper CWP28/21



Who Should Get Vaccinated? Individualized Allocation of Vaccines Over SIR Network

Toru Kitagawa[†]and Guanyi Wang[‡]

July 19, 2021

Abstract

How to allocate vaccines over heterogeneous individuals is one of the important policy decisions in pandemic times. This paper develops a procedure to estimate an individualized vaccine allocation policy under limited supply, exploiting social network data containing individual demographic characteristics and health status. We model the spillover effects of vaccination based on a Heterogeneous-Interacted-SIR network model and estimate an individualized vaccine allocation policy by maximizing an estimated social welfare (public health) criterion incorporating these spillovers. While this optimization problem is generally an NP-hard integer optimization problem, we show that the SIR structure leads to a submodular objective function, and provide a computationally attractive greedy algorithm for approximating a solution that has a theoretical performance guarantee. Moreover, we characterise a finite sample welfare regret bound and examine how its uniform convergence rate depends on the complexity and riskiness of the social network. In the simulation, we illustrate the importance of considering spillovers by comparing our method with targeting without network information.

Keywords: Vaccine allocation, Statistical treatment choice, Submodularity, SIR model, Social network, Spillovers.

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

[†]Department of Economics, University College London. Email: t.kitagawa@ucl.ac.uk

We would like to thank Andrew Chesher, Raffaella Giacomini, Aureo de Paula, Martin Weidner and the participants at 2020 UCL Economics department upgrade seminar. We are also benefited from the comments of the participants in the 2021 NASMES, AMES, and IAAE conferences. We thank Jeff Rowley for excellent research assistance. We gratefully acknowledge financial support from ERC grant (number 715940) and the ESRC Centre for Microdata Methods and Practice (CeMMAP) (grant number RES-589-28-0001).

1 Introduction

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

Since the start of COVID-19 pandemic, governments around the world have gone to great lengths to collect network data in which one can trace who is contacting whom. Motivated by these observations, we study how to estimate optimal individualized allocations of vaccines under capacity constraint, using micro-level social network data. Data is informative about the covariates of *N* units, their health status, and their associated neighbors. Using insample information, we evaluate the risk to each unit, calculated from its own covariates and spillovers from its heterogeneous neighbors, using an *individualized Susceptible-Infectious-Recovered* model. The purpose of vaccine allocation is to maximize public health, by selecting units to be vaccinated. Obtaining an optimal assignment is, however, challenging since whether a treatment is optimal for an individual depends on which treatments are given to her neighbors. This implies that the search for an optimal allocation has to be performed over the entire network jointly, not individually.

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

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

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

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

To assess how uncertainty in the SIR parameter estimates affect the welfare performance of the estimated policy, we derive a uniform upper bound of the welfare regret of our vaccine allocation rule and its convergence rate with respect to the size of the sample used for obtaining the SIR parameter estimates. The uniform upper bound of regret depends upon two things. Firstly, n, which is the sample size of the separate dataset used to estimate the SIR parameters. Secondly, the ratio of the network data sample size N to the maximum number of neighbors N_M plus the minimum between the number of infected units N_I and the number of available vaccine doses d (i.e., $(d \min\{N_M, d\} + 2dN_M + \min\{N_I, d\})/N$). As N_M and N_I grow, the complexity and risk of the social network increase, which can reduce the welfare regret performance of the estimated vaccine allocation rule.

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

1.1 Related Literature

Our work contributes to the literature on statistical treatment rules, which was first introduced into econometrics by Manski (2004). The optimal treatment allocation regime has been studied in many fields, such as medical statistics (Zhao et al., 2012, 2015), operational research (Loiola et al., 2007) and economics. Following the pioneering works of Hannan (1957) and Savage (1951),¹ researchers in econometrics and machine learning often use regret to evaluate the performance of decision rules. The recent literature of statistical treatment rules includes Dehejia (2005), Hirano and Porter (2009), Stoye (2009, 2012), Tetenov (2012), Bhattacharya and Dupas (2012), Kitagawa and Tetenov (2018), Zhou et al. (2018), Manski (2019), Kasy and Sautmann (2019), Athey and Wager (2020), Kock et al. (2020), Mbakop and Tabord-Meehan (2021), Manski and Tetenov (2021), Sakaguchi (2021), and Kitagawa et al. (2021) among others. The planner's objective function in the majority of these works is a sum of individual outcomes under the no-interference assumption (i.e., Stable Unit Treatment Value Assumption of Rubin (1974)). This assumption does not hold in this study because of the network spillover effects that are present. To our knowledge, there are only two other papers that also consider the network setting in statistical treatment choice, which are Viviano (2019) and Ananth (2020). These two papers assume the availability of pilot data from RCT studies performed over networks in order to form empirical welfare criteria. Their frameworks are not restricted to the SIR setting of the current paper and cover spillover structures commonly assumed in social science applications. In contrast, our approach forms welfare estimates by imposing the HI-SIR model structure and plugging in values of the primitive spillover parameters that are estimated or calibrated in some external study (e.g., Baqaee et al. (2020)). Another notable difference is that we consider allocation policies that are not constrained other than via the capacity constraint, while Viviano (2019) and Ananth (2020) assume the class of implementable allocation policies has a finite VC-dimension to control overfitting to the training RCT sample.

The SIR model was originally proposed by Kermack and McKendrick (1927), and is now the workhorse model in the epidemiological literature. Many extended versions have been studied in epidemiological analyses, such as the Susceptible-Infected-Susceptible model (Nåsell, 1996) and the Susceptible-Exposed-Infected-Recovered model (Li and Muldowney, 1995). During the global pandemic, an epidemological literature has sprung up within economics. Atkeson (2020) and Stock (2020) introduced the SIR model into economics to study the implications of the current pandemic on the US economy. We introduce heterogeneity into the SIR model, which is similar to what Acemoglu et al. (2020) does in studying the Multi-Risk SIR model. That paper assumes, however, that the infection rate after the release of a vaccine equals zero, which means it does not consider the vaccine allocation problem. Our work contributes to the current literature by studying micro-level vaccine as-

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

signment rules in a heterogeneous SIR model with network information. In contrast, the existing works analyzing vaccine allocation rules focus on solving for the optimal proportion of vaccinated units in the population (Pastor-Satorras and Vespignani (2002), Manski (2010, 2017)). Chen et al. (2020) analyzes vaccine allocation using a heterogeneous SIR model, while they consider vaccine allocation policies at the group-level rather than the individual-level.

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

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

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

2 Setup and Identification

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

Vaccine

$$1^{st} \text{ period}$$
 2^{nd} period
 $\uparrow A$ $[X_i, H_{0i}]_{i=1}^N$ $[H_{1i}]_{i=1}^N$ $t=2$ time

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

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

Assumption 2.1. (Undirected Relationships) The interference graph is undirected. i.e., $A_{ij} = A_{ji}$.

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

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

2.1 Heterogeneous-Interacted-SIR model

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

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

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

$$\mathbb{1}_{\{H_{1i}=S\}} + \mathbb{1}_{\{H_{1i}=I\}} + \mathbb{1}_{\{H_{1i}=R\}} + \mathbb{1}_{\{H_{1i}=D\}} = 1.$$
(2)

Without the vaccine, the state can move from susceptible to infected, then to either recovery or death. Now, we consider the setting after introducing the vaccine. Generally, vaccination has two purposes: the first is limiting the spread of disease, and the second is treatment. Vaccination builds up the immune system, which leads to recovery. However, the effective-ness of vaccination (i.e., the percentage of vaccinated units that recover) is not clear. For simplicity, we assume that assumption 2.2 holds.

Assumption 2.2. (*PT*) Perfect Treatment: A vaccinated unit enters the Recovered state, regardless of its previous state (i.e., $Pr(H_{1i} = R | v_i = 1) = 1$).

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

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

$$q_{i} = \left[\beta_{11} \frac{\sum_{j \in N_{i}} I_{j}(1-v_{j})a_{j}}{|N_{i}|} + \beta_{12} \frac{\sum_{j \in N_{i}} I_{j}(1-v_{j})b_{j}}{|N_{i}|}\right] \cdot a_{i} \\ + \left[\beta_{21} \frac{\sum_{j \in N_{i}} I_{j}(1-v_{j})a_{j}}{|N_{i}|} + \beta_{22} \frac{\sum_{j \in N_{i}} I_{j}(1-v_{j})b_{j}}{|N_{i}|}\right] \cdot b_{i},$$
(3)

where $\beta_{sk} = -\kappa_s \ln(1 - c_{sk})$, c_{sk} is the probability of successful disease transmission following a contact between group *s* and group *k* (i.e., c_{11} measures the transmission probability from one unit to another within the young group, c_{12} is the corresponding probability of transmission from a unit in the old group to a unit in the young group, with similar definitions for c_{21} and c_{22}), and κ_s is the average number of contacts in group *s* at each time period. β_{sk} describes the effective contact rate of the disease between group *s* and group *k*. The derivation of equation 3 can be found in the appendix.

In the above expression, $I_i(1 - v_i)$ means a susceptible individual can only be infected

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

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

$$p_i = 1 - a_i(\gamma_1 + \delta_1) - b_i(\gamma_2 + \delta_2).$$
(4)

Since the probability of recovery and death purely depend on personal physical fitness,³ there is no interactive part in equation 4. The transition probability to the infected state is then:

$$P_{Ii}(\mathbf{v}) \equiv \Pr(H_{1i} = I | X, \mathbf{v}, A, H_0) = S_i q_i \cdot (1 - v_i) + I_i p_i \cdot (1 - v_i).$$
(5)

In the above expression, the probability of an unvaccinated unit being infected has two components. The first is the probability of a healthy unit being infected. The second is the probability of staying in the infected state for those infected in the first period. Under Assumption 2.2, a vaccinated unit has zero probability of being infected. Similarly, the transition probability to the susceptible state is:

$$P_{Si}(\mathbf{v}) \equiv \Pr(H_{1i} = S | X, \mathbf{v}, A, H_0) = \left[1 - v_i - q_i(1 - v_i)\right] \cdot S_i.$$
 (6)

An unvaccinated unit can only exit the susceptible state by infection. Therefore, the probability of staying in the susceptible state decreases with the risk parameter β_{sk} , which depends on the number of infected neighbors and the number of contacts with them. The remaining two states do not rely on the network structure. First, the transition probability to the recovered state :

$$P_{Ri}(\mathbf{v}) \equiv \Pr(H_{1i} = R | X, \mathbf{v}, A, H_0) = v_i + \left[R_i + I_i(a_i\gamma_1 + b_i\gamma_2) \right] \cdot (1 - v_i).$$
(7)

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

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

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

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

death, which occurs with probability

$$P_{Di}(\mathbf{v}) \equiv \Pr(H_{1i} = D | X, \mathbf{v}, A, H_0) = I_i(a_i \delta_1 + b_i \delta_2) \cdot (1 - v_i).$$
(8)

2.2 Optimal Vaccine Allocation Problem

In Emanuel et al. (2020), a group of medical ethics experts suggest a successful vaccine is needed to reduce death and morbidity from infection, and is also needed for the restoration of economic and social activity. Following that suggestion, we choose our baseline outcome variable as the weighted average of the probability of being healthy in the second period. The idea of using weighted probability is to allow a flexible policy target of the planner. For example, if the planner wants to incorporate the importance of economic recovery into the policy objective, she may want to weight more the probabilities of being healthy of those who can contribute more to the economic output. For instance, the planner could specify the weights on the individuals to depend on their individuals characteristics including working hours and other socioeconomic characteristics (i.e., $g_i = g(X_i)$).⁵ We assume the weight is non-negative for every unit. Taking these into consideration, equation (9) specifies the goal of the vaccine allocation policy as a constrained optimization problem:

$$\max_{\mathbf{v}\in\{0,1\}^N} \frac{1}{N} \sum_{i=1}^N g_i \sum_{h\in\{S,R\}} P_{hi}(\mathbf{v}),$$
(9)

s.t.

 $\sum_{i=1}^{N} v_i \le d,\tag{10}$

where

$$P_{hi}(\mathbf{v}) = \Pr(H_{1i} = h | X, \mathbf{v}, A, H_0), \tag{11}$$

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

In equation (9), P_{hi} is the heterogeneous state transition function, which describes the probability of $h \in \{S, R\}$ in the second period. This transition probability depends on the individuals' covariates and previous state including whether being vaccinated or not, and the associated network structure. We adopt the HI-SIR model to formulate the above transition function, which has been provided in the previous subsection.

One relevant question is: Will vaccine allocation change the network structure? Yes, it

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

would change the behaviour of vaccinated units. For example, vaccinated units prefer to go out as compared to unvaccinated units. Given this, the number of contacts at each time period κ_s and the network structure A would change after the vaccine allocation. Our framework allows the network structure to vary without affecting the optimal allocation of vaccines in a special case where *only the vaccinated units* change their behaviours. This is because under our perfect treatment assumption, the vaccinated units no longer spread the disease or be infected, and their behavioral changes do not affect the health statuses of the neighbors and themselves. On the other hand, our framework cannot accommodate a general case where the unvaccinated units also change their behaviours, since if so the heterogeneous SIR parameters in the objective function change in response to the vaccine allocation. To allow this scenario, we could incorporate uncertainty as to the values of κ_s and A in the second period, for instance, by optimizing an objective function that takes the expectation of the SIR parameters the adjacency matrices conditional on v. We do not, however, consider such an extension in this paper and leave this topic for future research.

3 Estimation

In order to measure the individual risk level using the HI-SIR model, we need to know the associated SIR parameters: transmission rate (i.e., β_{11} , β_{12} , β_{21} , β_{22}), and recovery rate (i.e., γ_1 , γ_2). Given that we cannot observe the true value of those parameters, it is infeasible to evaluate the objective function (9) based on the in-sample information of (H_0 , X) and A of the target network. We therefore assume access to a separate dataset with sample size n or an external study analyzing it, from which we can form estimates for these exogenous parameters. We construct an empirical version of the population welfare (9) and maximize over the feasible allocation policies. To reflect the precision of the SIR parameter estimates in the welfare performance of an estimated allocation rule, we explicitly take into account the sampling uncertainty of the parameter estimates in our derivation of the welfare regret upper bound.

3.1 Estimation of SIR Parameters

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

$$\mathcal{R}_0 = \beta \times \frac{1}{\gamma}.$$
 (12)

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

$$\mathcal{R}_{0sk} = \beta_{sk} \times \frac{1}{\gamma_s} \quad \forall s, k = 1, 2,$$
(13)

where \mathcal{R}_{0sk} is the number of infectious individuals in group *s* resulting from one sick person in group *k*. We need to estimate the average number of younger infectious and older infectious from one sick person in group 1 and group 2, and also the recovery rate in each group. Given these values, we can estimate $\beta_{11}, \beta_{12}, \beta_{21}, \beta_{22}$ from equation (13).

Remark 3.1. We do not discuss what is a desirable procedure for estimating the model parameters in this work, since the choice of estimator depends on the type of data (e.g., Seroprevalence data, Reported cases data, etc.). See Keeling and Rohani (2011) for further details. For the COVID-19 transmissions, estimation of homogeneous \mathcal{R}_0 and other SIR parameters has been performed in several papers including Fernández-Villaverde and Jones (2020), Ferguson et al. (2020), and Korolev (2021). They note the difficulty in calibrating critical parameters at an early stage of the pandemic due to the lack of credible data, which motivates partial identification analysis of Manski and Molinari (2020) and Stoye (2021). Our approach, however, assumes availability of credible point estimates and does not allow identified-set estimates for the SIR parameters. See Ellison (2020) and Akbarpour et al. (2020) for recent estimates of heterogeneous SIR parameters.

3.2 Quadratic Integer Programming

Plugging the parameter estimates into our **HI-SIR** model, we now have the sample analog of the population maximization problem (9), which is

$$\max_{\mathbf{v}\in\{0,1\}^N} \mathcal{W}_n(\mathbf{v}), \quad \text{s.t.} \quad \sum_{i=1}^N v_i \le d, \qquad \text{where}$$

$$\mathcal{W}_n(\mathbf{v}) = \frac{1}{N} \sum_{i=1}^N g_i \sum_{h \in \{S,R\}} \hat{P}_{hi}(\mathbf{v}). \qquad (14)$$

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

$$\mathcal{W}_{n}(\mathbf{v}) = \frac{1}{N} \sum_{i=1}^{N} g_{i} \underbrace{\left[v_{i} + \left[R_{i} + (a_{i}\hat{\gamma}_{1} + b_{i}\hat{\gamma}_{2})I_{i} \right] (1 - v_{i}) + M_{i}S_{i}(1 - v_{i}) \right]}_{\text{Probability of being healthy}}$$
(15)

where

$$M_{i} = 1 - \frac{\sum_{j=1}^{N} (\hat{\beta}_{11}a_{i}a_{j} + \hat{\beta}_{12}a_{i}b_{j} + \hat{\beta}_{21}b_{i}a_{j} + \hat{\beta}_{22}b_{i}b_{j})A_{ij}I_{j}(1 - v_{j})}{|N_{i}|}.$$
 (16)

For the probability of being healthy in equation (15), there are two linear terms and one quadratic term in v. The first term measures the direct effect of vaccination. A vaccinated unit is safe from infection with 100% probability. The last two terms describe the probability of being free of infection for unvaccinated units. Infected units naturally recover with probability $\{\gamma_1, \gamma_2\}$, which depends on their own characteristics. For those units who are already recovered in the first period, they are free from infection in the second period. The last component takes into account the indirect effect of vaccination. For susceptible units, the probability of being infected by their infected neighbors is summarized by the interaction term.

After removing all the constant parts in equation (15), we obtain a simplified objective function (i.e., $W_n(\mathbf{v}) = F_n(\mathbf{v}) + constant$):

$$F_n(\mathbf{v}) = \sum_{i=1}^N \hat{c}_i v_i + \frac{1}{N} \sum_{i=1}^N \frac{1}{|N_i|} \sum_{j=1}^N T_{ij} (v_i + v_j - v_i v_j),$$
(17)

where

$$\hat{c}_i = g_i \left[1 - R_i - (a_i \hat{\gamma}_1 + b_i \hat{\gamma}_2) I_i - S_i \right] / N,$$
(18)

$$T_{ij} = g_i \big(\hat{\beta}_{11} a_i a_j + \hat{\beta}_{12} a_i b_j + \hat{\beta}_{21} b_i a_j + \hat{\beta}_{22} b_i b_j \big) A_{ij} I_j S_i.$$
(19)

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

4 Optimization

4.1 Submodularity

We showed in the last section that we can formulate our objective function as QAP. This kind of problem is well known as an *NP-hard* and *NP-hard to approximate* problem (Cela, 2013).

In general, we cannot solve QAP in polynomial time, which is an issue in practice. We shall, however, show that the quadratic integer programming in our vaccine allocation problem can be linked to the submodular optimization problem. The benefit of submodularity is that there exist off-the-shelf algorithms that can solve a submodular minimization problem in exact polynomial time and *approximately* solve a submodular maximization problem with capacity constraint in polynomial time. The seminal result of Nemhauser et al. (1978) provides a universal bound for the quality of approximation as detailed below in Section 4.2.

Definition 4.1 (Submodular function). Let $\mathcal{N} = \{1, 2, ..., N\}$. A real-valued set-function $F: 2^{\mathcal{N}} \to \mathbb{R}$ is submodular if and only if, for all subsets $A, B \subseteq \mathcal{N}$, we have: $F(A) + F(B) \ge F(A \cap B) + F(A \cup B)$.

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

$$F_n(V) = \mathbf{v}^{\mathsf{T}} \hat{W} \mathbf{v} + \hat{C}^{\mathsf{T}} \mathbf{v} - \mathbf{1}_{N \times 1}^{\mathsf{T}} \hat{W} \mathbf{v} - \mathbf{v}^{\mathsf{T}} \hat{W} \mathbf{1}_{N \times 1},$$
(20)

where

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

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

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

The quadratic functional form of F_n shown in (20) can be linked to one classic submodular function called a *cut function*. Cut functions have been well studied in combinatorial optimization and graph theory. We apply some of the results from that literature (e.g., Bach (2011)).

Lemma 4.1. Let $\hat{W} \in \mathbb{R}^{N \times N}$ and $\hat{C} \in \mathbb{R}^N$. Then the set function $F_n \colon V \mapsto \mathbf{v}^{\mathsf{T}} \hat{W} \mathbf{v} + \hat{C}^{\mathsf{T}} \mathbf{v} - \mathbf{1}_{N \times 1}^{\mathsf{T}} \hat{W} \mathbf{v} - \mathbf{v}^{\mathsf{T}} \hat{W} \mathbf{1}_{N \times 1}$ is submodular if and only if $\hat{w}_{ij} \leq 0 \ \forall i \neq j$.

⁶See the proof of Theorem 4.1 for a formal proof for the non-decreasing property of F_n .

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

Theorem 4.1. The objective function $F_n(V)$ is a non-decreasing submodular function for any adjacency matrix, covariate values, and parameter estimates.

Theorem 4.1 is the key result in our paper. It describes two important properties of our objective function; monotonicity and submodularity. We exploit these two properties to justify the uses of greedy maximization algorithms shown in the next subsection.

4.2 Greedy Maximization Algorithm

Greedy maximization algorithms for submodular functions have been studied and frequently used for well over forty years. The performance guarantee of the algorithm that we study was first introduced by Nemhauser et al. (1978). This algorithm essentially uses the diminishing returns property of the submodular function. The idea is to iteratively select the most valuable element until the capacity constraint is reached. At each round, the algorithm evaluates $\mathcal{O}(N)$ functions to identify the marginal gain of each element. The number of rounds depends on the capacity constraint *d*. As a result, the computational complexity of the greedy algorithm is of order $\mathcal{O}(N \cdot d)$, well below the computational complexity of the brute-force search. Algorithm 1 presents the greedy maximization algorithm applied to maximization of the empirical welfare (20).

Algorithr	n 1:	Capacity	y Cons	trained	Greed	ly A	lgoritł	ım
-----------	------	----------	--------	---------	-------	------	---------	----

1: **Input**: Dataset $\{S_i, I_i, R_i, a_i, b_i\}_{i=1}^N, \{A_{ij}\}_{i,j=1}^N$, estimated parameters $\{\hat{\beta}_{11}, \hat{\beta}_{12}, \hat{\beta}_{21}, \hat{\beta}_{22}, \hat{\gamma}_1, \hat{\gamma}_2\}$, weight $\{g_i\}_{i=1}^N$ and capacity constraint d;

2: **Initialization**: Starting from the empty set $V = \emptyset$;

- if |V| < d then
 - 3: for each $i \in \mathcal{N} \setminus V$ do

4: Compute the marginal gain $F_n(V + \{i\}) - F_n(V)$;

5: Select *i* which maximizes the marginal gain and add it into the set V;

```
else
```

```
return the set V; end
```

In general, there is no performance guarantee of the greedy algorithm. However, as shown by Nemhauser et al. (1978) for a non-decreasing submodular function with cardinality constraint (i.e., capacity constraint in our case), the greedy maximization algorithm is guaranteed to yield an allocation rule $\hat{V} \in \mathcal{V}_d$ that satisfies $F_n(\hat{V}) \ge (1 - \alpha_d)F_n(\hat{V}^*)$, where $\hat{V}^* \in \mathcal{V}_d$ is a constrained optimum under the capacity constraint, and α_d is a positive constant that depends only on $d \ge 1$ and $\alpha_d \ge 1/e$ for all $d \ge 1$. This seminal result implies that the greedy maximization algorithm provides a universal optimization guarantee for non-decreasing submodular functions, $F_n(\hat{V}) \ge (1 - 1/e)F_n(\hat{V}^*) \approx 0.63F_n(\hat{V}^*)$. Since we show in Theorem 4.1 that our objective function is non-decreasing and submodular, we obtain the following theorem as an immediate corollary of our Theorem 4.1 and Nemhauser et al. (1978).

Theorem 4.2 (Nemhauser et al. 1978). Let $F_n : 2^{\mathcal{N}} \to \mathbb{R}$ be the simplified empirical welfare function as defined in (20) and $\hat{V}^* \in \arg \max_{V \in \mathcal{V}_d} F_n(V)$, $d \ge 1$. The greedy algorithm shown in Algorithm 1 outputs $\hat{V} \in \mathcal{V}_d$ such that

$$F_n(\hat{V}) \ge (1 - \alpha_d) F_n(\hat{V}^*) \ge (1 - 1/e) F_n(\hat{V}^*),$$
(23)

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

4.3 Targeting Constraint

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

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

$$\sum_{i: X_i \in G_1} v_i = \sum_{i=1}^N a_i v_i \le d_1, \qquad \sum_{i: X_i \in G_2} v_i = \sum_{i=1}^N b_i v_i \le d_2.$$
(24)

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

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

Definition 4.2 (Matroid). Let \mathcal{I} be a nonempty family of allowable subsets of \mathcal{N} . Then the tuple $(\mathcal{N}, \mathcal{I})$ is a matroid if it satisfies:

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

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

$$\mathcal{I} \equiv \{ V \colon V \subseteq \mathcal{N}, |V \cap \mathcal{N}_1| \le d_1, |V \cap \mathcal{N}_2| \le d_2 \}.$$
(25)

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

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

Algorithm 2: Targeting Constraint Greedy Algorithm

1: **Input**: Dataset $\{S_i, I_i, R_i, a_i, b_i\}_{i=1}^N, \{A_{ij}\}_{i,j=1}^N$, estimated parameters $\{\hat{\beta}_{11}, \hat{\beta}_{12}, \hat{\beta}_{21}, \hat{\beta}_{22}, \hat{\gamma}_1, \hat{\gamma}_2\}$, weight $\{g_i\}_{i=1}^N$, capacity constraint d, and targeting constraints d_1, d_2 ;

2: Initialization: Starting from the empty set $V = \emptyset$;

- if |V| < d then
 - 3: for each $i \in \mathcal{N} \setminus V$ do
 - 4: Compute the marginal gain $F_n(V + \{i\}) F_n(V)$;
 - 5: Sort *i* in order of decreasing marginal gain
 - 6. if $\sum_{j \in V} a_j + a_{i(1)} \le d_1 \cap \sum_{j \in V} b_j + b_{i(1)} \le d_2$ then
 - 7: Add the 1st element of i into V;

```
else
```

end

8: Repeat step 6 with remaining i;

else

```
return the set V;
```

end

Proposition 4.1 (Fisher et al. 1978). Let $F_n : 2^{\mathcal{N}} \to \mathbb{R}$ be the simplified empirical welfare function as defined in (20) and $\hat{V}^{**} \in \arg \max_{V \in \mathcal{I}} F_n(V)$. The greedy maximization algorithm shown in Algorithm 2 outputs $\hat{V}' \in \mathcal{I}$ such that

$$F_n(\hat{V}') \ge \frac{1}{2} F_n(\hat{V}^{**}).$$
 (26)

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

4.4 Perfect Treatment Assumption and Submodularity

Recall Assumption 2.2 (Perfect Treatment): A vaccinated unit enters the Recovered state, regardless of its previous state (i.e., $Pr(H_{1i} = R | v_i = 1) = 1$). There are three possible ways to relax this assumption:

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

In the first case, if the person is recovered at H_0 , she will spread the disease during the first period. In that case, the recovered neighbors of unit *i* will be taken into account by the infection rate q_i . This will not, however, change the sign of our weighting matrix, hence submodularity (by Theorem 4.1) still holds. In the second case, if unit *i* is recovered in the first period (i.e., $H_{0i} = R$), she could become susceptible in the second period (i.e., $H_{1i} = S$). Then, she may be infected in the next period (i.e., $H_{2i} = I$). However, we only consider a one time period setting in this work, which rules out this risk. In the third case, varying this percentage only affects the coefficient of the linear term in the objective function (i.e., \hat{c}_i in equation 17), which is irrelevant to submodularity.

5 Regret Bounds

Following Manski (2004) and the subsequent literature on statistical treatment rules, we use regret to evaluate the performance of our algorithm for vaccine allocation. Let $F : 2^{\mathcal{N}} \to \mathbb{R}$ be the population analogue of $F_n(\cdot)$ in (20), where the estimated parameters are replaced by the truth. The expected regret measures the average difference in the welfare between using the constrained optimal assignment rule $V^* \in \arg \max_{V \in \mathcal{V}_d} F(V)$ and using the constrained estimated greedy algorithm \hat{V} obtained from Algorithm 1:

$$F(V^*) - \mathbb{E}_{P^n} \left[F(\hat{V}) \right] = \mathbb{E}_{P^n} \left[F(V^*) - F(\hat{V}) \right] \ge 0,$$
(27)

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

In this work, we assume that consistent estimators of effective contact rate and recovery rate are available from other studies. Generally, there is no requirement on the estimator except that Assumption 5.1 needs to hold.

Assumption 5.1. Let $\hat{\beta}_{sk}$ denote the estimate of effective contact rate between group s and group k, and $\hat{\gamma}_s$ denote the estimate of recovery rate in group s. The following properties need to hold:

$$\mathbb{P}\left\{ \left| \hat{\beta}_{sk} - \beta_{sk} \right| \ge \epsilon \right\} \le 2e^{-2n\epsilon^2} \quad \forall s, k = 1, 2.$$
(28)

$$\mathbb{P}\left\{\left|\hat{\gamma}_{s}-\gamma_{s}\right| \geq \epsilon\right\} \leq 2e^{-2n\epsilon^{2}} \quad \forall s=1,2,$$
(29)

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

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

The estimators for the contact rates and recovery rates may come from different studies with different sample sizes. In this case, we can view n in Assumption 5.1 as the smallest sample size among the studies.

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

$$F(V^*) - F(\hat{V}) = \underbrace{F(V^*) - F_n(\hat{V}^*)}_{(1)} + \underbrace{F_n(\hat{V}^*) - F_n(\hat{V})}_{(2)} + \underbrace{F_n(\hat{V}) - F(\hat{V})}_{(3)},$$
(30)

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

First, we start from the derivation of the upper bound of (1). This part is similar to the approach in Kitagawa and Tetenov (2018). Before looking at V^* , consider the following

inequality, which holds for any $\widetilde{V} \in \mathcal{V}_d$:

$$F(\widetilde{V}) - F_n(\widehat{V}^*) \le F(\widetilde{V}) - F_n(\widetilde{V})$$

$$\left(\because F_n(\widehat{V}^*) \ge F_n(\widetilde{V}) \right)$$

$$\le \sup_{V \in \mathcal{V}_d} |F_n(V) - F(V)|.$$
(31)

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

$$F(V^*) - F_n(\hat{V}^*) \le \sup_{V \in \mathcal{V}_d} |F_n(V) - F(V)|.$$
 (32)

For the second component, we can obtain an upper bound by applying Theorem 4.2:

$$F_{n}(\hat{V}^{*}) - F_{n}(\hat{V}) \leq \frac{1}{e} F_{n}(\hat{V}^{*})$$

$$\leq \frac{1}{e} (F_{n}(\hat{V}^{*}) - F(\hat{V}^{*})) + \frac{1}{e} F(V^{*})$$

$$\leq \frac{1}{e} |F_{n}(\hat{V}^{*}) - F(\hat{V}^{*})| + \frac{1}{e} F(V^{*})$$

$$\leq \frac{1}{e} \sup_{V \in \mathcal{V}_{d}} |F_{n}(V) - F(V)| + \frac{1}{e} F(V^{*}).$$
(33)

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

$$F_n(\hat{V}) - F(\hat{V}) \le |F_n(\hat{V}) - F(\hat{V})| \le \sup_{V \in \mathcal{V}_d} |F_n(V) - F(V)|.$$
(34)

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

$$F(V^*) - F(\hat{V}) \le \left(2 + \frac{1}{e}\right) \sup_{V \in \mathcal{V}_d} \left|F_n(V) - F(V)\right| + \frac{1}{e}F(V^*).$$
(35)

Compared with the regret upper bound when one could compute \hat{V}^* , the regret upper bound shown in (35) has one additional term $\frac{1}{e} \sup_{V \in \mathcal{V}_d} |F_n(V) - F(V)| + \frac{1}{e}F(V^*)$. This additional term comes from equation (33) and captures the welfare loss induced by the use of greedy algorithm. As we characterize below, the first term converges to zero as $n \to \infty$ under Assumption 5.1, while the second term remains independent of the accuracy of the parameter estimates. A simulation study in Section 6 assesses the magnitude of the optimization error of the greedy algorithm, and shows numerically that the greedy algorithm yields an exact optimum for small network cases (N = 35) at least. Based on this, we believe that the optimization error term of the greedy algorithm is much smaller than the universal theoretical bound $\frac{1}{e}F(V^*)$.

In the partition matroid (targeting constraint) case, by applying Proposition 4.1 and

repeating the arguments to derive (35), we obtain

$$F(V^{**}) - F(\hat{V}') \le \frac{5}{2} \sup_{V \in \mathcal{I}} \left| F_n(V) - F(V) \right| + \frac{1}{2} F(V^{**}), \tag{36}$$

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

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

$$\begin{aligned} \left|F_{n}(V)-F(V)\right| &= \left|\mathbf{v}^{\mathsf{T}}(\hat{W}-W)\mathbf{v}+(\hat{C}^{\mathsf{T}}-C^{\mathsf{T}})\mathbf{v}-\mathbf{1}_{N\times1}^{\mathsf{T}}(\hat{W}-W)\mathbf{v}-\mathbf{v}^{\mathsf{T}}(\hat{W}-W)\mathbf{1}_{N\times1}\right| \\ &\leq \left|\mathbf{v}^{\mathsf{T}}(\hat{W}-W)\mathbf{v}\right|+\left|(\hat{C}^{\mathsf{T}}-C^{\mathsf{T}})\mathbf{v}\right|+\left|\mathbf{1}_{N\times1}^{\mathsf{T}}(\hat{W}-W)\mathbf{v}\right|+\left|\mathbf{v}^{\mathsf{T}}(\hat{W}-W)\mathbf{1}_{N\times1}\right| \quad (37) \\ &\leq \mathbf{v}^{\mathsf{T}}\left|\hat{W}-W\right|\mathbf{v}+\left|(\hat{C}^{\mathsf{T}}-C^{\mathsf{T}})\right|\mathbf{v}+\mathbf{1}_{N\times1}^{\mathsf{T}}\left|\hat{W}-W\right|\mathbf{v}+\mathbf{v}^{\mathsf{T}}\left|\hat{W}-W\right|\mathbf{1}_{N\times1}, \end{aligned}$$

where the absolute value of a matrix or vector stands for the element-wise absolute values. Therefore, we can decompose the maximal deviation $\sup_{V \in \mathcal{V}} |F_n(V) - F(V)|$ into four parts:

$$\sup_{V \in \mathcal{V}_{d}} \left| F_{n}(V) - F(V) \right| \leq \sup_{V \in \mathcal{V}_{d}} \mathbf{v}^{\mathsf{T}} \left| \hat{W} - W \right| \mathbf{v} + \sup_{V \in \mathcal{V}_{d}} \left| \hat{C}^{\mathsf{T}} - C^{\mathsf{T}} \right| \mathbf{v} + \sup_{V \in \mathcal{V}_{d}} \mathbf{1}_{N \times 1}^{\mathsf{T}} \left| \hat{W} - W \right| \mathbf{v} + \sup_{V \in \mathcal{V}_{d}} \mathbf{v}^{\mathsf{T}} \left| \hat{W} - W \right| \mathbf{1}_{N \times 1}.$$
(38)

Under Assumption 5.1, we can obtain an upper bound for the mean of each element in $\hat{W} - W$ and $\hat{C} - C$, as shown in the next lemma.

Lemma 5.1. Under Assumption 2.1, 2.2, and 5.1, we have

$$\mathbb{E}_{P^n} \left| \hat{w}_{ij} - w_{ij} \right| \le \sqrt{\frac{1 + \ln(2)}{2n}} \frac{A_{ij}g_i}{N}, \qquad \mathbb{E}_{P^n} \left| \hat{c}_i - c_i \right| \le \sqrt{\frac{1 + \ln(2)}{2n}} \frac{I_i g_i}{N}.$$
(39)

Combining this lemma with equations (35) and (38), we obtain the following theorem:

Theorem 5.1. Let $N_M = \max_{i \in \mathcal{N}} |N_i|$, N_I be the total number of infected units, and $g = \max_{i \in \mathcal{N}} g_i$. Under Assumptions 2.1, 2.2, and 5.1, we have

$$\mathbb{E}_{P^n}\left[F(V^*) - F(\hat{V})\right] \le \bar{C} \cdot \frac{g\left[d\min\{N_M, d\} + 2dN_M + \min\{N_I, d\}\right]}{N} \sqrt{\frac{1}{n}} + \frac{1}{e}F(V^*), \quad (40)$$

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

Proof of the above theorem is shown in the appendix. In Theorem 5.1, we provided a distribution-free upper bound on the expected regret. We show that the convergence rate of the upper bound depends on the network data sample size N and also the sample size n for estimating the SIR parameters. At the same time, the regret upper bound is increasing in the complexity and the riskiness of the network. The intuition is that our algorithm finds it harder to identify the most valuable units when the maximum number of edges and the

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

6 Simulation Exercises

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

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

Table 1: Summary of the SIR parameter values

In addition, we choose three different densities, 0.1, 0.5 and 1, in order to identify the effect of network complexity. Here, density = 1 means that the network is fully connected (i.e., complete graph). We choose *full* to understand the behaviour of our heuristic algorithm not only in the sparse network case but also in the densest case. We also compare three capacity constraints, d = 7%N, 10%N, 20%N, to evaluate the marginal performance gain of our greedy algorithm. We choose equal weight in the following comparisons. We, however, show the impact of changing weights on the number of vaccinated younger units in Table 5.

In the following sections, we compare our greedy algorithm with three familiar allocation rules. We first compare our algorithm with a brute force method in order to find the difference between the potentially sub-optimal greedy solution and the brute-force optimal solution. However, the number of possible combinations dramatically increases with the number of nodes and the capacity constraint. We cannot use a large number of agents to compute the brute force optimum in the simulation. Given this, in Section 6.2, we use a random assignment rule as a baseline to evaluate the performance of our algorithm in a large network setting. The third allocation rule that we compare our greedy algorithm with is an allocation rule which assigns the vaccine without considering network information. We compare the greedy algorithm with this third rule in Section 6.3.

Allocation Rule	Greedy Algorithm			Brute Force ⁸			
Capacity Constraint	$\overline{d=7\%N}$	d=10%N	d=20%N	d = 7% N	d=10%N	d = 20% N	
Parameter set 1							
N = 500, $density = 0.1$	0.60	0.65	0.77	0.60	0.65	0.77	
N 500 1 1 0 5	(0.21)	(0.22)	(0.26)	(0.21)	(0.22)	(0.26)	
N = 500, aensity = 0.5	0.4/	(0.51)	0.03	0.47	(0.51)	0.63	
	(0.40)	(0.39)	(0.39)	(0.40)	(0.39)	(0.39)	
N = 500, $density = 1$	0.33	0.3/	0.49	0.33	0.3/	0.49	
N. 000 1 1/ 0 1	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	
N = 800, density = 0.1	0.58	0.66	0.76	0.58	0.66	0.76	
	(0.23)	(0.25)	(0.26)	(0.23)	(0.25)	(0.26)	
N = 800, density = 0.5	0.44	0.52	0.62	0.44	0.52	0.62	
	(0.38)	(0.38)	(0.38)	(0.38)	(0.38)	(0.38)	
N = 800, density = 1	0.30	0.36	0.46	0.30	0.36	0.46	
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	
Parameter set 2							
N = 500, density = 0.1	0.59	0.64	0.77	0.59	0.64	0.77	
	(0.27)	(0.29)	(0.32)	(0.27)	(0.29)	(0.32)	
N = 500, density = 0.5	0.42	0.46	0.59	0.42	0.46	0.59	
	(0.49)	(0.49)	(0.49)	(0.49)	(0.49)	(0.49)	
N = 500, density = 1	0.25	0.29	0.41	0.25	0.29	0.41	
, ,	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	
N = 800, density = 0.1	0.57	0.65	0.76	0.57	0.65	0.76	
, ,	(0.28)	(0.30)	(0.31)	(0.28)	(0.30)	(0.31)	
N = 800, density = 0.5	0.40	0.47	0.58	0.40	0.47	0.58	
	(0.47)	(0.47)	(0.47)	(0.47)	(0.47)	(0.47)	
N = 800, density = 1	0.22	0.29	0.40	0.22	0.29	0.40	
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	

6.1 Comparing with Brute Force

Table 2: The value of welfare (the sum of probabilities of being healthy in the second period) averaged over 100 random networks (standard errors in parentheses). We use the Greedy Algorithm or the Brute Force algorithm to determine who in each network should be vaccinated.

Since Theorem 4.2 shows the gap between the optimal solution and the heuristic result is at most 37%, we want to explore this theoretical difference using numerical study. We list all the possible combinations and use brute force to search for the optimal solution given a

⁸We compare all possible combinations given the capacity constraint and select the set V that maximizes W_n .

manageable number of units. We specify the maximum number of units to be N = 35, which is limited by computer performance. As the number of nodes increases, the possible number of combinations grows exponentially. The memory requirement and running time become impractical in a more realistic case. We recognize that the results from a small network may not be accurate in a large network setting, but help us to understand the regret of our greedy algorithm to some degree. We summarize the in-sample welfare W_n of these two approaches in Table 2.

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

6.2 Comparing With Random Assignment

In this section, we use a random assignment rule to define the baseline of vaccine allocation. We randomly draw an allocation 10,000 times and calculate the average value of the outcome variable. Random allocation is one common assignment rule for policymakers. The purpose of this simulation is to learn about the improvement of our greedy allocation rule. In order to evaluate its performance in a relatively large network setting, we choose N = 500 and 800. Table 3 records the main differences in terms of in-sample welfare between these two methods.

From Table 3, we find that the performance of both methods decreases with the number of edges, which is also true for the first comparison. As the number of edges increase, the greedy algorithm finds it harder to identify who is relatively crucial in the network, which supports our interpretation of Theorem 5.1 in the previous section. This effect becomes more pronounced as the capacity constraint is relaxed. In the most extreme case, when everyone is connected with each other, the performance of our method is still better than the random assignment rule. This performance gap widens with the capacity constraint. We also find that the average welfare increases by 12% when the capacity constraint increases by 0.1N. Moreover, this improvement is robust with respect to the variation of number of nodes and the changes of density levels of network. The number of nodes decreases the performance of our method in a sparse network setting. For N = 800, welfare in the densest network is 14% lower than the welfare with *density* = 0.5, no matter which capacity constraint and parameter set we use.

Allocation Rule	Greedy Algorithm			Random Assignment ⁹			
Capacity Constraint	d=7%N	d=10%N	d=20%N	d = 7% N	d=10%N	d = 20% N	
Parameter set 1							
N = 500, density = 0.1	0.61	0.65	0.77	0.57	0.59	0.66	
	(0.00)	(0.01)	(0.01)	(0.00)	(0.00)	(0.00)	
N = 500, density = 0.5	0.61	0.64	0.76	0.57	0.59	0.66	
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	
N = 500, density = 1	0.61	0.64	0.76	0.57	0.59	0.66	
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	
N = 800, density = 0.1	0.59	0.63	0.75	0.55	0.57	0.64	
	(0.04)	(0.04)	(0.04)	(0.04)	(0.04)	(0.04)	
N = 800, density = 0.5	0.48	0.51	0.63	0.44	0.46	0.53	
	(0.07)	(0.07)	(0.07)	(0.07)	(0.07)	(0.07)	
N = 800, density = 1	0.34	0.37	0.49	0.30	0.32	0.39	
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	
Parameter set 2							
N = 500, density = 0.1	0.60	0.64	0.77	0.56	0.59	0.66	
	(0.01)	(0.01)	(0.01)	(0.00)	(0.00)	(0.00)	
N = 500, density = 0.5	0.60	0.64	0.76	0.56	0.59	0.66	
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	
N = 500, density = 1	0.60	0.64	0.76	0.56	0.59	0.66	
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	
N = 800, density = 0.1	0.58	0.62	0.74	0.54	0.56	0.63	
	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	
N = 800, density = 0.5	0.44	0.48	0.60	0.40	0.43	0.50	
	(0.08)	(0.08)	(0.08)	(0.08)	(0.08)	(0.08)	
N = 800, $density = 1$	0.26	0.30	0.42	0.23	0.25	0.32	
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	

Table 3: The value of welfare (the sum of probabilities of being healthy in the second period) averaged over 100 random networks (standard errors in parentheses). We use the Greedy Algorithm or the Random allocation to determine who in each network should be vaccinated.

If we look at the random assignment rule in Table 3, its performance is much worse than the performance of the greedy algorithm. This difference increases when the complexity of and the number of nodes in the network increase. The performance of the random assignment rule improves as we relax the capacity constraint. However, this improvement is only about 7% when the capacity constraint increases by 0.1N. Compared with the greedy algorithm, random assignment is less effective. Given its scarcity, we waste considerable resources by randomly assigning the vaccine. Looking at the situation of full edges, the performance of random allocation is inferior. The ratio of the welfare attained by random allocation to the welfare attained by the greedy algorithm is illustrated in Figure 1. This ratio increases slowly with the number of edges and deceases with the number of nodes in the network. In addition, the ratio decreases in an obvious way with the number of vaccines that are available.

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



Figure 1: Comparison between Greedy Algorithm and Random allocation

6.3 Comparing With Targeting Without Network Information

Allocation Rule	C	Greedy Algorith	hm	TWNI ¹⁰		
Capacity Constraint	d = 7%N	d=10%N	d=20%N	d = 7%N	d=10%N	d = 20% N
Parameter set 1						
N = 500, density = 0.1	0.61	0.65	0.77	0.57	0.59	0.65
	(0.00)	(0.01)	(0.01)	(0.00)	(0.01)	(0.01)
N = 500, density = 0.5	0.61	0.64	0.76	0.57	0.59	0.65
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
N = 500, density = 1	0.61	0.64	0.76	0.57	0.59	0.65
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
N = 800, density = 0.1	0.59	0.63	0.75	0.56	0.58	0.65
	(0.04)	(0.04)	(0.04)	(0.04)	(0.04)	(0.04)
N = 800, density = 0.5	0.48	0.51	0.63	0.44	0.47	0.54
	(0.07)	(0.07)	(0.07)	(0.07)	(0.07)	(0.07)
N = 800, density = 1	0.34	0.37	0.49	0.30	0.33	0.40
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Parameter set 2						
N = 500, density = 0.1	0.60	0.64	0.77	0.56	0.58	0.65
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
N = 500, density = 0.5	0.60	0.64	0.76	0.56	0.58	0.65
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
N = 500, $density = 1$	0.60	0.64	0.76	0.56	0.58	0.65
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
N = 800, density = 0.1	0.58	0.62	0.74	0.55	0.57	0.64
	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)
N = 800, density = 0.5	0.44	0.48	0.60	0.41	0.43	0.50
	(0.08)	(0.08)	(0.08)	(0.08)	(0.08)	(0.08)
N = 800, $density = 1$	0.26	0.30	0.42	0.23	0.26	0.33
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)

Table 4: The value of welfare (the sum of probabilities of being healthy in the second period) averaged over 100 random networks (standard errors in parentheses). We use the Greedy Algorithm or the Targeting Without Network Information allocation to determine who in each network should be vaccinated

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

Generally, TWNI assigns treatment based on personal characteristics. In this study, we only have one covariate: *age*. This means either the old group receives the vaccine or the young group receives the vaccine. Under the previous setting (i.e., older people are more likely to be infected and to die), group 2 will consume the entire vaccine allocation. Given different capacity constraints, this assignment rule selects units to be vaccinated from group 2 until the upper bound is reached. Table 4 indicates the results for TWNI allocation are similar to those for random allocation. In addition, despite the outcome value varying with the SIR parameters, the sizable improvement from using network information to allocate vaccination is quite robust to variations in the size and density of network.



Figure 2: Comparison between Greedy Algorithm and Targeting Without Network Information

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

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

from considering network information is robust to variation of the SIR parameters, and an allocation rule which ignores spillovers waste a sizeable proportion of a scarce resource.

In Table 5, we illustrate the impact on the percentage of vaccinated younger units by varying the weight choice g_i (In this simulation exercise, we choose equal weight for the units in same group). If we assign weight $g_1 = 1.5$ for G_1 , we find all the vaccines are consumed by younger units. Comparing with the equal weight case, this number changes dramatically. Moreover, we find our greedy algorithm offers more vaccines to younger units in the case of parameter set 2 than parameter set 1, i.e., when the transmission rate parameters are higher within and across the groups.

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

Table 5: The percentage of vaccinated younger units in the second period under the vaccine allocation policies obtained by Greedy Algorithm, averaged over 100 random networks. We choose three different sets of weights in this comparison

7 Conclusion

In this work, we have introduced a novel method to estimate individualized vaccine allocation rules under network interference. We introduce the heterogeneous-interacted-SIR model to specify the spillover effects of infectious disease. We show that the welfare objective function of the vaccine allocation problem is non-decreasing and submodular, and so is its empirical analogue formed by plugging in the estimates of the SIR parameters. Based on this specific diminishing returns property, we provide a greedy algorithm with performance guarantee under two different exogenous constraints, which can easily accommodate various targets that policymakers commonly face in reality. Moreover, we show that this algorithm implies an upper bound for regret that converges uniformly at $O(n^{-1/2})$. Using simulation, we point out the importance of considering network information in the allocation problem. Several open questions and extensions are worth considering in future work. First, this paper considered a one-time vaccine allocation. We did not consider if there are multiple allocation periods, and how to decide the allocation dynamically. A relevant important question is how to jointly optimize allocations and timing of first- and second-doses of vaccines, as recently discussed for Covid-19 vaccines in Maier et al. (2021), Tuite et al. (2021), and Wang et al. (2021)). Moreover, we do not study how the vaccine allocation rule impacts on the outcome variables after multiple periods. As discussed in Bu et al. (2020), changes to the network structure should be considered in a dynamic setting. Second, we only compare the greedy algorithm with the brute-force optimum in a small network. Other than the universal bounds of Theorem 5.1, we do not know the performance of our method relative to the optimal solution in the large network data setting. Third, we did not impose any other constraints than the capacity constraint and the targeting constraint. For interpretability and fairness, we may want to additionally restrict the policy rule as a simple function of observed covariates. We regard these as interesting questions that are worthy of consideration.

References

- ACEMOGLU, D., V. CHERNOZHUKOV, I. WERNING, AND M. D. WHINSTON (2020): "A multirisk SIR model with optimally targeted lockdown," Tech. rep., National Bureau of Economic Research.
- AKBARPOUR, M., C. COOK, A. MARZUOLI, S. MONGEY, A. NAGARAJ, M. SACCAROLA, P. TEBALDI, S. VASSERMAN, AND H. YANG (2020): "Socioeconomic network heterogeneity and pandemic policy response," Tech. rep., National Bureau of Economic Research.
- ANANTH, A. (2020): "Optimal Treatment Assignment Rules on Networked Populations," *working paper*.
- ARONOW, P. M., C. SAMII, ET AL. (2017): "Estimating average causal effects under general interference, with application to a social network experiment," *The Annals of Applied Statistics*, 11, 1912–1947.
- ATHEY, S., D. ECKLES, AND G. W. IMBENS (2018): "Exact p-values for network interference," *Journal of the American Statistical Association*, 113, 230–240.
- ATHEY, S. AND S. WAGER (2020): "Policy learning with observational data," Econometrica.
- ATKESON, A. (2020): "What will be the economic impact of covid-19 in the us? rough estimates of disease scenarios," Tech. rep., National Bureau of Economic Research.
- BACH, F. (2011): "Learning with submodular functions: A convex optimization perspective," *arXiv preprint arXiv:1111.6453*.

- BAKSHY, E., J. M. HOFMAN, W. A. MASON, AND D. J. WATTS (2011): "Everyone's an influencer: quantifying influence on twitter," in *Proceedings of the fourth ACM international conference on Web search and data mining*, 65–74.
- BANERJEE, A., A. G. CHANDRASEKHAR, E. DUFLO, AND M. O. JACKSON (2013): "The diffusion of microfinance," *Science*, 341.
- BAQAEE, D., E. FARHI, M. J. MINA, AND J. H. STOCK (2020): "Reopening Scenarios," Tech. rep., National Bureau of Economic Research.
- BASSE, G., A. FELLER, AND P. TOULIS (2019): "Randomization tests of causal effects under interference," *Biometrika*, 106, 487–494.
- BHATTACHARYA, D. AND P. DUPAS (2012): "Inferring welfare maximizing treatment assignment under budget constraints," *Journal of Econometrics*, 167, 168–196.
- BU, F., A. E. AIELLO, J. XU, AND A. VOLFOVSKY (2020): "Likelihood-based Inference for Partially Observed Epidemics on Dynamic Networks," *Journal of the American Statistical Association*, 1–17.
- CELA, E. (2013): *The quadratic assignment problem: theory and algorithms*, vol. 1, Springer Science & Business Media.
- CHEN, W., C. WANG, AND Y. WANG (2010): "Scalable influence maximization for prevalent viral marketing in large-scale social networks," in *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*, 1029–1038.
- CHEN, X., M. LI, D. SIMCHI-LEVI, AND T. ZHAO (2020): "Allocation of COVID-19 Vaccines Under Limited Supply," *Available at SSRN 3678986*.
- CUNNINGHAM, W. H. (1985): "Minimum cuts, modular functions, and matroid polyhedra," *Networks*, 15, 205–215.
- DEHEJIA, R. H. (2005): "Program evaluation as a decision problem," *Journal of Econometrics*, 125, 141–173.
- DOMINGOS, P. AND M. RICHARDSON (2001): "Mining the network value of customers," in Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining, 57–66.
- ELLISON, G. (2020): "Implications of heterogeneous SIR models for analyses of COVID-19," Tech. rep., National Bureau of Economic Research.
- EMANUEL, E. J., G. PERSAD, A. KERN, A. BUCHANAN, C. FABRE, D. HALLIDAY, J. HEATH, L. HERZOG, R. LELAND, E. T. LEMANGO, ET AL. (2020): "An ethical framework for global vaccine allocation," *Science*, 369, 1309–1312.

- FERGUSON, N., D. LAYDON, G. NEDJATI-GILANI, N. IMAI, K. AINSLIE, M. BAGUELIN, S. BHATIA, A. BOONYASIRI, Z. CUCUNUBÁ, G. CUOMO-DANNENBURG, ET AL. (2020): "Report 9: Impact of non-pharmaceutical interventions (NPIs) to reduce COVID19 mortality and healthcare demand," *Imperial College London*, 10, 77482.
- FERNÁNDEZ-VILLAVERDE, J. AND C. I. JONES (2020): "Estimating and Simulating a SIRD Model of COVID-19 for Many Countries, States, and Cities," Tech. rep., National Bureau of Economic Research.
- FISHER, M. L., G. L. NEMHAUSER, AND L. A. WOLSEY (1978): "An analysis of approximations for maximizing submodular set functions—II," in *Polyhedral combinatorics*, Springer, 73– 87.
- GOLDENBERG, J., B. LIBAI, AND E. MULLER (2001): "Talk of the network: A complex systems look at the underlying process of word-of-mouth," *Marketing letters*, 12, 211–223.
- GRAHAM, B. AND A. DE PAULA (2020): *The Econometric Analysis of Network Data*, Academic Press.
- GRANOVETTER, M. (1978): "Threshold models of collective behavior," American journal of sociology, 83, 1420–1443.
- HANNAN, J. (1957): "APPROXIMATION TO BAYES RISK IN REPEATED PLAY," Contributions to the Theory of Games, 21, 97.
- HIRANO, K. AND J. R. PORTER (2009): "Asymptotics for statistical treatment rules," *Econometrica*, 77, 1683–1701.
- HOEFFDING, W. (1963): "Probability inequalities for sums of bounded random variables," *Journal of the American Statistical Association*, 58, 13–30.
- KASY, M. AND A. SAUTMANN (2019): "Adaptive treatment assignment in experiments for policy choice," *CESifo Working Paper*.
- KEELING, M. J. AND P. ROHANI (2011): *Modeling infectious diseases in humans and animals*, Princeton University Press.
- KEMPE, D., J. KLEINBERG, AND É. TARDOS (2003): "Maximizing the spread of influence through a social network," in Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining, 137–146.
- KERMACK, W. O. AND A. G. MCKENDRICK (1927): "A contribution to the mathematical theory of epidemics," *Proceedings of the royal society of london. Series A, Containing papers of a mathematical and physical character*, 115, 700–721.

- KITAGAWA, T., S. SAKAGUCHI, AND A. TETENOV (2021): "Constrained Classification and Policy Learning," *arXiv preprint arXiv:2106.12886*.
- KITAGAWA, T. AND A. TETENOV (2018): "Who should be treated? empirical welfare maximization methods for treatment choice," *Econometrica*, 86, 591–616.
- KLINE, B. AND E. TAMER (2020): "Econometric analysis of models with social interactions," in *The Econometric Analysis of Network Data*, Elsevier, 149–181.
- KOCK, A. B., D. PREINERSTORFER, AND B. VELIYEV (2020): "Functional sequential treatment allocation," *Journal of the American Statistical Association*, 1–36.
- KOOPMANS, T. C. AND M. BECKMANN (1957): "Assignment problems and the location of economic activities," *Econometrica: journal of the Econometric Society*, 53–76.
- KOROLEV, I. (2021): "Identification and Estimation of the SEIRD Epidemic Model for COVID-19," *Journal of econometrics*, 220, 63–85.
- KRAUSE, A. AND D. GOLOVIN (2014): "Submodular function maximization." .
- LEE, L.-F., X. LIU, E. PATACCHINI, AND Y. ZENOU (2020): "Who is the key player? A network analysis of juvenile delinquency," *Journal of Business & Economic Statistics*, 1–9.
- LEUNG, M. P. AND H. R. MOON (2019): "Normal Approximation in Large Network Models," *Available at SSRN 3377709*.
- LI, M. Y. AND J. S. MULDOWNEY (1995): "Global stability for the SEIR model in epidemiology," *Mathematical biosciences*, 125, 155–164.
- LI, S. AND S. WAGER (2020): "Random Graph Asymptotics for Treatment Effect Estimation under Network Interference," *arXiv preprint arXiv:2007.13302*.
- LOIOLA, E. M., N. M. M. DE ABREU, P. O. BOAVENTURA-NETTO, P. HAHN, AND T. QUERIDO (2007): "A survey for the quadratic assignment problem," *European journal of operational research*, 176, 657–690.
- LUGOSI, G. (2002): "Pattern classification and learning theory," in *Principles of nonparametric learning*, Springer, 1–56.
- MAIER, B. F., A. BURDINSKI, A. H. ROSE, F. SCHLOSSER, D. HINRICHS, C. BETSCH, L. KORN,
 P. SPRENGHOLZ, M. MEYER-HERMANN, T. MITRA, ET AL. (2021): "Potential benefits of delaying the second mRNA COVID-19 vaccine dose," *arXiv preprint arXiv:2102.13600*.
- MANSKI, C. F. (2004): "Statistical treatment rules for heterogeneous populations," *Econometrica*, 72, 1221–1246.

—— (2010): "Vaccination with partial knowledge of external effectiveness," *Proceedings* of the National Academy of Sciences, 107, 3953–3960.

—— (2013): "Identification of treatment response with social interactions," *The Econometrics Journal*, 16, S1–S23.

—— (2017): "Mandating vaccination with unknown indirect effects," *Journal of Public Economic Theory*, 19, 603–619.

(2019): "Treatment choice with trial data: statistical decision theory should supplant hypothesis testing," *The American Statistician*, 73, 296–304.

- MANSKI, C. F. AND F. MOLINARI (2020): "Estimating the COVID-19 infection rate: Anatomy of an inference problem," *Journal of Econometrics*.
- MANSKI, C. F. AND A. TETENOV (2021): "Statistical Decision Properties of Imprecise Trials Assessing Coronavirus Disease 2019 (COVID-19) Drugs," *Value in Health*, 24, 641–647.
- MBAKOP, E. AND M. TABORD-MEEHAN (2021): "Model selection for treatment choice: Penalized welfare maximization," *Econometrica*, 89, 825–848.
- MIAO, Y. (2010): "Concentration inequality of maximum likelihood estimator," *Applied mathematics letters*, 23, 1305–1309.
- NÅSELL, I. (1996): "The quasi-stationary distribution of the closed endemic SIS model," *Advances in Applied Probability*, 895–932.
- NEMHAUSER, G. L., L. A. WOLSEY, AND M. L. FISHER (1978): "An analysis of approximations for maximizing submodular set functions—I," *Mathematical programming*, 14, 265–294.
- PASTOR-SATORRAS, R. AND A. VESPIGNANI (2002): "Immunization of complex networks," *Physical review E*, 65, 036104.
- PESARAN, H. AND C. F. YANG (2020): "Matching Theory and Evidence on COVID-19 using a Stochastic Model of Epidemics on Networks," *IAAE Webinar*.
- RUBIN, D. B. (1974): "Estimating causal effects of treatments in randomized and nonrandomized studies." *Journal of educational Psychology*, 66, 688.
- SAKAGUCHI, S. (2021): "Estimation of Optimal Dynamic Treatment Assignment Rules under Policy Constraint," *arXiv preprint arXiv:2106.05031*.
- SAVAGE, L. J. (1951): "The theory of statistical decision," *Journal of the American Statistical association*, 46, 55–67.

- SÄVJE, F., P. M. ARONOW, AND M. G. HUDGENS (2017): "Average treatment effects in the presence of unknown interference," *arXiv preprint arXiv:1711.06399*.
- STOCK, J. H. (2020): "Data gaps and the policy response to the novel coronavirus," Tech. rep., National Bureau of Economic Research.
- STOYE, J. (2009): "Minimax regret treatment choice with finite samples," Journal of Econometrics, 151, 70–81.
- ——— (2012): "Minimax regret treatment choice with covariates or with limited validity of experiments," *Journal of Econometrics*, 166, 138–156.
- ——— (2021): "Bounding Infection Prevalence by Bounding Selectivity and Accuracy of Tests: With Application to Early COVID-19," *arXiv preprint arXiv:2008.06178*.
- TETENOV, A. (2012): "Statistical treatment choice based on asymmetric minimax regret criteria," *Journal of Econometrics*, 166, 157–165.
- TUITE, A. R., L. ZHU, D. N. FISMAN, AND J. A. SALOMON (2021): "Alternative dose allocation strategies to increase benefits from constrained COVID-19 vaccine supply," *Annals of internal medicine*.
- VAZQUEZ-BARE, G. (2020): "Causal spillover effects using instrumental variables," *arXiv* preprint arXiv:2003.06023.
- VIVIANO, D. (2019): "Policy targeting under network interference," *arXiv preprint arXiv:1906.10258*.
- WANG, X., Z. DU, K. JOHNSON, S. FOX, M. LACHMANN, J. S. MCLELLAN, AND L. A. MEYERS (2021): "The impacts of COVID-19 vaccine timing, number of doses, and risk prioritization on mortality in the US," *medRxiv*.
- ZHAO, Y., D. ZENG, A. J. RUSH, AND M. R. KOSOROK (2012): "Estimating individualized treatment rules using outcome weighted learning," *Journal of the American Statistical Association*, 107, 1106–1118.
- ZHAO, Y.-Q., D. ZENG, E. B. LABER, AND M. R. KOSOROK (2015): "New statistical learning methods for estimating optimal dynamic treatment regimes," *Journal of the American Statistical Association*, 110, 583–598.
- ZHOU, Z., S. ATHEY, AND S. WAGER (2018): "Offline multi-action policy learning: Generalization and optimization," *arXiv preprint arXiv:1810.04778*.

Appendix A The Transmission Term

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

$$1 - q_{i} = (1 - c_{11})^{\frac{\kappa_{1} \sum_{j \in N_{i}} I_{j}(1 - v_{j})a_{j}a_{i}}{|N_{i}|}} \cdot (1 - c_{12})^{\frac{\kappa_{1} \sum_{j \in N_{i}} I_{j}(1 - v_{j})b_{j}a_{i}}{|N_{i}|}} \cdot (1 - c_{22})^{\frac{\kappa_{2} \sum_{j \in N_{i}} I_{j}(1 - v_{j})b_{j}b_{i}}{|N_{i}|}}.$$
(41)

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

$$q_i = 1 - e^{-z}, (42)$$

where

$$z = \frac{\beta_{11}}{|N_i|} \sum_{j \in N_i} I_j (1 - v_j) a_j a_i + \frac{\beta_{12}}{|N_i|} \sum_{j \in N_i} I_j (1 - v_j) b_j a_i + \frac{\beta_{21}}{|N_i|} \sum_{j \in N_i} I_j (1 - v_j) a_j b_i + \frac{\beta_{22}}{|N_i|} \sum_{j \in N_i} I_j (1 - v_j) b_j b_i.$$
(43)

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

$$q_i \simeq z. \tag{44}$$

Appendix B Lemmas

B.1 Preliminary Lemma

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

Lemma B.1 (Proposition 6.3 Bach (2011)). Let $Q \in \mathbb{R}^{p \times p}$, $q \in \mathbb{R}^{p}$, and $\mathcal{N} = \{1, 2, ..., p\}$. For $A \in 2^{\mathcal{N}}$, define $1_{A} = (1_{1 \in A}, ..., 1_{p \in A})'$. The function $F \colon A \mapsto q^{\mathsf{T}} 1_{A} + \frac{1}{2} \mathbf{1}_{A}^{\mathsf{T}} Q \mathbf{1}_{A}$ is submodular if and only if all off-diagonal elements of Q are non-positive.

Lemma B.2 (Theorem 2.2 Cunningham (1985)). Function F is a cut function if and only if:

For any three disjoint subsets A, B, C of S,

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

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

Lemma B.3 (Hoeffding's inequality Hoeffding (1963)). Let $X_1, ..., X_n$ be independent bounded random variables such that X_i falls in the interval $[a_i, b_i]$ with probability one. Denote their sum by $S_n = \sum_{i=1}^n X_i$. Then for any $\epsilon > 0$ we have

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

and

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

B.2 Proof of Lemma 4.1

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

Proof. The first step is to show our objective function is a cut function based on Lemma B.2. In our case, simply consider three arbitrary disjoint sets $A, B, C \subseteq \mathcal{N}$.

$$F_n(A) = \sum_{j \in A} \{ \hat{w}_{jj} + \hat{c}_j \} + \sum_{i \neq j \in A} \{ \hat{w}_{ij} \} - \sum_{j \in A} \sum_{m=1}^N \{ \hat{w}_{mj} + \hat{w}_{jm} \},$$
(45)

$$F_n(B) = \sum_{j \in B} \{ \hat{w}_{jj} + \hat{c}_j \} + \sum_{i \neq j \in B} \{ \hat{w}_{ij} \} - \sum_{j \in B} \sum_{m=1}^N \{ \hat{w}_{mj} + \hat{w}_{jm} \},$$
(46)

$$F_{n}(\{A \cup B\}) = \sum_{j \in \{A \cup B\}} \{\hat{w}_{jj} + \hat{c}_{j}\} + \sum_{i \neq j \in \{A \cup B\}} \{\hat{w}_{ij}\} - \sum_{j \in \{A \cup B\}} \sum_{m=1}^{N} \{\hat{w}_{mj} + \hat{w}_{jm}\}$$

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

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

$$= F_{n}(A) + F_{n}(B).$$
(47)

Therefore, we have:

$$F_n(\{A \cup B\} \cup C) = F_n(\{A \cup B\}) + F_n(C)$$

= $F_n(A) + F_n(B) + F_n(C)$ (48)

Combining the previous results, we get:

$$F_n(\{A \cup B\} \cup C) = F_n(A \cup B) + F_n(A \cup C) + F_n(B \cup C) - F_n(A) - F_n(B) - F_n(C) - F_n(\emptyset),$$
(49)

since $F_n(\emptyset) = 0$.

Now, we have shown that $F_n(V)$ is a cut function. The next step is to find the sufficient and necessary conditions for submodularity of the cut function. Lemma B.1 indicates, for any cut function which can be written as a quadratic function plus a linear part, submodularity holds if and only if all off-diagonal elements of the weighting matrix are non-positive. That requires $\hat{w}_{ij} \leq 0, \forall i \neq j$.

B.3 Proof of Lemma 5.1

Under Assumption 2.1, 2.2, and 5.1, we have

$$\mathbb{E}_{P^n} \left| \hat{w}_{ij} - w_{ij} \right| \le \sqrt{\frac{1 + \ln(2)}{2n}} \frac{A_{ij}g_i}{N}, \qquad \mathbb{E}_{P^n} \left| \hat{c}_i - c_i \right| \le \sqrt{\frac{1 + \ln(2)}{2n}} \frac{I_i g_i}{N}.$$
(50)

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

$$\hat{w}_{ij} - w_{ij} = \frac{S_i g_i A_{ij} I_j}{|N_i| N} \Big[(\beta_{11} - \hat{\beta}_{11}) a_i a_j + (\beta_{12} - \hat{\beta}_{12}) a_i b_j + (\beta_{21} - \hat{\beta}_{21}) b_i a_j + (\beta_{22} - \hat{\beta}_{22}) b_i b_j \Big].$$
(51)

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

$$\mathbb{E}_{P^{n}} \left| \hat{w}_{ij} - w_{ij} \right| = \mathbb{E}_{P^{n}} \left| \frac{S_{i}g_{i}A_{ij}I_{j}}{|N_{i}|N} \left[(\hat{\beta}_{11} - \beta_{11})a_{i}a_{j} + (\hat{\beta}_{12} - \beta_{12})a_{i}b_{j} + (\hat{\beta}_{21} - \beta_{21})b_{i}a_{j} + (\hat{\beta}_{22} - \beta_{22})b_{i}b_{j} \right] \right|$$

$$\leq \frac{S_{i}g_{i}A_{ij}I_{j}a_{i}a_{j}}{|N_{i}|N} \mathbb{E}_{P^{n}} \left| \hat{\beta}_{11} - \beta_{11} \right| + \frac{S_{i}g_{i}A_{ij}I_{j}a_{i}b_{j}}{|N_{i}|N} \mathbb{E}_{P^{n}} \left| \hat{\beta}_{12} - \beta_{12} \right|$$

$$+ \frac{S_{i}g_{i}A_{ij}I_{j}b_{i}a_{j}}{|N_{i}|N} \mathbb{E}_{P^{n}} \left| \hat{\beta}_{21} - \beta_{21} \right| + \frac{S_{i}g_{i}A_{ij}I_{j}b_{i}b_{j}}{|N_{i}|N} \mathbb{E}_{P^{n}} \left| \hat{\beta}_{22} - \beta_{22} \right|.$$
(52)

Since β_{sk} is the effective contact rate of the disease between group s and k, it is naturally

bounded in [0, 1]. We can apply Lemma B.3 to get the upper bound of each component:

$$\mathbb{P}\left\{ \left| \hat{\beta}_{sk} - \beta_{sk} \right| \ge \epsilon \right\} \le 2e^{-2n\epsilon^2} \quad \forall s, k = 1, 2.$$
(53)

Now we can bound $\mathbb{E}(|\hat{\beta} - \beta|)$. Recall that for any nonnegative random variable Y, $\mathbb{E}(Y) = \int_0^\infty \mathbb{P}(Y \ge t) dt$. Hence, for any a > 0,

$$\mathbb{E}(|\hat{\beta} - \beta|^2) = \int_0^\infty \mathbb{P}(|\hat{\beta} - \beta|^2 \ge t)dt$$

= $\int_0^a \mathbb{P}(|\hat{\beta} - \beta|^2 \ge t)dt + \int_a^\infty \mathbb{P}(|\hat{\beta} - \beta|^2 \ge t)dt$ (54)
 $\le a + \int_a^\infty \mathbb{P}(|\hat{\beta} - \beta|^2 \ge t)dt.$

Equation (53) implies that $\mathbb{P}(|\hat{\beta} - \beta| \ge \sqrt{t}) \le 2e^{-2nt}$. Hence,

$$\mathbb{E}(|\hat{\beta} - \beta|^2) \leq a + \int_a^\infty \mathbb{P}(|\hat{\beta} - \beta|^2 \geq t)dt$$

$$= a + \int_a^\infty \mathbb{P}(|\hat{\beta} - \beta| \geq \sqrt{t})dt$$

$$\leq a + 2\int_a^\infty e^{-2nt}dt$$

$$= a + \frac{e^{-2na}}{n}.$$
(55)

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

$$\mathbb{E}(|\hat{\beta} - \beta|^2) \le \frac{\ln(2)}{2n} + \frac{1}{2n} = \frac{1 + \ln(2)}{2n}.$$
(56)

Therefore, we have

$$\mathbb{E}(|\hat{\beta} - \beta|) \le \sqrt{(\mathbb{E}(|\hat{\beta} - \beta|^2))} \le \sqrt{\frac{1 + \ln(2)}{2n}}.$$
(57)

Plugging this upper bound back to equation (52), we get

$$\mathbb{E}\left|\hat{w}_{ij} - w_{ij}\right| \leq \sqrt{\frac{1 + \ln(2)}{2n}} (a_i S_i A_{ij} a_j I_j + a_i S_i A_{ij} b_j I_j + b_i S_i A_{ij} a_j I_j + b_i S_i A_{ij} b_j I_j) \frac{g_i}{|N_i| N}$$

$$(\because |N_i| \geq 1 \text{ and by treating 0 neighbor as equal to 1})$$

$$\leq \sqrt{\frac{1 + \ln(2)}{2n}} \frac{A_{ij} g_i}{N}.$$
(58)

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

$$\hat{c}_i - c_i = -(\hat{\gamma}_1 - \gamma_1) \frac{a_i I_i g_i}{N} - (\hat{\gamma}_2 - \gamma_2) \frac{b_i I_i g_i}{N}.$$
(59)

Take the absolute value and expectation of both sides,

$$\mathbb{E}|\hat{c}_{i} - c_{i}| = \mathbb{E}\left| (\hat{\gamma}_{1} - \gamma_{1}) \frac{a_{i}I_{i}g_{i}}{N} + (\hat{\gamma}_{2} - \gamma_{2}) \frac{b_{i}I_{i}g_{i}}{N} \right|$$

$$\leq \mathbb{E}|\hat{\gamma}_{1} - \gamma_{1}| \frac{a_{i}I_{i}g_{i}}{N} + \mathbb{E}|\hat{\gamma}_{2} - \gamma_{2}| \frac{b_{i}I_{i}g_{i}}{N}.$$
(60)

With the same idea as for β , γ is also bounded in [0, 1]. By using Lemma B.3, we get

$$\mathbb{E}|\hat{\gamma} - \gamma| \le \sqrt{\frac{1 + \ln(2)}{2n}} \quad \forall \hat{\gamma} = \hat{\gamma}_1, \hat{\gamma}_2.$$
(61)

Plugging this upper bound back into equation 60, we get

$$\mathbb{E}|\hat{c}_{i} - c_{i}| \leq \sqrt{\frac{1 + \ln(2)}{2n}} (a_{i} + b_{i}) \frac{I_{i}g_{i}}{N} = \sqrt{\frac{1 + \ln(2)}{2n}} \frac{I_{i}g_{i}}{N}.$$
(62)

Appendix C Proofs for Theorems

C.1 Proof of Theorem 4.1

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

Proof. Recall

$$F_n(V) = \mathbf{v}^{\mathsf{T}} \hat{W} \mathbf{v} + \hat{C}^{\mathsf{T}} \mathbf{v} - \mathbf{1}_{N \times 1}^{\mathsf{T}} \hat{W} \mathbf{v} - \mathbf{v}^{\mathsf{T}} \hat{W} \mathbf{1}_{N \times 1}.$$
(63)

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

$$\tilde{F}_n(V) = \tilde{\mathbf{v}}^{\mathsf{T}} \hat{W} \tilde{\mathbf{v}} + \hat{C}^{\mathsf{T}} \tilde{\mathbf{v}} - \mathbf{1}_{N \times 1}^{\mathsf{T}} \hat{W} \tilde{\mathbf{v}} - \tilde{\mathbf{v}}^{\mathsf{T}} \hat{W} \mathbf{1}_{N \times 1}.$$
(64)

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

$$\frac{\partial F_n(V)}{\partial \tilde{\mathbf{v}}} = \tilde{\mathbf{v}}^{\mathsf{T}} \hat{W}^{\mathsf{T}} + \tilde{\mathbf{v}}^{\mathsf{T}} \hat{W} + \hat{C}^{\mathsf{T}} - \mathbf{1}_{N \times 1}^{\mathsf{T}} \hat{W} - \mathbf{1}_{N \times 1}^{\mathsf{T}} \hat{W}^{\mathsf{T}}$$
$$= \underbrace{(\tilde{\mathbf{v}}^{\mathsf{T}} - \mathbf{1}_{N \times 1}^{\mathsf{T}})}_{\leq 0} \underbrace{\hat{W}}_{\leq 0}^{\mathsf{T}} + \underbrace{(\tilde{\mathbf{v}}^{\mathsf{T}} - \mathbf{1}_{N \times 1}^{\mathsf{T}})}_{\leq 0} \underbrace{\hat{W}}_{\leq 0} + \underbrace{\hat{C}}_{\geq 0}^{\mathsf{T}} \geq 0.$$
(65)

Given all the elements in $\frac{\partial \tilde{F}_n(V)}{\partial \tilde{\mathbf{v}}}$ are non-negative, this non-decreasing property also holds under the integer increment in every element of \mathbf{v} . Therefore, $F_n(V)$ is a non-decreasing set function. Combining this with Lemma 4.1, we complete the proof.

C.2 Proof of Theorem 5.1

Let $N_M = \max_{i \in \mathcal{N}} |N_i|$, N_I be the total number of infected units, and $g = \max_{i \in \mathcal{N}} g_i$. Under Assumptions 2.1, 2.2, and 5.1, we have

$$\mathbb{E}_{P^n}\left[F(V^*) - F(\hat{V})\right] \le \bar{C} \cdot \frac{g\left[d\min\{N_M, d\} + 2dN_M + \min\{N_I, d\}\right]}{N} \sqrt{\frac{1}{n}} + \frac{1}{e}F(V^*), \quad (66)$$

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

Proof.

$$\mathbb{E}_{P^{n}}\left[\sup_{V\in\mathcal{V}_{d}}\left|F_{n}(V)-F(V)\right|\right] \leq \mathbb{E}_{P^{n}}\left[\sup_{V\in\mathcal{V}_{d}}\mathbf{v}^{\mathsf{T}}\left|\hat{W}-W\right|\mathbf{v}\right] + \mathbb{E}_{P^{n}}\left[\sup_{V\in\mathcal{V}_{d}}\left|\hat{C}^{\mathsf{T}}-C^{\mathsf{T}}\right|\mathbf{v}\right] \\
+ \mathbb{E}_{P^{n}}\left[\sup_{V\in\mathcal{V}_{d}}\mathbf{1}_{N\times1}^{\mathsf{T}}\left|\hat{W}-W\right|\mathbf{v}\right] + \mathbb{E}_{P^{n}}\left[\sup_{V\in\mathcal{V}_{d}}\mathbf{v}^{\mathsf{T}}\left|\hat{W}-W\right|\mathbf{1}_{N\times1}\right] \\
= \sup_{V\in\mathcal{V}_{d}}\mathbf{v}^{\mathsf{T}}\mathbb{E}_{P^{n}}\left|\hat{W}-W\right|\mathbf{v} + \sup_{V\in\mathcal{V}_{d}}\mathbb{E}_{P^{n}}\left|\hat{C}^{\mathsf{T}}-C^{\mathsf{T}}\right|\mathbf{v} \\
+ \sup_{V\in\mathcal{V}_{d}}\mathbf{1}_{N\times1}^{\mathsf{T}}\mathbb{E}_{P^{n}}\left|\hat{W}-W\right|\mathbf{v} + \sup_{V\in\mathcal{V}_{d}}\mathbf{v}^{\mathsf{T}}\mathbb{E}_{P^{n}}\left|\hat{W}-W\right|\mathbf{1}_{N\times1}.$$
(67)

From equation (67), $\mathbb{E}_{P^n} \left[\sup_{V \in \mathcal{V}_d} |F_n(V) - F(V)| \right]$ can be decomposed into four components. Since v only contains $\{0, 1\}$ and the absolute value must be non-negative, V that maximizes each component under capacity constraint must select units with a greater number of edges, as compare to those that are not selected. We define the maximum number of edges for each unit as N_M . Hence, the number of edges for selected units must be lower or equal to N_M . Next, we look at each term in equation 67 separately. Using Lemma 5.1, the first term is bounded as:

$$\sup_{V \in \mathcal{V}_{d}} \mathbf{v}^{\mathsf{T}} \mathbb{E}_{P^{n}} \left| \hat{W} - W \right| \mathbf{v} \leq \sup_{V \in \mathcal{V}_{d}} \sqrt{\frac{1 + \ln(2)}{2n}} \sum_{i \in V} \sum_{j \in V} \frac{A_{ij}g_{i}}{N}$$

$$\leq \sqrt{\frac{1 + \ln(2)}{2n}} \sum_{i \in V} \frac{g_{i} \min\{N_{M}, d\}}{N}$$

$$\left(\because \sum_{j \in V} A_{ij} \leq \min\{N_{M}, d\} \quad \forall i \in \mathcal{N} \right)$$

$$\leq \sqrt{\frac{1 + \ln(2)}{2n}} \frac{dg \cdot \min\{N_{M}, d\}}{N}.$$
(68)

The second term is bounded as:

$$\sup_{V \in \mathcal{V}_d} \mathbb{E}_{P^n} \left| \hat{C}^{\mathsf{T}} - C^{\mathsf{T}} \right| \mathbf{v} \le \sup_{V \in \mathcal{V}_d} \sqrt{\frac{1 + \ln(2)}{2n}} \sum_{i \in V} \frac{I_i g_i}{N} \le \sqrt{\frac{1 + \ln(2)}{2n}} \frac{g \min\{N_I, d\}}{N}.$$
(69)

The third term is bounded as:

$$\sup_{V \in \mathcal{V}_d} \mathbf{1}_{N \times 1}^{\mathsf{T}} \mathbb{E}_{P^n} \left| \hat{W} - W \right| \mathbf{v} \leq \sup_{V \in \mathcal{V}_d} \sqrt{\frac{1 + \ln(2)}{2n}} \sum_{i \in \mathcal{N}} \sum_{j \in V} \frac{A_{ij}g_i}{N}$$

$$\leq \sqrt{\frac{1 + \ln(2)}{2n}} \frac{\sum_{j \in V} gN_M}{N}$$

$$= \sqrt{\frac{1 + \ln(2)}{2n}} \frac{dgN_M}{N}.$$
(70)

The fourth term is bounded as:

$$\sup_{V \in \mathcal{V}_{d}} \mathbf{v}^{\mathsf{T}} \mathbb{E}_{P^{n}} \left| \hat{W} - W \right| \mathbf{1}_{N \times 1} \leq \sup_{V \in \mathcal{V}_{d}} \sqrt{\frac{1 + \ln(2)}{2n}} \sum_{i \in V} \sum_{j \in \mathcal{N}} \frac{A_{ij}g_{i}}{N}$$

$$\leq \sqrt{\frac{1 + \ln(2)}{2n}} \frac{\sum_{i \in V} gN_{M}}{N}$$

$$= \sqrt{\frac{1 + \ln(2)}{2n}} \frac{dgN_{M}}{N}.$$
(71)

_

Combining the bounds of the four terms, we get

$$\mathbb{E}_{P^{n}}\left[\sup_{V\in\mathcal{V}_{d}}\left|F_{n}(V)-F(V)\right|\right] \leq \sqrt{\frac{1+\ln(2)}{2n}} \frac{dg\min\{N_{M},d\}}{N} + 2\sqrt{\frac{1+\ln(2)}{2n}} \frac{dgN_{M}}{N} + \sqrt{\frac{1+\ln(2)}{2n}} \frac{g\min\{N_{I},d\}}{N} = \frac{dg\min\{N_{M},d\} + 2dgN_{M} + g\min\{N_{I},d\}}{N} \sqrt{\frac{1+\ln(2)}{2n}}.$$
(72)

Therefore, we have from equation (35)

$$\mathbb{E}_{P^{n}}[F(V^{*}) - F(\hat{V})] \leq \left(2 + \frac{1}{e}\right) \frac{dg \min\{N_{M}, d\} + 2dgN_{M} + g \min\{N_{I}, d\}}{N} \sqrt{\frac{1 + \ln(2)}{2n}} \\ + \frac{1}{e}F(V^{*}) \\ = \left(2 + \frac{1}{e}\right) \frac{g[d \min\{N_{M}, d\} + 2dN_{M} + \min\{N_{I}, d\}]}{N} \sqrt{\frac{1 + \ln(2)}{2n}} \\ + \frac{1}{e}F(V^{*})$$
(73)

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