# Better Vaccination Strategies for Better People

Po-An Chen*
University of Southern California
poanchen@usc.edu

Mary David†
University of Southern California
marysdavid@gmail.com

David Kempe‡
University of Southern California
dkempe@usc.edu

## ABSTRACT

In this paper, we study the vaccination of graphs against the outbreak of infectious diseases, in the following natural model generalizing a model by Aspnes et al.: An infectious disease breaks out at a random node of the graph and propagates along the edges of the graph. Vaccinated nodes cannot be infected, nor pass on the infection, whereas all other nodes do. The decisions on which nodes get vaccinated must be made before the random outbreak location is known. There is a cost associated with vaccination and a different cost with getting infected.

In this model, we provide two results. First, we improve the approximation guarantee for finding the best vaccination strategy from $O(\log^{1.5} n)$ to $O(\log z)$ (where $z$ is the support size of the outbreak distribution), by rounding a natural linear program with region-growing techniques. Second, we analyze the impact of autonomy on the part of the nodes: while a benevolent authority may suggest which nodes should be vaccinated, nodes may opt out after being chosen. We analyze the "Price of Opting Out" in this sense under partially altruistic behavior. Individuals base their decisions on their own cost and the societal cost, the latter scaled by some factor $\beta$. If the altruism parameter is $\beta = 0$, it is known that the Price of Anarchy and Price of Stability can be $\Theta(n)$. We show that with positive altruism, Nash Equilibria may not exist, but the Price of Opting Out is at most $1/\beta$ (whereas the Price of Anarchy can remain at $\Theta(n)$).

**Categories and Subject Descriptors:** G.2.3 [Mathematics of Computing]: Discrete Mathematics–Applications

**General Terms:** Economics, Theory, Algorithms

*Supported in part by ONR Young Investigator Award N00014-08-1-0733.

†Supported in part by a USC WiSE Fellowship and a Google Anita Borg Scholarship.

‡Supported in part by NSF CAREER award 0545855, and ONR Young Investigator Award N00014-08-1-0733.


*EC'10,* June 7–11, 2010, Cambridge, Massachusetts, USA.


**Keywords:** altruism, selfishness, vaccination, inoculation, stability, anarchy, approximation

## 1. INTRODUCTION

Recent epidemics of the Avian Flu and Swine Flu, among others, have reinforced the dramatic vulnerability of our society to outbreaks of epidemic diseases. Similarly, recent worms and viruses (such as Storm or Conficker) have shown us how severe consequences can be for massive infections of nodes in a computer network. Protecting social and computer networks from such outbreaks is a task of paramount social and economic importance.

Strategies for protecting a network fall roughly into two categories: preventive and reactive. Reactive strategies attempt to isolate nodes of the network (individuals or machines) once they have been diagnosed with an infection. Prophylactic vaccinations or inoculations protect nodes so that they will in the future not be affected by an outbreak[1]. In this paper, we focus on preventive strategies, i.e., the decisions which nodes in a network should be vaccinated.

As was evident from the extensive coverage of vaccinations against H1N1, there are many factors complicating the allocation of vaccines. Among the most prominent ones are (1) supply shortages, limiting the number of individuals who can be vaccinated, and (2) node autonomy: individuals make their own decisions on whether to get inoculated, which may conflict with the socially optimal strategy. The former naturally leads to optimization problems for allocating the limited amounts of vaccine, while the latter raises natural questions about the inefficiency of game-theoretic outcomes in such settings.

We study these issues in a natural model for network vaccinations, generalizing a model proposed by Aspnes et al. [4]. In this model, vaccinated nodes can never contract the disease, so they are effectively removed from the network. After all vaccination decisions are made, the disease will break out at a node $v$ chosen according to a known probability distribution $p_v$ and infect all nodes reachable in the network (with the vaccinated nodes removed). There is a cost of $C_v$ associated with node $v$ being vaccinated and a cost of $L_v$ for $v$ being infected. (A formal description of the model is given in Section 2.) The optimization goal is to find a set of nodes to vaccinate that minimizes the expected total cost of all nodes.

[1]In practice, the protection may not be perfect; considering inoculations which succeed only with a certain probability is an interesting direction for future work.

Our first result in this paper is an improved approximation algorithm for the goal of minimizing the total cost. The previous best algorithm was an $O(\log^{1.5} n)$ approximation due to Aspnes et al. [4], based on repeated greedy invocations of the Arora-Rao-Vazirani Sparsest Cut Algorithm [3]. In Section 4, we give an Integer Program for the optimization problem of minimizing the expected number of infections, under a hard constraint on the total number of vaccinated nodes. Using region-growing techniques [11, 10], we derive an $(O(1), O(\log z))$ bicriteria approximation (where $z$ is the size of the support of the distribution $p$), violating the vaccination constraint by a factor $O(\log z)$. Up to constant factors, this bicriteria result, applied with $z = 1$, also subsumes the MinSBCC result of [14]. We give an example showing that the integrality gap of the LP matches our guarantees, even in the bicriteria setting. Using the bicriteria approximation as a subroutine, we then obtain the following theorem.

THEOREM 1.1. *There is a polynomial-time $O(\log z)$ approximation algorithm for minimizing the expected social cost.*

In reality, the decision of whether to get vaccinated usually lies with the individual nodes, whose interests do not necessarily align with the social goal of minimizing the total cost. Individuals tend to undervaccinate when they are not concerned with the impact of their action on other nodes. It is therefore natural to investigate how inefficient "societally stable" states can become as a result of individual decisions and their externalities. Indeed, Aspnes et al. [4] already showed that each instance of the inoculation game has at least one pure Nash Equilibrium and that the Price of Anarchy [16] can be $\Theta(n)$ in the worst case (but no worse).

The $\Theta(n)$ lower bound relies on the fact that individuals are entirely selfish in their vaccination decisions, and completely unaware of — or indifferent to — how their decisions may affect others in the network. Time and again, this assumption of selfishness has been found to be violated in controlled experiments [18, 19]. In a recent paper, Meier et al. [20] analyzed the impact of *friendship* on stable outcomes. In their model, a node's utility is the sum of its own cost and a $\beta$ fraction of the cost of all its neighbors. Meier et al. show that for some graphs, this notion of friendship leads to significantly more efficient equilibria, while for others, the improvement is small.

In this paper, we instead consider a notion of *altruism* first mentioned by Ledyard [18] and studied recently in the context of traffic routing [6]. In this model, an individual's utility is the convex combination of his own cost (with weight $1-\beta$) and the cost of *all* other nodes (with weight $\beta$). Thus, our notion models a general feeling of altruism or responsibility for the welfare of society, and the parameter $\beta$ captures how strong this feeling is compared to selfish incentives. An alternative interpretation of this model is that the cost of the disease is "socialized" to an extent, e.g., that nodes' health insurance rates increase if others catch the disease. The parameter $\beta$ then captures how steeply the individuals are penalized for others' diseases.

Interestingly, in the inoculation game with altruism, pure Nash Equilibria need not always exist[2], as we show in Section 2. Even when they do exist, Nash Equilibria can sometimes be as bad as in the model without altruism. Interestingly, however, while the Price of Stability [2] without altruism can also be $\Theta(n)$, a similar notion improves dramatically with the altruistic model.

Since Nash Equilibria may not exist, the notion of "Price of Stability" does not apply directly. Mixed Nash Equilibria are not a natural solution concept here, as vaccination decisions tend to be permanent or very long-term. We therefore instead consider an "Opt-Out" dynamic, and correspondingly define the *Price of Opting Out*. A benevolent authority suggests an initial vaccination assignment $S_0$. The nodes in $S_0$ can choose to opt out of being vaccinated, in any order. However, no node $v \notin S_0$ may opt to become vaccinated. (Precise definitions are given in Section 2.3.) This models a scenario in which individuals may choose to avoid being vaccinated due to various concerns, but the authority will not revise plans vis-à-vis nodes not originally included in the vaccination plan. Our main theorem (stated formally and proved in Section 3) is then the following:

THEOREM 1.2. *The Price of Opting Out is at most $1/\beta$.*

Thus, in a sense of somewhat limited autonomy among the nodes, our theorem establishes a $1/\beta$ bound on the social inefficiency introduced by individual nodes' decisions. Together with the $\Theta(n)$ bounds for Price of Anarchy, and Price of Stability without altruism, this result can be interpreted as saying that coordination of vaccination strategies or socialization of healthcare costs alone may not lead to societally desirable outcomes if individual nodes can override suggestions. However, the combination of both, i.e., socializing costs *and* starting with a carefully chosen assignment, may lead to significantly more desirable outcomes, even when individuals get to override the suggested vaccination strategies.

Naturally, it would be desirable to strengthen Theorem 1.2 to the outcomes after arbitrary best-response dynamics. Since best-response dynamics may cycle for inoculation games with partial altruism (see Section 2), we cannot focus on stable states alone, but would have to consider all states reachable via best-response dynamics. Notice that such a result would be much stronger than traditional Price of Stability bounds.

## 1.1 Related Work

A number of recent studies have analyzed the spread of worms or viruses on Internet-like topologies by focusing on characterizing the epidemic threshold (the transmission rate at which the disease goes from dying out quickly to infecting a large share of the network) for models such as small-world graphs [23] and preferential attachment models [5, 17]. The epidemic threshold is related to graph properties such as degree distribution, spectral radius and isoperimetric constants [7]. Based on these observations, Dezső and Barabási [8] suggest the vaccination of high-degree nodes in power-law random graphs as a way of increasing the epidemic threshold and thereby reducing the spread of diseases. Similar heuristics with analysis under random graph models with given degree distributions are also presented in [15]

The model of Aspnes et al. [4] has been extended in several ways. As discussed above, Meier et al. [20] consider the addition of friendship. Moscibroda et al. [21] instead consider

[2]As far as we are aware, the authors of [20] do not establish whether the model with friendship always guarantees pure Nash Equilibria.

malicious Byzantine players who may misrepresent their actions with an intent to harm other players. (Naturally, this model is more suited to computer networks than social networks.) Perhaps surprisingly, such *malice* can sometimes lead to societally more desirable outcomes, due to the fear of other players. Recently, Diaz et al. [9] showed that the same "windfall of malice" can be achieved with a *mediator*. A mediator is a trusted third party that suggests actions to each player; the players retain free will and can ignore the mediator's suggestions.

Besides the random outbreak model, several other natural models of disease outbreaks have been studied in the recent literature. If the source node $s$ (or node set) of the initial outbreak is known, and the goal is to minimize the number of infected nodes, the problem is equivalent to the MIN-SIZE BOUNDED CAPACITY CUT problem studied by Hayrapetyan et al. [14]. They gave a $(1/\lambda, 1/(1-\lambda))$ bicriteria approximation algorithm for this problem, which — up to constant factors — is subsumed by our bicriteria approximation result. If the goal is instead to maximize the number of uninfected nodes, i.e., the number of nodes *not* in a component with $s$, the problem is called MAX-SIZE BOUNDED CAPACITY CUT [14, 22]. At optimality, the two problems are equivalent, but the known approximation results for MAX-SIZE BOUNDED CAPACITY CUT are weaker.

If there is a timing component to the infection process, i.e., in each time step, the algorithm can vaccinate $k$ nodes, and the infection spreads one hop in the network, then the problem is known as the FIREFIGHTER problem [1]. For different optimization versions of this problem, Anshelevich et al. [1] recently analyzed approximation algorithms and hardness results.

## 2. PRELIMINARIES

We first describe our generalization of the basic model of Aspnes et al. [4]. Our generalization allows the vaccination cost, infection cost, and probability of initial infection to vary among nodes. We then extend the model to include a notion of altruism. We show that with altruism, there are instances without pure Nash Equilibria. We therefore propose a notion of "opting out" from vaccinations, and define the Price of Opting Out.

### 2.1 Basic Model

The social or computer network is represented by an undirected graph $G = (V, E)$ of $n$ nodes, each of which can either be vaccinated or unvaccinated.

If node $v$ chooses to vaccinate, its cost is $C_v$. $C_v$ combines such factors as monetary cost, side effects, pain, loss of income, inertia, etc. In other words, it subsumes all factors that may cause $v$ to not want to be vaccinated.

Once all vaccination decisions are made, one node becomes infected. Node $v$ is chosen as the initially infected node with probability $p_v$. (Thus, we assume that $\sum_v p_v = 1$.) We let $z$ denote the size of the support of $p$, i.e., the number of nodes $v$ with $p_v > 0$.

From the initially infected node, the infection spreads along edges of the graph to all unvaccinated nodes. However, no vaccinated nodes can become infected or pass on the infection. Let $S$ be the set of vaccinated nodes, and $\Gamma_1, \ldots, \Gamma_k$ the connected components of $G \setminus S$. If node $v$ is unvaccinated and in component $\Gamma_i$, its probability of infection is $\sum_{u\in\Gamma_i} p_u$. The cost for becoming infected is $L_v$, leading to an expected cost of $L_v \cdot \sum_{u\in\Gamma_i} p_u$ if node $v$ chooses to stay unvaccinated. As with $C_v$, $L_v$ captures all factors (such as pain, monetary cost, etc.) that make becoming infected undesirable for $v$.

Since each node $v$ in component $\Gamma_i$ has expected cost $L_v \cdot \sum_{u\in\Gamma_i} p_u$, the total social cost with set $S$ vaccinating is

$$P(S) = \sum_{v\in S} C_v + \sum_i \sum_{u\in\Gamma_i} p_u \sum_{v\in\Gamma_i} L_v. \quad (1)$$

We use $S^*$ to denote the optimum set of nodes to vaccinate, i.e., the set minimizing $P(S)$.

While it is socially optimal to minimize $P(S)$, individual nodes' preferences may not align with this objective. An individual node will choose the strategy (be vaccinated or not be vaccinated) based only on its own tradeoff. That is, a selfish node will vaccinate if $C_v \le L_v \cdot \sum_{u\in\Gamma_i} p_u$, and not vaccinate otherwise. Let

$$p_v(S) = \begin{cases} C_v & \text{if } v \in S \\ L_v \cdot \sum_{u\in\Gamma_i} p_u & \text{if } v \notin S, v \in \Gamma_i \end{cases} \quad (2)$$

denote the cost that node $v$ experiences based on all players' vaccination decisions. Since players will act selfishly, this scenario leads naturally to a game termed the *inoculation game*.

Aspnes et al. [4] already established that there always is a pure-strategy Nash Equilibrium in this setting[3], and showed a linear lower bound on the Price of Anarchy. Recall that the Price of Anarchy [16] is defined as the ratio between the social cost at the worst Nash Equilibrium and at the social optimum. By contrast, the Price of Stability [2] is the ratio between the social cost at the *best* Nash Equilibrium and at the social optimum.

PROPOSITION 2.1. *Both the Price of Anarchy and the Price of Stability can be $\Theta(n)$.*

**Proof.** A simple example is a star graph with $C_v = C$, $L_v = L$ and $p_v = 1/n$ for every node $v$, and $C = L + \epsilon$. In the unique Nash Equilibrium, no player vaccinates, while the socially optimal solution vaccinates the center node of the star. The respective social costs are $nL$ and $C + (1-1/n) \cdot L$. ■

### 2.2 Altruism

The basic model introduced above assumes that individuals are completely selfish and do not take into account the effects of their actions on other nodes. This assumption is frequently found to be violated in public goods and other experiments [18, 19], where participants act somewhat altruistically or spitefully, even in single-shot experiments with complete strangers. We therefore study the inoculation game under a model of partial altruism. Our model is based on a suggestion of Ledyard [18, p. 154] and has been studied before in the context of traffic routing [6]. The altruism level of the nodes is denoted by $\beta$.

DEFINITION 2.2 (PERCEIVED COST). *The* perceived cost *of a $\beta$-altruistic node is the convex combination*

$$p_v^{(\beta)}(S) = (1-\beta) \cdot p_v(S) + \beta \cdot P(S). \quad (3)$$

[3]Strictly speaking, their results were for uniform probabilities $p_v$ and costs $C_v, L_v$, but they carry over directly.

Thus, the perceived cost of a partially altruistic node is the convex combination of the individual cost (selfish part) and the social cost (altruistic part). As usual, a node will choose the strategy (vaccinate or do not vaccinate) which minimizes the perceived cost. The tradeoff is characterized by the following proposition, which can be obtained by simple rearranging.

Proposition 2.3. *Let $S \subseteq V \setminus \{v\}$ be the set of other nodes vaccinating, and $\Gamma$ the component of $G \setminus S$ containing $v$. Let $\Gamma_1, \ldots, \Gamma_k$ be the subcomponents of $\Gamma$ resulting from removing $v$ from $\Gamma$. Then, $v$ will prefer to be vaccinated if and only if*

$$\begin{aligned} C_v &\leq (1-\beta)\sum_{u\in\Gamma} p_u \cdot L_v \\ &+ \beta \cdot (\sum_{u\in\Gamma} p_u \sum_{u'\in\Gamma} L_{u'} - \sum_i \sum_{u\in\Gamma_i} p_u \sum_{u'\in\Gamma_i} L_{u'}). \end{aligned}$$

Remark 2.4. *Our definition of perceived cost is similar to the notion of* friendship *used by Meier et al. [20]. In their case, the altruistic part does not consider the cost of* all *nodes, but just that of the* neighbors *of $v$ in $G$. Thus, they model more the incentives due to friendship in a social network, while our model captures more a general notion of altruism toward all others.*

*By interpreting the altruistic term as a* monetary *cost (instead of perceived cost due to altruism), we can also consider our model as one of* socialized health care costs*: nodes in the network have to pay a $\beta$ fraction of the cost incurred by other nodes, e.g., in the form of health care premiums. This payment should provide additional incentives for nodes to be vaccinated, as spreading a disease to others will eventually lead to higher costs for them as well.*

Proposition 2.5. *There is an instance such that for every $\beta$, the Price of Anarchy is $\Theta(n)$.*

**Proof.** For the complete bipartite graph $K_{2,n-2}$, if $C_v = C$, $L_v = L$, $p_v = 1/n$ for every node $v$, and $C = L + \epsilon$, the state in which no node is vaccinated is a Nash Equilibrium regardless of the value of $\beta$. The calculation of the ratio is the same as in Proposition 2.1. (Notice that for $\beta$ large enough, the state with both nodes on the left side vaccinated is also a Nash Equilibrium, and the Price of Stability is thus smaller than $\Theta(n)$.) ■

While the inoculation game with selfish players is a potential game and thus possesses pure Nash Equilibria [4], the introduction of partial altruism changes the situation significantly.

Proposition 2.6. *There exist instances of the inoculation game with partial altruism in which there is no pure Nash Equilibrium.*

**Proof.** Consider the graph in Figure 1 with two nodes $u, v$ and cliques of the indicated sizes. Whenever an edge is shown, there is an edge from $u$ (or $v$) to *all* nodes in the corresponding clique. On the left, there are 14 cliques of size 10000 each. The altruism value is $\beta = \frac{1}{4}$.

It is a fairly straightforward calculation to show that
(1) no node besides $u$ or $v$ ever wants to be vaccinated,
(2) $v$ wants to be vaccinated if and only if $u$ is vaccinated,

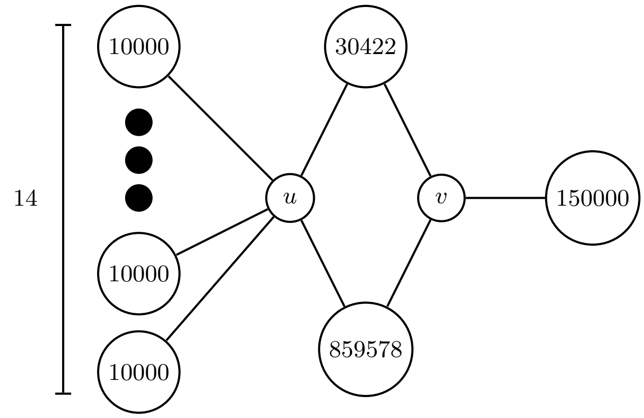


**Figure 1: A graph without Nash Equilibrium**

and (3) $u$ wants to be vaccinated if and only if $v$ is *not* vaccinated.
Thus, this instance encodes a "Matching Pennies" type of game and has no pure Nash Equilibrium. ■

## 2.3 Opting Out

In light of Proposition 2.6, the standard notion of the Price of Stability [2] is not defined for pure Nash Equilibria in our game. Mixed Nash Equilibria are not a natural solution concept in inoculation games, because the decision of whether or not to be vaccinated tends to be permanent or very long-term. Nevertheless, it is of interest to analyze the effect that nodes' autonomy with respect to vaccination decisions has on the social cost.

Since the main concern with individual autonomy is *undervaccination* of the network (see the discussion at the end of Section 3), we consider a natural model of *opting out*. A benevolent authority suggests a set of nodes $S_0$ to vaccinate, such as the optimal solution $S^*$ or an approximation. Subsequently, nodes that were targeted for vaccination have the option to override this decision, e.g., by not showing up for their vaccination. However, we do not allow nodes $v \notin S_0$ to override the decision and become vaccinated instead. Since the resulting dynamic is monotone in the number of vaccinated nodes, it will always converge to a final set of vaccinated nodes. However, this set of nodes may depend on the order in which nodes opt out.

For a starting set $S_0$, we define $\mathcal{R}(S_0)$ to be the collection of all node sets $S$ such that the opting-out dynamic, starting from $S_0$, will eventually reach $S$. Formally, we define $\mathcal{R}(S_0)$ as follows:

Definition 2.7. *(1) $S_0 \in \mathcal{R}(S_0)$, and (2) If $S \in \mathcal{R}(S_0)$, and $v \in S$ prefers being unvaccinated, given that all nodes in $S \setminus \{v\}$ are vaccinated, then $S \setminus \{v\} \in \mathcal{R}(S_0)$.*

Note that we do not require each $S \in \mathcal{R}(S_0)$ itself to be stable; we also include in $\mathcal{R}(S_0)$ sets such that in later steps, further nodes will opt out. We then define the *Price of Opting Out* to be the worst-case ratio between the social cost of any set $S \in \mathcal{R}(S^*)$ and the social cost at $S^*$. Thus, the Price of Opting Out captures the increase in cost due to giving nodes the authority to opt out of vaccinations.

Our notion of the Price of Opting Out bears some similarity with the "Price of Sinking" defined by Goemans et al. [13] in the context of routing games and valid utility games. However, they considered the strongly connected sink component of the best-response graph, and considered bounds on the expected cost under the stationary distribution of a random walk. Our goal is to obtain bounds on *each* reachable state. Naturally, it is a question for future work to consider not only opt-out dynamics, but states reached by arbitrary sequences of best responses.

# 3. THE PRICE OF OPTING OUT

In this section, we bound the price of opting out, by analyzing any state that can be reached from an initially vaccinated set $S_0$ by a sequence of opt-out moves. Our main theorem is the following:

THEOREM 3.1. *If $S$ is obtained from $S_0$ by a sequence of opt-out moves, then $P(S) \leq \frac{1}{\beta} \cdot P(S_0)$.*

Theorem 3.1 thus in a sense captures the Price of Limited Autonomy: letting individuals choose not to be vaccinated when an optimal (or near-optimal) solution prescribes that they should be.

**Proof.** Let $\{v_1, \ldots, v_\ell\} = S_0 \setminus S$ be the set of all nodes who have opted out of being vaccinated, in the order in which they opted out. Let $S_t = S \cup \{v_{t+1}, \ldots, v_\ell\}$ be the set of nodes still vaccinated after $t$ nodes have opted out, and $\Gamma_1^{(t)}, \ldots, \Gamma_{k_t}^{(t)}$ the connected components of $G \setminus S_t$. In particular, $\Gamma_1^{(0)}, \ldots, \Gamma_{k_0}^{(0)}$ are the connected components of the initial vaccinated set $S_0$. Define

$$\Phi(t) \quad := \quad \sum_{v \in S_t} C_v + \beta \cdot \sum_i \sum_{u \in \Gamma_i^{(t)}} p_u \sum_{u' \in \Gamma_i^{(t)}} L_{u'}.$$

We will prove by induction that for all $t$, we have

$$\Phi(t) \quad \leq \quad P(S_0). \tag{4}$$

The base case $t = 0$ holds because $\beta \leq 1$, simply substituting the definition of social cost. Consider step $t$ in which node $v = v_t$ decides to opt out of vaccinating. By ceasing to be vaccinated, $v$ merges one or more components $\Gamma_i^{(t)}$ for $i \in M$, forming one new component $\Gamma_j^{(t+1)}$. Then,

$$\begin{aligned} \Phi(t+1) \quad &= \quad \Phi(t) - C_v + \beta \cdot \left(\sum_{u \in \Gamma_j^{(t+1)}} p_u \sum_{u' \in \Gamma_j^{(t+1)}} L_{u'}\right. \\ &- \quad \left.\sum_{i \in M} \sum_{u \in \Gamma_i^{(t)}} p_u \sum_{u' \in \Gamma_i^{(t)}} L_{u'}\right). \end{aligned}$$

By Proposition 2.3, the fact that $v$ chooses to opt out of vaccinating implies that

$$\begin{aligned} C_v \quad &\geq \quad (1-\beta) \sum_{u \in \Gamma_j^{(t+1)}} p_u \cdot L_v \\ &\quad +\beta \cdot \left(\sum_{u \in \Gamma_j^{(t+1)}} p_u \sum_{u' \in \Gamma_j^{(t+1)}} L_{u'}\right. \\ &\qquad \left.- \sum_{i \in M} \sum_{u \in \Gamma_i^{(t)}} p_u \sum_{u' \in \Gamma_i^{(t)}} L_{u'}\right) \\ &\geq \quad \beta \cdot \left(\sum_{u \in \Gamma_j^{(t+1)}} p_u \sum_{u' \in \Gamma_j^{(t+1)}} L_{u'}\right. \\ &\qquad \left.- \sum_{i \in M} \sum_{u \in \Gamma_i^{(t)}} p_u \sum_{u' \in \Gamma_i^{(t)}} L_{u'}\right). \end{aligned}$$

Substituting this inequality into $\Phi(t+1)$ shows that $\Phi(t+1) \leq \Phi(t)$, and the claim now follows by induction.

After all $v_t$ have opted out, the total social cost is

$$\begin{aligned} P(S) \quad &= \quad \sum_{v \in S} C_v + \sum_i \sum_{u \in \Gamma_i^{(\ell)}} p_u \sum_{u' \in \Gamma_i^{(\ell)}} L_{u'} \\ &\leq \quad \frac{1}{\beta} \cdot \Phi(\ell) \\ &\leq \quad \frac{1}{\beta} \cdot P(S_0), \end{aligned}$$

where the last step followed from the claim we proved by induction. ■

By applying Theorem 3.1 to the optimum set $S^*$, we obtain the following corollary:

COROLLARY 3.2. *The Price of Opting Out is at most $1/\beta$.*

Notice, however, that Theorem 3.1 is more general. In particular, it applies to approximately optimal starting sets $S_0$. While computing $S^*$ itself is NP-complete, we show in Section 4 how to find an $O(\log z)$ approximation to the social cost. Theorem 3.1 then guarantees that if we start with the set $S_0$ vaccinated, after allowing the nodes to opt out, the social cost will be within a factor $O(1/\beta \cdot \log z)$ of optimal, and stable to further opting out of nodes.

We also remark that Theorem 3.1 is tight.

PROPOSITION 3.3. *There are instances with altruism $\beta$ where the Price of Opting Out is $1/\beta$.*

**Proof.** Let $\beta > 0$ be arbitrary, and consider a star graph with $n$ nodes. Let $p_v = 1/n$, $L_v = 1$ and $C_v = 1 + \beta(n + 1/n - 2)$ for every node $v$. Then, an easy calculation shows that opting out always leads to a solution with no nodes vaccinated, for a total cost of $n$. On the other hand, the optimum solution vaccinates the center node of the star, giving a total cost of $1+\beta(n+1/n-2)+1-1/n = \beta n+O(1)$. As $n \to \infty$, the Price of Opting Out then converges to $1/\beta$. ■

Instead of the Price of Opting Out, one could study the Price of Opting In. There are two reasons why this is not as natural an approach: (1) It is less realistic that one could force individuals to undergo vaccinations, and (2) The Price of Opting In is always 1. For an easy calculation shows that if $\beta$-altruistic nodes prefer to switch their status to "vaccinated", the social cost always decreases.

Naturally, the most general dynamic one would want to study would combine opting in and opting out, considering any sequence of best-response steps. Ideally, one would want to prove a $1/\beta$ bound for *all* states reachable by such best responses. However, analyzing a full best-response dynamic in this sense appears to be quite challenging, and it is possible that it reaches states much less efficient than $1/\beta \cdot P(S^*)$.

# 4. APPROXIMATING SOCIAL COST

In this section, we present an improved approximation algorithm for the problem of minimizing the social cost. The previous best algorithm was an $O(\log^{1.5} n)$ approximation due to Aspnes et al. [4]. Their algorithm used repeated greedy invocations of the Sparsest Cut algorithm of Arora et al. [3], and was thus not particularly practical.

THEOREM 4.1.

1. *There is a polynomial-time $(O(1), O(\log z))$ bicriteria approximation algorithm for the problem of minimizing the expected cost of infected nodes, subject to a constraint on the cost of vaccinated nodes. The $O(1)$ factor applies to the cost of infected nodes, while the $O(\log z)$ applies to the cost of vaccinated nodes.*

2. *There is a polynomial-time $O(\log z)$ approximation algorithm for the problem of minimizing the expected social cost.*

Both of our algorithms are based on rounding the following natural linear program for the first problem (i.e., minimizing the cost of infected nodes subject to the constraint that the cost of vaccinated nodes be at most $\gamma$):

$$
\begin{array}{lll}
\text{Minimize} & \sum_{u\in V} p_u \sum_{v\in V}(x_{u,v}\cdot L_v) & \\
\text{subject to} & \sum_{v\in V}(y_v\cdot C_v)\le\gamma & \\
& x_{u,u}=1-y_u & \text{for all } u\\
& x_{u,v}\ge x_{u,w}-y_v & \text{for all } u,\ \text{edges } (w,v)\\
& y_v\in\{0,1\} & \text{for all } v\\
& x_{u,v}\ge 0 & \text{for all } u,v
\end{array}
$$

In the above Integer Linear Program, $y_v$ is a decision variable encoding whether node $v$ is vaccinated/removed ($y_v = 1$) or not ($y_v = 0$). $x_{u,v}$ is 1 if there is a path between $u$ and $v$ avoiding all vaccinated nodes (and 0 if not). That is, $x_{u,v}$ is 1 iff an infection starting at $u$ would also infect $v$. $p_u$ is the probability that node $u$ will be initially infected. Notice that the objective function captures the expected cost of infected nodes by linearity of expectation.

The first constraint limits the total cost of nodes that can be vaccinated. The second and third constraint, together with the minimization objective, ensure that the $y_v$ and $x_{u,v}$ are consistent, i.e., that they satisfy the intended definition of the $x_{u,v}$.

Note that when $C_v = C$ and $L_v = L$ for all nodes $v$, and all nodes have a uniform probability of infection, we recover the sum-of-squares partitioning problem of Aspnes et al. [4], namely, to minimize the expected number of infected nodes given a constraint on the number of vaccinated nodes.

## 4.1 The LP Rounding Algorithm

As usual, we relax the integrality constraint to $y_v \in [0, 1]$, so that the LP can be solved in polynomial time. The $y_v$ can then be interpreted as the "lengths" of nodes. From these lengths of nodes, we can then use shortest path distances to define a hemimetric $d$ (satisfying non-negativity and the triangle inequality), by setting

$$d_{u,v} \quad := \quad \min_{P \text{ is a } u-v \text{ path}} \sum_{w\in P, w\neq v} y_w.$$

With this definition, we then have

$$x_{u,v} \quad = \quad \max(0, 1-d_{u,v}-y_v). \tag{5}$$

Our goal is to use (a modification of) the region-growing techniques of Garg, Vazirani, and Yannakakis [11] to round the fractional LP values in order to obtain a set of vaccinated nodes. The hemimetric we defined is akin to a "spreading metric" [10], and our analysis bears a lot of similarity with that of [10, 11]. We define the ball around $u$ of radius $r$ as $B_u(r) := \{v \mid d_{u,v} \le r\}$, and its *boundary* (the set of nodes partially inside $B_u(r)$) as $\delta_u(r) := \{v \mid d_{u,v} \le r \le d_{u,v}+y_v\}$. Notice that for each radius $r \le 1$, the set $\delta_u(r)$ forms a cut separating $u$ from all nodes $v$ with $d_{u,v} > r$.

Next, we define the volume of a ball. The fraction to which a node $v$ is inside a ball of radius $r$ is defined as $q_v = 1$ if $d_{u,v} + y_v \le r$, $q_v = 0$ if $d_{u,v} > r$, and $q_v = \frac{r-d_{u,v}}{y_v}$ if $d_{u,v} \le r \le d_{u,v} + y_v$. Then, the *volume* of a ball is $V(B_u(r)) := \sum_{v\in B_u(r)} y_v q_v C_v$.

Let $Z$ be the support of the distribution $p_v$, i.e., the set of all $v$ such that $p_v > 0$, and $z := |Z|$. With these definitions in place, we can now define the rounding algorithm.

**Algorithm 1** LP Rounding Algorithm

```
Solve the LP to obtain the hemimetric d_{u,v}.
while V(G) ∩ Z ≠ ∅ do
    Let u ∈ V(G) ∩ Z be an arbitrary node with p_u > 0.
    Choose r_u ≤ ρ to minimize (Σ_{v∈δ_u(r_u)} C_v) / (V(B_u(r_u)) + V(B_u(ρ))/z).
    Set S := S ∪ δ_u(r_u).
    Remove from G all of B_u(r_u) and δ_u(r_u).
end while
Output the set S of nodes to be vaccinated.
```

Notice that $\rho < \frac{1}{2}$ is a parameter that can be used to trade off between the violations of the vaccination constraint and infection objective.

The key step in analyzing Algorithm 1 is a region growing lemma for node-weighted multicuts. Such a lemma has been known as folklore and is occasionally attributed to [12], although no proof appears there.

LEMMA 4.2. *For any node $u$, any $\nu > 0$, and radius bound $\rho$, there is a radius $r_u \le \rho$ such that*

$$\sum_{v\in\delta_u(r)} C_v \quad \le \quad \frac{1}{\rho}\cdot\ln\left(\frac{V(B_u(\rho))+\nu}{\nu}\right)\cdot(V(B_u(r_u))+\nu).$$

*Furthermore, $r_u$ can be found in polynomial time by examining all distances $r \le \rho$ such that $r = d_{u,v}$ for some node $v$, and retaining the one minimizing $\frac{\sum_{v\in\delta_u(r)} C_v}{V(B_u(r))+\nu}$.*

**Proof.** We assume for contradiction that the search over all $r \le \rho$ fails to find a radius $r_u$ meeting the claim of the lemma. Since the sets only change at discrete points equaling $d_{u,v}$ for some $v$, this implies that the claimed inequality fails to hold for all $r \le \rho$.

First, notice that $\frac{d}{dr}V(B_u(r)) = \sum_{v\in\delta_u(r)} C_v$ for all $r$, so the assumption implies that $\frac{\frac{d}{dr}V(B_u(r))}{V(B_u(r))+\nu} > \frac{1}{\rho}\cdot\ln\left(\frac{V(B_u(\rho))+\nu}{\nu}\right)$ for all $r \in [0, \rho]$. By taking an integral over $r$ from 0 to $\rho$ on both sides, we obtain

$$
\begin{array}{rcl}
\int_0^\rho \frac{\frac{d}{dr}V(B_u(r))}{V(B_u(r))+\nu}dr & > & \int_0^\rho \frac{1}{\rho}\cdot\ln\left(\frac{V(B_u(\rho))+\nu}{\nu}\right)dr\\
& = & \ln\left(\frac{V(B_u(\rho))+\nu}{\nu}\right).
\end{array}
$$

But the lefthand side evaluates to $[\ln(V(B_u(r)) + \nu)]_0^\rho = \ln\left(\frac{V(B_u(\rho))+\nu}{\nu}\right)$, which gives a contradiction, thus completing the proof. ■

Using this lemma, we can now prove the following bicriteria approximation guarantee:

THEOREM 4.3. *When Algorithm 1 terminates, it satisfies the following two properties:*

1. *The total cost of vaccinated nodes in $S$ is at most $\frac{2\ln(z+1)}{\rho}\cdot\gamma$.*
2. *The expected cost of infected nodes is at most $\frac{1}{1-2\rho}$ times the objective function value of the LP.*

**Proof.** 1. Consider the set of nodes $\delta_u(r_u)$ removed (i.e., chosen for vaccination) in some iteration. Note that the number of iterations in the rounding algorithm is upper-bounded by $z$ (and $z = n$ if all nodes have a nonzero probability of infection). By applying Lemma 4.2 with $\nu = V(B_u(\rho))/z$, we obtain that

$$\sum_{v\in\delta_u(r)} C_v \quad \le \quad \frac{\ln(z+1)}{\rho}\cdot\left(V(B_u(r_u))+\frac{V(B_u(\rho))}{z}\right).$$

Let $U$ be the set of all nodes which were chosen as centers of balls at some point during the algorithm. To upper-bound the total vaccination cost of removed nodes, we sum over all nodes $u \in U$. Since all nodes in $B_u(r_u)$ and $\delta_u(r_u)$ are removed from $G$, their weight is only counted for one region, and we obtain that

$$\begin{aligned} \sum_{v \in S} C_v &\leq \frac{\ln(z+1)}{\rho} \cdot \sum_{u \in U} (V(B_u(r_u)) + \frac{V(B_u(\rho))}{z}) \\ &\leq \frac{2\ln(z+1)}{\rho} \cdot \gamma. \end{aligned}$$

In the last inequality, we used that $V(B_u(\rho)) \leq \gamma$ for all $u$, and that $\sum_u V(B_u(r_u)) \leq \gamma$ by disjointness of the regions. Both of these inequalities follow immediately from the first constraint of the LP.

2. To analyze the objective function, consider any two nodes $v, w$ in the same component of $G \setminus S$. Thus, both $v$ and $w$ are in the same ball $B_u(r_u)$ for some $u \in U$.

   Because neither node is vaccinated, neither is in $\delta_u(r_u)$, so we obtain that $d_{u,v} + y_v \leq r_u \leq \rho$ and $d_{u,w} + y_w \leq r_u \leq \rho$. Because $d_{u,v} + y_v = d_{v,u} + y_u$, this implies that $d_{v,u} \leq \rho$, so by triangle inequality, $d_{v,w} \leq d_{v,u} + d_{u,w} \leq 2\rho - y_w$. Together with Equality (5), this implies that $x_{v,w} = \max(0, 1 - d_{v,w} - y_w) \geq 1 - 2\rho$. Thus, the objective function after rounding is

$$\begin{aligned} & \sum_{u \in U} \sum_{v \in B_u(r_u) \setminus \delta_u(r_u)} p_v \sum_{w \in B_u(r_u) \setminus \delta_u(r_u)} L_w \\ \leq\; & \frac{1}{1-2\rho} \cdot \sum_{u \in U} \sum_{v \in B_u(r_u) \setminus \delta_u(r_u)} p_v \\ & \qquad \sum_{w \in B_u(r_u) \setminus \delta_u(r_u)} x_{v,w} \cdot L_w \\ \leq\; & \frac{1}{1-2\rho} \cdot \sum_{v \in V} p_v \sum_{w \in V} x_{v,w} \cdot L_w \\ =\; & \frac{1}{1-2\rho} \cdot \mathrm{OPT}_{\mathrm{LP}}. \end{aligned}$$

   This proves the approximation guarantee for the objective function. ■

Notice that in the special case that $z = 1$, our problem becomes the MinSBCC problem [14]: To minimize the number of nodes on the $s$-side of a cut, given a budget for vaccinations. For this problem, [14] had presented a $(2, 2)$-bicriteria approximation algorithm. With $\rho = \frac{1}{4}$, we obtain a $(4\ln(2), 2) \approx (2.77, 2)$ bicriteria approximation, thus nearly recovering the specialized result of [14].

## 4.2 Minimizing the Total Social Cost

We next show how to modify Algorithm 1 to minimize the objective function of social cost as defined in Section 2. To do so, we change the objective function of the LP to "Minimize $\sum_{u \in V} p_u \sum_{v \in V} (x_{u,v} \cdot L_v) + \gamma$", and leave the constraints of the LP the same, treating $\gamma$ as a variable now. After solving the LP, we run Algorithm 1 on the resulting hemimetric, as before. Using Theorem 4.3, we then obtain the following:

COROLLARY 4.4. *The algorithm runs in polynomial time and approximates the expected social cost to within a factor* $\max(\frac{1}{1-2\rho}, \frac{2\ln(z+1)}{\rho}) = O(\log z)$.

**Proof.** Let $\gamma$ be the total vaccination cost of the fractional LP solution, and $\Lambda$ the total infection cost of the fractional LP solution. So the total LP cost is $\gamma + \Lambda$. Theorem 4.3 guarantees that for the rounded solution $S$, the total vaccination cost is $\sum_{u \in S} C_u \leq \frac{2\ln(z+1)}{\rho} \cdot \gamma$, and the total infection cost at most $\frac{1}{1-2\rho} \cdot \Lambda$. Adding these terms implies the claim. ■

## 4.3 Integrality Gap of the LP

In this section, we show that the bicriteria analysis in Theorem 4.3 is essentially tight, by giving a class of example graphs on which the integrality gap matches the approximation guarantee of Algorithm 1 up to constants, even in a bicriteria sense.

Let $G$ be a node expander graph of maximum degree $\Delta$, i.e., $G$ satisfies that any set $S$ with $|S| \leq n/2$ has at least $|S|$ neighbors. The existence of such a graph (for large enough, but constant, $\Delta$) can be proved easily using the probabilistic method. Set $p_v = \frac{1}{n}$ and $C_v = 1$ for all vertices $v$, and $\gamma = \frac{n}{2\Delta}$ for the integer LP.

Now consider a solution $S$ vaccinating at most $\gamma$ nodes, and let $\Gamma_1, \ldots, \Gamma_k$ be the connected components of $G \setminus S$. We will show that there is at least one $i$ such that $|\Gamma_i| \geq n/2$.

For assume that this were not the case. Then, the expansion property can be applied to each $\Gamma_i$, guaranteeing that each has at least $|\Gamma_i|$ neighbors. By definition, all neighbors of $\Gamma_i$ must lie in $S$, and by the degree bound, each node $v \in S$ can be the neighbor of at most $\Delta$ nodes $u \notin S$. So

$$\Delta|S| \;\geq\; \sum_i |\Gamma_i| \;=\; n - |S|,$$

or $|S| \geq \frac{1}{\Delta+1} \cdot n$, which contradicts the assumption that $|S| \leq \gamma = \frac{1}{2\Delta} \cdot n$. Therefore, there must be at least one component $\Gamma_i$ of size at least $n/2$, and the expected number of infected nodes is at least $n/4$. (With probability at least $\frac{1}{2}$, at least $n/2$ nodes get infected.)

Next, we consider the objective value of a fractional solution to the LP. Since we want to prove bounds in a bicriteria sense, we tighten the first constraint by a factor $b \geq 1$, i.e., we require that $\sum_{v \in V} y_v \leq \gamma/b$. Consider the fractional solution which assigns uniformly $y_v = \frac{\gamma}{bn} = \frac{1}{2\Delta b}$. If nodes $u, v$ are $h$ hops away from each other in $G$, assign $x_{u,v} = \max(0, 1 - h \cdot \frac{1}{2\Delta b})$. This clearly defines a feasible solution to the tightened LP.

To evaluate the objective function, focus on one node $u$, and consider $\sum_v x_{u,v}$. We can group the nodes $v$ by increasing distance from $u$. By the degree bound, there are at most $\Delta^h$ nodes $v$ at distance $h$ hops from $u$. Furthermore, if $u$ and $v$ are at distance at least $2\Delta b$, we have that $x_{u,v} = 0$ in the above definition. We therefore obtain that

$$\begin{aligned} \sum_v x_{u,v} &\leq \sum_{h=0}^{2\Delta b} \Delta^h \cdot (1 - h \cdot \frac{1}{2\Delta b}) \\ &= \frac{\Delta^{1+2\Delta b} - 1}{\Delta - 1} + \frac{1}{2\Delta b} \cdot \frac{\Delta + \Delta^{1+2\Delta b}((\Delta-1)\cdot(2\Delta b) - 1)}{(\Delta-1)^2} \\ &= \Theta(\Delta^{2\Delta b}). \end{aligned}$$

Summing over all $n$ and multiplying by $p_u$ gives us that the fractional LP objective value is $\Theta(\Delta^{2\Delta b})$. Thus, the integrality gap of the objective function is $\Theta(\frac{n}{\Delta^{2\Delta b}})$. In particular, so long as $b = o(\log n)$, the integrality gap is polynomial in $z$ (and also polynomial in $n$). Thus, the factor $\Theta(\log z)$ incurred in the number of vaccinated nodes by our rounding algorithm is necessary so long as we desire a better-than-polynomial approximation in the expected number of infected nodes.

# 5. CONCLUSIONS

In this paper, we presented improved approximation algorithms and bounds on the Price of Opting Out for a network inoculation game, in the presence of altruism. We believe that the Price of Opting Out result in particular has interesting consequences in terms of policy: it suggests that if individuals have the freedom to opt out of suggested vaccinations, then neither coordination of strategies nor socialization of costs alone will lead to an efficient outcome, yet the combination of both gives outcomes of much lower societal cost.

Naturally, many questions remain for future work. Most directly, our Price of Opting Out result should be generalized (if possible) to a more general notion of autonomy. Since pure Nash Equilibria may not exist in general, a natural (and very strong) result would be to show that *all* states reachable from the optimum by any sequence of individual best responses have the same cost bound. Such sequences could include arbitrary choices to vaccinate or not to vaccinate matching nodes' preferences. Such a result would significantly strengthen the policy implications of our results, since in most scenarios, individuals do have the freedom to decide whether or not they want to be vaccinated. Furthermore, while mixed Nash Equilibria are not an ideal solution concept for the type of game we study here, it would nevertheless be interesting from a theory point of view what they look like, and how efficient or inefficient they are.

Stronger bounds could also be obtained under additional assumptions about the network structure. For instance, most social networks have bounded degrees. Indeed, we can show that even in the basic model of Aspnes et al. without altruism, the Price of Anarchy is bounded by $\sqrt{n\Delta}$ if all degrees are bounded by $\Delta$ (whereas the general bound is $\Theta(n)$). The exact impact on the Price of Opting Out or a generalization constitutes an interesting direction for future work. Similarly, it would be interesting to study the impact of other graph parameters.

The approximation guarantees we derive in Section 4 are essentially best possible for the particular linear program we use. However, it would be desirable to see whether there are corresponding lower bounds in terms of approximation hardness, or whether different techniques might lead to improved approximations.

More generally, the model proposed by Aspnes et al. (whose generalization we study here) is somewhat simplistic. It assumes that each edge of the graph will deterministically transmit the infection, and that each vaccination will deterministically protect the node. Assigning (known) probabilities to both types of events would be much more natural, but most likely lead to a significantly more difficult optimization problem.

Finally, our analysis assumes that all nodes know the full topology of the network. This is certainly not true in social networks, and it would be interesting to formulate a natural model of partial knowledge, and analyze its impact on the behavior of individuals in the network.

### *Acknowledgments*

We would like to thank Elliot Anshelevich, Shanghua Teng, and Mihalis Yannakakis for helpful discussions. Apologies to Venkat Guruswami (author of "Better Extractors for Better Codes?") and Adam Kalai (author of "Better Computers for Better People") for usurping their titles.

## 6. REFERENCES

[1] Elliot Anshelevich, Deeparnab Chakrabarty, Ameya Hate, and Chaitanya Swamy. Approximability of the firefighter problem: Computing cuts over time. In *Proc. 20th Intl. Symp. on Algorithms and Computation*, pages 974–983, 2009.

[2] Elliot Anshelevich, Anirban Dasgupta, Jon Kleinberg, Eva Tardos, Thomas Wexler, and Tim Roughgarden. The price of stability for network design with fair cost allocation. In *Proc. 45th IEEE Symp. on Foundations of Computer Science*, pages 295–304, 2004.

[3] Sanjeev Arora, Satish Rao, and Umesh Vazirani. Expander flows, geometric embeddings and graph partitioning. In *Proc. 35th ACM Symp. on Theory of Computing*, 2004.

[4] James Aspnes, Kevin L. Chang, and Alexksandr Yamposlkiy. Inoculation strategies for victims of viruses and the sum-of-squares partition problem. *Journal of Computer and System Sciences*, 72(6):1077–1093, 2006.

[5] Albert-László Barabási, Réka Albert, and Hawoong Jeong. Mean-field theory for scale-free random networks. *Physica A*, 272:173–187, 1999.

[6] P. Chen and D. Kempe. Altruism and selfishness in traffic routing. In *Proc. 9th ACM Conf. on Electronic Commerce*, pages 140–149, 2008.

[7] Fan Chung. *Spectral Graph Theory.* American Mathematical Society, 1997.

[8] Zoltán Dezső and Albert-László Barabási. Halting viruses in scale-free networks. *Physical Review E*, 65, 2002.

[9] Josep Díaz, Dieter Mitsche, Navin Rustagi, and Jared Saia. On the power of mediators. In *Proc. 5th Workshop on Internet and Network Economics (WINE)*, pages 455–462, 2009.

[10] Guy Even, Joseph Naor, Satish Rao, and Baruch Schieber. Fast approximate graph partitioning algorithms. *SIAM J. on Computing*, 28(6):2187–2214, 1999.

[11] Naveen Garg, Vijay V. Vazirani, and Mihalis Yannakakis. Approximate max-flow min-(multi)cut theorems and their applications. *SIAM J. on Computing*, 25:235–251, 1996.

[12] Naveen Garg, Vijay V. Vazirani, and Mihalis Yannakakis. Multiway cuts in node weighted graphs. *Journal of Algorithms*, 50:49–61, 2004.

[13] Michel X. Goemans, Vahab Mirrokni, and Adrian Vetta. Sink equilibria and convergence. In *Proc. 46th IEEE Symp. on Foundations of Computer Science*, pages 142–154, 2005.

[14] A. Hayrapetyan, D. Kempe, M. Pál, and Z. Svitkina. Unbalanced graph cuts. In *Proc. 13th European Symp. on Algorithms*, pages 191–202, 2005.

[15] Matthew O. Jackson. *Social and Economic Networks.* Princeton University Press, 2008.

[16] Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. In *Proc. 17th Annual Symp. on Theoretical Aspects of Computer Science*, pages 404–413. Springer, 1999.

[17] Ravi Kumar, Prabhakar Raghavan, Sridhar Rajagopalan, D. Sivakumar, Andrew Tomkins, and Eli Upfal. Stochastic models for the web graph. In *Proc. 41st IEEE Symp. on Foundations of Computer Science*, pages 57–65, 2000.

[18] John Ledyard. Public goods: A survey of experimental

resesarch. In J. Kagel and A. Roth, editors, *Handbook of Experimental Economics*, pages 111–194. Princeton University Press, 1997.

[19] David Levine. Modeling altruism and spitefulness in experiments. *Review of Economic Dynamics*, 1:593–622, 1998.

[20] Dominic Meier, Yvonne Anne Oswald, Stefan Schmid, and Roger Wattenhofer. On the windfall of friendship: Inoculation strategies on social networks. In *Proc. 9th ACM Conf. on Electronic Commerce*, pages 294–301, 2008.

[21] Thomas Moscibroda, Stefan Schmid, and Roger Wattenhofer. When selfish meets evil: Byzantine players in a virus inoculation game. In *Proc. 25th ACM Symp. on Principles of Distributed Computing*, pages 35–44, 2006.

[22] Zoya Svitkina and Eva Tardos. Min-max multiway cut. In *Proc. 7th Intl. Workshop on Approximation Algorithms for Combinatorial Optimization Problems*, 2004.

[23] Duncan Watts and Steven Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393:440–442, 1998.