To Delay or Not: Temporal Vaccination Games on Networks

Abhijin Adiga*, Srini Venkat* and Anil Vullikanti*†
*Biocomplexity Institute of Virginia Tech
†Department of Computer Science, Virginia Tech
Email: {abhijin,vsriniv,akumar}@vbi.vt.edu

***Abstract*—Interventions such as vaccinations or installing anti-virus software are common strategies for controlling the spread of epidemics and malware on complex networks. Typically, nodes decide whether to implement such an intervention independently, depending on the costs they incur. A node can be protected by *herd immunity*, if enough other nodes implement such an intervention, making the problem of determining strategic decisions for vaccination a natural game-theoretical problem. There has been a lot of work on vaccination and network security game models, but all these models assume the vaccination decisions are made at the start of the game. However, in practice, a lot of individuals defer their vaccination decision, and the reasons for this behavior are not well understood, especially in network models.**

In this paper, we study a novel repeated game formulation, which considers vaccination decisions over time. We characterize Nash equilibria and the social optimum in such games, and find that a significant fraction of vaccinations might be deferred, in general. This depends crucially on the network structure, and the information and the vaccination delay. We show that finding Nash equilibria and the social optimum are NP-hard in general, and we develop an approximation algorithm for the social optimum whose approximation guarantee depends on the delay.

I. Introduction

The spread of epidemics on social contact networks and malware in computer networks is commonly modeled by diffusion processes on networks, in which the infection spreads from a node to its neighbors. The threat of malware is becoming increasingly critical, as the number of connected devices grows rapidly. Similarly, infectious diseases continue to pose a significant public health problem, despite significant advances in medicines. There exist effective interventions to control the spread of epidemics and malware, e.g., by taking vaccinations in the case of diseases, and installing antivirus software patches in the case of malware. However, these interventions incur a certain cost for the individual (e.g., the economic cost of the vaccine or the patch). Further, an individual can get protected without any intervention if enough nodes in the network are protected— this is referred to as *herd immunity* in mathematical epidemiology, and is a natural setting for a game-theoretical analysis. This has been a very active area of research both in epidemiology and network security, e.g. [3], [2], [13], [4], [17], [14], [7], [10], [9], [19], [20].

An important limitation of all these approaches is that they typically consider vaccination decisions only at the start of the epidemic, in a simultaneous game setting, as discussed later in Section VI. In practice, very few people get vaccinated early in the season, despite active campaigns by public health agencies, such as the CDC. Vaccination rates increase as the epidemic spreads, and understanding this remains a big challenge. There has been some work on temporal vaccination strategies in the mathematical epidemiology literature, e.g., [19], but this is based on differential equation models, and does not consider realistic network structure. A recent work [11] studied the case of developing a vaccine distribution plan to suppress a pandemic on a network, thus considering a temporal graph. However, the paper focuses on centralized policymaking, while our paper is the first study of temporal vaccination decision making by individuals in the complex network.

Our main contributions are summarized below.

1) We develop a repeated game formulation, Temporal-Vaccination, which considers vaccination decisions at multiple time steps. We characterize the structure of Nash equilibria (NE) in such games, and find that they depend very crucially on the network structure and the delay parameters. Further, we show that NE need not always exist, and deciding if there exists a NE is NP-hard, in general.
2) Both the NE and optimal strategies exhibit interesting temporal structure. Even if vaccination decisions are allowed to be made at multiple time steps, we show that all decisions are made either at time 0 (before the start of the epidemic), or the first time $T > 0$ when the next round is played, if the source of the infection is known before T. Further, there can be significant variation in the number of nodes that choose to vaccinate at time T, instead of at time 0. Additionally, we find that even if no nodes are vaccinated at time 0 in the NE, the social optimum might choose to have a significant fraction at the start.
3) Computing the social optimum turns out to be a challenging stochastic optimization problem. We show that it is NP-complete, and develop an approximation algorithm which gives a strategy with cost at most $2T$ times the optimum, where T is as defined above.
4) We study the characteristics of NE in different real and synthetic networks empirically. We use best response strategies, which converge to NE very fast, in general. We find the number of nodes which get vaccinated at the

start is very sensitive to the ratio of the vaccination and infection costs. Further, a significant fraction of nodes defer their vaccination decision. We also find that high degree nodes appear to be more likely to get vaccinated initially.

II. Preliminaries

Setting. We consider the spread of a *highly infectious disease* on a graph $G = (V, E)$, modeled as a simple discrete time **SI** model of an epidemic, where S and I denote Susceptible and Infectious states (see, e.g., [1]). If a node v is infected at time t, all of its uninfected neighbors v' will be infected at time $t+1$, unless it is vaccinated at or before $t+1$. Let C_v^t denote the cost for node v to get vaccinated at time t. A node need not decide to get vaccinated at time $t=0$, but might defer the decision to a future time. If node v gets infected, we assume it incurs a cost L_v, such that that $L_v > C_v^t$ for all t, i.e., infection is costlier than vaccination. We assume that the vaccine has 100% efficacy and starts protecting the node immediately. We discuss these assumptions later in Section VII. We refer to Table I for some additional notation and definitions needed for the rest of the paper.

A. Multi-stage game formulation

We formally define TemporalVaccination as a multi-stage game. We denote a game instance by $(G, \mathcal{T}, \mathbf{C}, \mathbf{L})$, where $G = (V, E)$ is a graph on n vertices, $\mathcal{T} = \{t_0 = 0, \cdots, t_k\}$ is a set of times at which vaccination decision can be made, $\mathbf{C} = \{C_v^t \mid v \in V, t \in \mathcal{T}\}$ is the set of vaccination costs, and $\mathbf{L} = \{L_v \mid v \in V\}$ is the set of infection costs. Each node $v \in V$ is a player. The strategy function is denoted by Y. The strategy for node v at time t, given that the source of infection is s, is $Y(v, s, t) \in \{0, 1\}$. Since the source is not revealed at $t = 0$, we denote the strategy at time 0 by $Y(v, \cdot, 0)$. Let $\mathbf{Y}$ be the set of all strategy functions. Y_t corresponds to the strategies at time t, i.e., $\{Y(v, s, t) : v, s \in V\}$. Let $Y_{<t}$ correspond to the strategies until time t, i.e., $\{Y(v, s, t') : v, s \in V, t' < t\}$.

The TemporalVaccination$(G, \mathcal{T}, \mathbf{C}, \mathbf{L})$ game is played in the following manner:

1) At time $t = 0$, all the nodes play a simultaneous vaccination game to decide whether to get vaccinated or not. If node v gets vaccinated at this time, we denote this by $Y(v, \cdot, 0) = 1$. Note that the vaccine takes effect immediately.
2) A randomly chosen node $s \in V$ is selected to be the source of the epidemic. We assume that if $Y(s, \cdot, 0) = 1$, it remains immune, and the epidemic does not start. If $Y(s, \cdot, 0) = 0$, then s gets infected and the infection spreads to each uninfected neighbor in subsequent times. We also assume perfect information, so all nodes know the identity of source s, and the entire network.
3) For each $t = 1, 2, \ldots$, we have the following two steps:
   (a) If $t \in \mathcal{T}$, a simultaneous vaccination game is played at time t, and each node v decides whether to get vaccinated at this time or not— this is denoted by $Y(v, s, t) \in \{0, 1\}$, with 1 denoting vaccination.
   (b) Let I_{t-1} denote the set of nodes which are infected at time $t-1$. For each node $u \in I_{t-1}$, each uninfected neighbor $v \in N(u)$, $v \notin I_{t-1}$ will get infected at time t, unless v is already vaccinated or vaccinates at t. Recall that vaccination takes precedence over infection in our model. Define set $I_t = I_{t-1} \cup \{v : \ v \text{ gets infected at time } t\}$ to be the new set of infected nodes.
4) The game stops at time t if there are no more uninfected nodes that can be infected from their neighbors, and there are no more vaccination games to be played, i.e., $t' \notin \mathcal{T}$ for all $t' \geq t$. Each node v incurs cost L_v if it ever got infected, i.e., $v \in I_t$. It incurs cost $C_v^{t'}$ if it got vaccinated at time $t' \leq t$. The overall cost for node v is the expectation over all possible choices of the source.

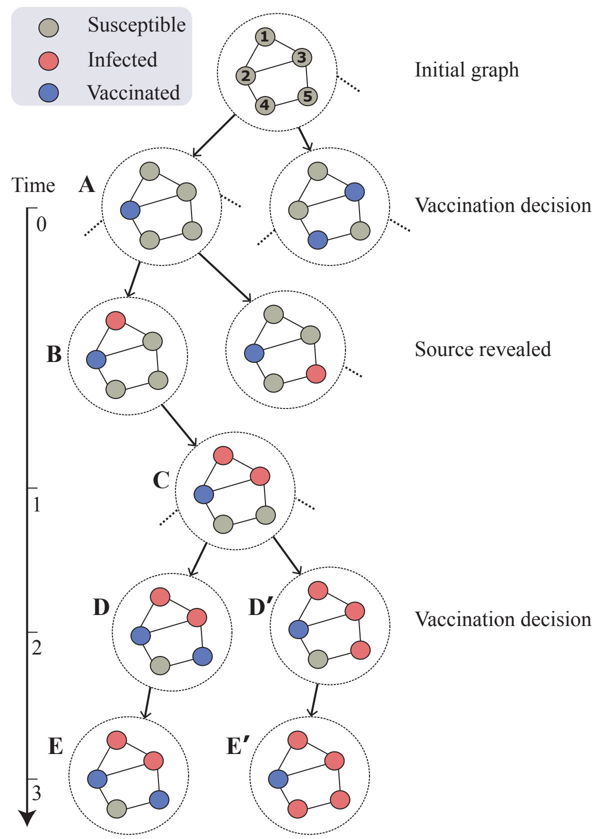

Fig. 1: Example illustrating the different rounds of the TemporalVaccination game on a graph with 5 nodes, with $\mathcal{T} = \{0, 2\}$. At $t = 0$, some nodes are vaccinated (blue). After this, the infection source is revealed (red). Since $1 \notin \mathcal{T}$, there is no vaccination decision (Step 3(a) of the game specification) at time 1, and instead only the epidemic spreading step (Step 3(b)) occurs. Next, since $2 \in \mathcal{T}$, vaccination decisions are taken, and thus both 3(a) and 3(b) occur. In configuration D node 5 chooses to vaccinate, while in D' it chooses not to vaccinate.

The configurations representing the states of players during the course of the TemporalVaccination game can be viewed in the form of a tree— note that this is not the same as the game tree in the extensive form representation, and we describe it here only for clarifying the structure of the game. The configuration at some time t of the game can be specified by the tuple $(\{Y_{<t}(v,s,t) : v \in V\}, I_t)$, where $\{Y_{<t}(v,s,t) : v \in V\}$ represents all the vaccination decisions in the game rounds played so far and I_t is the current set of infected nodes. The root node of this tree is the tuple (ϕ, ϕ), since initially there are no infections, and no vaccination game has been played. The root node has 2^n child configurations (the "level 1" configurations) of the form (Y_0, ϕ), each corresponding to a set $V' \subset V$ of nodes that get vaccinated at time 0. Each level-1 configuration has n children (the "level 2" configurations), $(Y_0, \{s\})$, each corresponding to a random source s. From this point on, the tree expands based on Step 3 of the game specification above. If $t \in \mathcal{T}$, then step 3(a) happens and there are 2^n child configurations, each corresponding to a subset $V' \subset V$ of nodes that get vaccinated at time t. After this, step 3(b) occurs, where the set I_t is updated and each configuration corresponds to $(Y_{<t}, I_t)$. The costs of all the nodes are specified for the leaf configurations of this tree, as discussed in Step 4 of the game specification.

Example. A fragment of the tree of configurations is shown in Figure 1. Configuration **A** can be written down as (Y_0, ϕ) where Y_0 is such that $Y(2, \cdot, 0) = 1$ and $Y(v, \cdot, 0) = 0$ for all $v \neq 2$. There are 2^n such configurations in level 1, which are children of the root (ϕ, ϕ). The child configuration marked **B** corresponds to the tuple $(Y_0, \{1\})$; there are n such configurations which are children of **A**. **C** represents a level 3 configuration, where the disease spreads at time $t = 1$. Configurations **D** and **D′** depict the vaccination decisions (and subsequent epidemic spread) that occur at time 2, with **E** and **E′** respectively showing the final states.

Simplified uniform cost game. In the rest of the paper, we will frequently consider a simplified game instance where $\mathcal{T} = \{0, T\}$ and every vertex has the same infection cost L and vaccination costs C^0 and C^T. We will use the term $(G, T, \{C^0, C^T\}, L)$ to denote such instances. Further, when $C^0 = C^T = C$, we will use (G, T, C, L).

Cost model. As discussed above, the cost of a node depends on the final configuration— node v incurs cost L_v if it ever got infected, and it incurs cost C_v^t if it got vaccinated time t. The overall cost for a node is the expected cost over all the random sources (which is the only source of stochasticity in the model). Formally, we define the cost incurred by v under the strategy profile $Y(.)$ as:

$$\text{cost}(v, Y(.)) = C_v^0 Y(v, \cdot, 0) + \sum_{s \in V} \frac{1}{n} \Big(\sum_t C_v^t Y(v,s,t) + L_v \mathbf{1}_{s \to v | Y(.)} \Big),$$

where $\mathbf{1}_{s \to v | Y(.)} = 1$ if v gets infected due to s under strategy profile $Y(.)$. We define $\text{cost}(Y(.)) = \sum_v \text{cost}(v, Y(.))$.

B. Nash equilibria and social optimum

For a strategy profile $Y(\cdot)$, let $Y_{-v}(\cdot)$ be the strategy profile for all the remaining players. We say that a strategy $Y(\cdot)$ is a Nash equilibrium (NE) [16] if for each $v \in V$: $\text{cost}(v, Y'(.)) \geq \text{cost}(v, Y(.))$ where $Y'(.)$ is any strategy profile such that $Y'_{-v}(\cdot) = Y_{-v}(\cdot)$, i.e., $Y'(\cdot)$ has the same strategies as $Y(\cdot)$ for all other players $v' \neq v$. In other words, no player v can reduce its expected cost by unilaterally changing its strategy, given that the other players' strategies are fixed.

TABLE I: Notation summary

Network	
$G = (V, E)$	simple undirected graph with vertex set V and edge set E
n	number of vertices in G
$G[V']$	subgraph induced by $V' \subseteq V$
$N(v, G)$	neighborhood of v in G
$d_G(s, i)$	shortest path distance between nodes s and i in the graph G
$B(G, v, l)$	$\{u : d_G(u, v) \leq l\}$ i.e., l-ball
$P(G, v, l)$	$\{u : d_G(u, v) = l\}$ i.e., l-perimeter
Strategy	
$\mathcal{T}$	set of time instants at which the game is played
T	$\min\{t \in \mathcal{T} : t > 0\}$, smallest positive time in $\mathcal{T}$
$\mathbf{Y}$	set of all strategy profiles
$Y(.)$	a strategy profile
$Y(v, s, t)$	strategy for node v at time t, given source s; takes values from $\{0, 1\}$; $Y(v, \cdot, 0)$ for time 0
Y_t	strategies at time t, $\{Y(v, s, t) : v, s \in V\}$
$Y_{<t}$	strategies until time t, $\{Y(v, s, t') : v, s \in V, t' < t\}$
$Y_{-v}(\cdot)$	strategies of all players except v
$V_0(Y)$	$\{v \in V : Y(v, \cdot, 0) = 1\}$
Costs	
C_v^t	Cost for node v to vaccinate at time t
$\mathbf{C}$	$\{C_v^t \mid v \in V, t \in \mathcal{T}\}$
L_v	Cost of infection for node v
$\mathbf{L}$	$\{L_v \mid v \in V\}$
$\text{cost}(v, t)$	cost incurred by v at time t
Game	
$(G, \mathcal{T}, \mathbf{C}, \mathbf{L})$	TemporalVaccination game instance
$(G, T, \{C^0, C^T\}, L)$	two-stage uniform cost game
(G, T, C, L)	further simplified game instance with $C^0 = C^T$

We define the *social optimum* as a strategy $Y(\cdot)$ that has the minimum cost, over the space of all possible strategies—this is not necessarily (and is not usually) a pure NE. Therefore, the cost of a pure NE relative to the social cost is an important measure, and the maximum such ratio over all possible pure NE is known as the *price of anarchy* [12].

III. Characterization of NE in the TemporalVaccination game

We start with a characterization of pure NE, which will be repeatedly used in our subsequent discussions. Section III-A will cover the hardness of computing a NE. This will be followed by bounds on the Price of Anarchy. In Section III-C and III-D, we study the structure of the NE and social optimum of complete graphs and Erdős-Rényi graphs respectively.

Lemma 3.1: For a game instance $(G, \mathcal{T}, \mathbf{C}, \mathbf{L})$, let $T = \min\{t \in \mathcal{T} : t > 0\}$ be the smallest positive time in $\mathcal{T}$. Let $V_0(Y) = \{i \mid Y(i, \cdot, 0) = 1\}$, the set of nodes which vaccinated at time 0. For any $i \in V$, let $G_i = G[V - V_0(Y) \cup \{i\}]$. A strategy profile $Y(\cdot)$ is a pure NE iff

(1) $\forall i \in V$, $Y(i, \cdot, 0) + \sum_t Y(i, s, t) \leq 1$, for each $s \in V$.

(2) Every node $i \in V$ such that $Y(i, \cdot, 0) = 1$ satisfies $|B(G_i, i, T-1)| \frac{L_i}{n} + |P(G_i, i, T)| \frac{C_i^T}{n} > C_i^0$.

(3) Every node $i \in V$ such that $Y(i,\cdot,0) = 0$ satisfies $|B(G_i,i,T-1)|\frac{L_i}{n} + |P(G_i,i,T)|\frac{C_i^T}{n} \leq C_i^0$.

(4) Every node $i \in V$ such that $Y(i,s,T) = 1$ satisfies $|B(G_i,i,T-1)|\frac{L_i}{n} + |P(G_i,i,T-1)|\frac{C_i^T}{n} \leq C_i^0$ and $d(G_i,s,i) = T$.

Proof: Statement (1) follows from the fact that a node becomes immune the first time it gets vaccinated. Therefore, vaccinating again only increases its cost, and will not happen in a NE. Consider any insecure node i at time 0 (i.e., $Y(i,\cdot,0) = 0$). If $d_{G_i}(s,i) < T$, node i will be infected even before it gets a chance to vaccinate at T, and so it has $Y(i,s,T) = 0$. If $d_{G_i}(s,i) = T$, then node i is better off getting vaccinated, since $C_i^T < L_i$. Therefore, all nodes at distance T from s in $G[V - V_0(Y)]$ get vaccinated at time T.

From the above discussion, the expected cost of not vaccinating at time 0 is $\Pr(i$ will get infected whether it vaccinates at time T or not$)\times L_i + \Pr(i$ will vaccinate at time T and become secure$)\times C_i^T = |B(G_i,i,T-1)|\frac{L_i}{n} + |P(G_i,i,T)|\frac{C_i^T}{n}$. A node vaccinates at time 0 if the expected cost of not vaccinating is greater than C_i^0 (which implies statement (2)). If this expected cost is less than C_i^0, then the node does not get vaccinated at time 0, i.e., $Y(i,\cdot,0) = 0$ (which implies statement (3)). In this case, it vaccinates at time T (i.e., $Y(i,s,T) = 1$) iff it is at distance T from s, which implies statement (4). ■

Lemma 3.1 implies that once the source is revealed, all nodes make their vaccination decisions at the earliest possible time. Therefore, it suffices to study the TEMPORALVACCINATION game with $\mathcal{T} = \{0, T\}$.

Corollary 3.1: Let $Y(\cdot)$ be a pure NE for a TEMPORALVACCINATION instance $(G, \mathcal{T}, \mathbf{C}, \mathbf{L})$. Let $T = \min\{t \in \mathcal{T} : t > 0\}$ be the smallest positive time in $\mathcal{T}$. Then, we have $Y_i(t) = 0$ for all $t > T$.

It is easy to verify that Lemma 3.1 and Corollary 3.1 hold for mixed NE as well.

A. Existence of pure NE

Consider an instance $(G, \mathcal{T}, \mathbf{C}, \mathbf{L})$ where $\mathcal{T} = \{0, T\}$ and for each $v \in V$, $C_v^T = 0$. This game belongs to the family GNS(δ) studied in [13], where each node makes a decision to vaccinate or not at time 0, and the infection spreads from a random source node to a distance of upto δ. Here, $\delta = T - 1$. In [13], it was shown that for $\delta = 1$, pure NE always exists, while for $1 < \delta < \infty$, there are instances of GNS(δ) for which pure NE does not exist, and that determining if an instance has a pure NE is NP-complete. Therefore, we have the following result.

Lemma 3.2: The problem of determining if an instance $(G, \mathcal{T}, \mathbf{C}, \mathbf{L})$ has a pure NE is NP-complete. For the case $T = 2$, and $C_i^T = 0$ for all nodes $i \in V(G)$, a pure NE always exists.

Since pure NE don't always exist, we consider sufficient conditions in which they exist. Given $(G, \mathcal{T}, \mathbf{C}, \mathbf{L})$, we define an auxiliary graph $\widehat{G}$ as follows: Let $\mathcal{T} = \{0, T\}$. The vertex set of $\widehat{G}$, $\widehat{V}$ comprises of all vertices v which satisfy $\text{cost}(v, G) = |B(G,v,T-1)|\frac{L_i}{n} + |P(G,v,T)|\frac{C_v^T}{n} > C_v^0$. These are precisely the nodes which have the incentive to vaccinate at time 0. We draw an edge between two vertices u and v in $\widehat{G}$ if and only if $B(G,u,T) \cap B(G,v,T) \neq \varnothing$, i.e., u is adjacent to v if and only if removing one reduces the incentive to vaccinate for the other. Now we show the following:

Lemma 3.3: For an instance $(G, \mathcal{T}, \mathbf{C}, \mathbf{L})$ where, $\mathcal{T} = \{0, T\}$, a pure NE exists if $\widehat{G}$ is bipartite.

Proof: First, we note that if v is not adjacent to w in $\widehat{G}$, then, removing w does not change cost, i.e., $\text{cost}(v, G - \{w\}) = \text{cost}(v, G)$. Also, if the cost is above C_v^0 even after removing all the neighbors of v in $\widehat{G}$, then, it implies that v will vaccinate at time 0 in every pure NE. Such nodes can be ignored in our analysis. Therefore, from now on we will assume that G does not have such nodes and for every v, $\text{cost}(v, G - N(v, \widehat{G})) \leq C_v^0$. Suppose $\widehat{G}$ is a bipartite graph with bipartition $(\widehat{V}_1, \widehat{V}_2)$. We now show that the strategy Y satisfying $Y(v,\cdot,0) = 1$ if $v \in \widehat{V}_1$, and $Y(v,\cdot,0) = 0$ if $v \in V \setminus \widehat{V}_1$ corresponds to a pure NE. If a vertex $v \in \widehat{V}_1$ changes its strategy from secure to insecure, then, since $N(v, \widehat{G}) \cap \widehat{V}_1 = \varnothing$, it implies that $B(G(Y_{-v}), v, T) = B(G, v, T)$. Therefore, $\text{cost}(v, G(Y')) = \text{cost}(v, G) > C_v^0 \geq \text{cost}(v, G(Y))$. For a vertex $v \in \widehat{V}_2$, since $N(v, \widehat{G}) \subseteq \widehat{V}_1$, $G(Y)$ does not contain $N(v, \widehat{G})$. Hence by the earlier assumption, $\text{cost}(v, G(Y)) = \text{cost}(v, G - N(v, \widehat{G})) \leq C_v^0$. Therefore, v has no incentive to secure itself. The nodes in $V(G) \setminus V(\widehat{G})$ satisfy $\text{cost}(v, G(Y)) \leq \text{cost}(v, G) \leq C_v^0$. Therefore, they have no incentive to vaccinate either. Further, for every $v \in V$, we set $Y(v,s,T) = 1$ if and only if $v \in P(G[V \setminus \widehat{V}_1], s, T)$ and $Y(v,\cdot,0) = 0$. From Lemma 3.1, it follows that Y is a pure NE. ■

B. Bounds on the Price of Anarchy

Now, we will show that the price of anarchy for the TEMPORALVACCINATION game can be as high as $\Theta(n)$.

Lemma 3.4: Let C and L be any real numbers satisfying $0 < C < L$ and let $r = \lceil L/C \rceil$ and let $T > 1$. There exists a graph on $n > 2r^2T$ vertices such that the price of anarchy of the simplified TEMPORALVACCINATION instance (G, T, C, L) is at least $\frac{(n-2r^2T)^2}{n(2T+1)r}$. In particular, if C, L and T are constants, then, the price of anarchy is $\Theta(n)$.

Proof: The graph G is constructed as follows. The vertex set V is the disjoint union $\{\ell\} \cup \bigcup_{i=1}^{r} K_i \cup P_i$ where each $K_i = \{c_{i1}, \ldots, c_{i(n-2r^2T)/r}\}$ induces a cycle on the vertices c_{ij}, and each $P_i = \{p_{i1}, \ldots, p_{iT}\}$ induces the path $p_{i1}p_{i2}\cdots p_{iT}$. In addition, each p_{i1} is adjacent to all vertices in K_i, and each p_{iT} is adjacent to ℓ. Note that there is a deficit of $r^2T - rT - 1$ nodes. One way to rectify this is to add a path with these many nodes. This will not affect the proof in any way. Therefore, we will ignore this set of nodes for the rest of the proof. Let $\mathcal{P} = \bigcap_{i=1}^{r} P_i$ and $\mathcal{K} = \bigcap_{i=1}^{r} K_i$.

First, we will show that for every $v \in V(G)$, $\text{cost}(v, G) = |B(G,v,T-1)|\frac{L}{n} + |P(G,v,T)|\frac{C}{n} \leq C$. If $v \in K_i$, then,

$B(G,v,T-1) \subseteq K_i \cup P_i$ and $P(G,v,T)=\{p_{iT}\}$. Therefore, $\mathrm{cost}(v,G) = \left(\frac{n-2r^2T}{r}+T\right)\frac{L}{n}+\frac{C}{n} \leq \frac{(n-1)L}{nr}+\frac{C}{n} \leq C$. Suppose $v \in P_i$. If $v=p_{iT}$ then, $B(G,v,T-1) \subseteq \mathcal{P} \cup \{\ell\}$ and $P(G,v,T)=K_i \cup \{p_{j(T-1)} \mid j \neq i\}$. If $v=p_{ik}$, $k \neq T$, then, $B(G,v,T-1) \subseteq K_i \cup \mathcal{P} \cup \{\ell\}$ and $P(G,v,T) = \{p_{j(T-k)} \mid j \neq i\}$. In either case, $\mathrm{cost}(v,G) \leq \left(\frac{n-2r^2T}{r}+rT+r+1\right)\frac{L}{n} \leq C$. If $v=\ell$, then, $B(G,v,T-1) \subseteq \mathcal{P} \cup \{\ell\}$ and $P(G,v,T)=\{p_{iT} \mid 1 \leq i \leq r\}$. Therefore, $\mathrm{cost}(v,G) \leq (rT+1)\frac{L}{n} < C$. This implies that at time $t=0$, no vertex gets vaccinated and therefore, there always exists a pure NE for this game where the strategy of each vertex only depends on the source. Hence, the cost of the NE is only due to the second stage of vaccination.

Let s be the source. Now, we will compute a lower bound for the expected cost. The sets of vertices which are infected and vaccinated respectively are $B(G,s,T-1)$ and $P(G,s,T)$. It is sufficient to consider the case $s \in \mathcal{K}$. Suppose $s \in K_i$ for some i. Then, $B(G,s,T-1)=K_i \cup P_i \setminus \{p_{iT}\}$ and $P(G,s,T)=\{p_{iT}\}$. The cost of the NE is $\left(\frac{n-2r^2T}{r}+T-1\right)L+1 > \frac{n-2r^2T}{r}L$. Hence,

$$\mathrm{cost}(Y) > \Pr(s \in \mathcal{K})\frac{n-2r^2T}{r}L = \left(\frac{n-2r^2T}{n}\right)\frac{n-2r^2T}{r}L,$$

which is bounded by $\frac{\left(n-2r^2T\right)^2L}{nr}$. Now, we bound the cost of the social optimum from above. Suppose at time $t=0$, we vaccinate p_{i1} for all $i=1,\ldots,r$ and ℓ. Then, the residual graph is disconnected, and comprises of the following components: (1) cycles induced by K_i and (2) paths induced by $P_i \setminus \{p_{i1}\}$. It is easy to see that at time $T-1$, the number of nodes infected is at most $2(T-1)+1$, and the number of nodes to be vaccinated at $t=T$ to stop further spread is at most two. Hence, the cost of the social optimum is at most $\big(2(T-1)+1\big)L+2C \leq (2T+1)L$. Hence, proved. ■

C. Equilibria in complete graph

For the complete graph, we note that if $T>1$, then any node which did not vaccinate at time $t=0$ will get infected if the source is an infected node. The result below gives a lower bound on the number of nodes that vaccinate at $t=0$ in any NE when G is a complete graph. Its proof is omitted for brevity.

Lemma 3.5: For the instance $(G,\mathcal{T},\mathbf{C},\mathbf{L})$ where $G=(V,E)$ is a complete graph on n nodes and $T>1$ where, $T=\min\{t \in \mathcal{T} \mid t>0\}$, the number of nodes that vaccinate at $t=0$, n_0 is at least $\min_{v\in V}\left(1-\frac{C_v^0}{L_v}\right)n$.

For the two-stage uniform cost game, we can show the following stronger result.

Lemma 3.6: For the instance (G,T,C,L) where $G=(V,E)$ is a complete graph on n nodes, in every pure NE strategy $n_0=\left\lceil\left(1-\frac{C}{L}\right)n\right\rceil$ nodes vaccinate at $t=0$. Further, if $T>1$, the PoA is $\approx \left(1-\frac{C}{4L}\right)^{-1}$ for $C,L \ll n$.

Proof: By Lemma 3.5, every pure NE strategy should have at least $\left(1-\frac{C}{L}\right)n$ nodes vaccinating at $t=0$. If $n_0=\left\lceil\left(1-\frac{C}{L}\right)n\right\rceil$, then, for every remaining node, the expected cost is $\frac{n-n_0}{n}L \leq \frac{1}{n}\left(n-\left(1-\frac{C}{L}\right)n\right)L \leq C$. Also, making a secure node insecure will increase this cost for every remaining insecure node. Hence, this is a pure NE and we have proved the first part. Now we will consider the Price of Anarchy. When $T>1$, the expected cost of any strategy is $n_0C+\frac{n-n_0}{n}(n-n_0)L = n_0C+\frac{(n-n_0)^2}{n}L$, with minima at $n_0=\left(1-\frac{C}{2L}\right)n$. Therefore, the cost of social optimum is $\left(1-\frac{C}{2L}\right)nC+\left(\frac{C^2}{4L^2}\right)nL = nC\left(1-\frac{C}{4L}\right)$. By substituting $n_0=\left\lceil\left(1-\frac{C}{L}\right)n\right\rceil$ above, the cost of any pure NE is $\approx nC$. Therefore, PoA is $\left(1-\frac{C}{4L}\right)^{-1}$ ■

D. Equilibria in Erdős-Rényi graphs

Here, we study the structure of pure NE on instances of the $G(n,p)$ model. We consider two regimes: (R1) $\frac{1}{n} < p < \frac{\log n}{n}$ and $T \leq \frac{\log(\frac{nC}{4L\log n})}{1.5+\log np}$, and (R2) $p > \frac{c\log n}{n}$ and $T \leq \frac{\log(\frac{nC}{c'(L+npC)})}{\log np}$ where $c \leq 2$ is a constant and c' is only a function of c. We say that an event $A(n)$ occurs *almost surely* if $\Pr(A(n)) \to 1$ as $n \to \infty$.

Lemma 3.7: Consider the two-stage, uniform cost game (G,T,C,L) where, $G \in G(n,p)$. In both regimes (R1) and (R2), almost surely, pure NE exist, and any such strategy Y has $Y(v,\cdot,0)=0$ for all $v \in V(G)$. Further, the expected cost of a pure NE Y satisfies $\mathrm{cost}(Y) \leq 2T^3\log n(np)^T(L+C)$ for (R1) and $\mathrm{cost}(Y) \leq c'(np)^T(L+C)$ for (R2) almost surely.

Proof: Recall that the expected cost incurred by v for not vaccinating at time 0 is $\frac{1}{n}\big(|B(G,v,T-1)|L+|P(G,v,T)|C\big)$. First, we will prove the results concerning (R1). We will use the following result from [5]: For $np>1$, any $v \in V$, $|B(G,v,\ell)| \leq 2\ell^3(np)^\ell\log n$ and $|P(G,v,\ell) \leq 2\ell^2(np)^\ell\log n$ with probability at least $1-o(n^{-1.5})$. Henceforth, we will assume that for all v, this condition is satisfied. Note that, by union bound, this happens with probability $1-o(n^{-0.5})$. Applying this, the expected cost incurred by v for not vaccinating at time 0 is

$$\leq \frac{1}{n}\big(2(T-1)^3(np)^{T-1}(\log n)L+2T^2(np)^T(\log n)C\big)$$
$$\leq \frac{2T^3(np)^T\log n(L+C)}{n} < \frac{2(L+C)\log n}{n}e^{3\log T+T\log(np)}.$$

In the second inequality, we use the assumption $np>1$. Since, in (R1) $T \leq \frac{\log(\frac{nC}{2(L+C)\log n})}{1.5+\log np}$, and $3\log T+T(\log(np)) \leq T(1.5+\log(np))$, the cost is at most C. Therefore, from Lemma 3.1, it follows that $Y(v,\cdot,0)=0$ with probability $1-o(n^{-0.5})$ for all $v \in V(G)$. Again, with probability $1-o(n^{-0.5})$, the expected cost of any pure NE strategy is,

$$\begin{aligned}\mathrm{cost}(Y) &\leq \mathbb{E}_s[|B(G,s,T-1)|]L+\mathbb{E}_s[|P(G,s,T)|]C \qquad (1)\\ &\leq 2(T-1)^3(np)^{T-1}(\log n)L+2T^2(np)^T(\log n)C\\ &\leq 2T^3(np)^T(L+C)\log n\,.\end{aligned}$$

We proceed similarly for (R2). Here, we use the following result from [5]: For $p > \frac{c\log n}{n}$, where $c \leq 2$ is a constant, for any $v \in V(G)$, $|B(G,v,\ell)| \leq c'(np)^\ell$ and $|P(G,v,\ell) \leq c'(np)^\ell$ with probability at least $1-o(n^{-1.5})$, where c' is only a function of c. Again, this implies that (by union bound) with probability at least $1-o(n^{-0.5})$, these conditions hold for every $v \in V(G)$. The expected cost incurred by v

for not vaccinating at time 0 is at most $\frac{1}{n}\big(c'(np)^{T-1}L + c'(np)^T C\big) \le \frac{c'(np)^T(L+C)}{n}$. Since for (R2), $T \le \frac{\log(\frac{nC}{c'(L+C)})}{\log np}$, it follows that the cost is at most C. Hence, $Y(v,\cdot,0) = 0$ almost surely. As in the previous case, the cost of the pure NE Y is $\text{cost}(Y) \le c'(np)^{T-1}L + c'(np)^T C = c'(np)^T(L + C)$. ■

Social optimum: Let n_0 denote the number of nodes which are vaccinated at time 0. We show the following:

Lemma 3.8: In the two-stage, uniform cost game (G, T, C, L) where, $G \in G(n,p)$ and $p > \frac{c \log n}{n}$, the number of nodes vaccinated at time 0 in any optimal strategy is $n_0 \le n - \frac{1}{p}\big(\frac{nC}{2c'L(T+1)}\big)^{1/T}$, almost surely, where c' is a function of c.

Proof: Let G' denote the residual graph obtained after removing the vaccinated nodes at time 0. The cost of the optimum strategy Y^{OPT} is $\text{cost}(Y^{\text{OPT}}) = n_0 C + \mathbb{E}_s[|B(G', s, T-1)|]L + \mathbb{E}_s[|P(G', s, T)|]C$. Note that $G' \in G(n - n_0, p)$. Therefore, $\text{cost}(Y^{\text{OPT}})$

$$= n_0 C + \sum_{s \in V(G')} \frac{1}{n}\big(|B(G', s, T-1)|]L + |P(G', s, T)|C\big)$$

$$\le n_0 C + \frac{n - n_0}{n}\big(c'((n-n_0)p)^{T-1}L + c'((n-n_0)p)^T C\big)$$

$$< n_0 C + \frac{c'(n-n_0)^{T+1}p^T(L+C)}{n}.$$

Note that the last expression is a convex function in n_0 with minima at $n - \frac{1}{p}\big(\frac{nC}{c'(L+C)(T+1)}\big)^{1/T}$. ■

IV. Social optimum

We first show that computing the social optimum of the TemporalVaccination game is NP-complete, and then we develop an approximation algorithm using the approach of two-stage stochastic optimization.

Lemma 4.1: Computing the social optimum of a TemporalVaccination instance $(G, \mathcal{T}, \mathbf{C}, \mathbf{L})$ is NP-complete.

Proof: For a given strategy $Y(\cdot)$, we can estimate the expected cost in polynomial time, since there are only n random choices, corresponding to the random source. Therefore the decision version of this problem, i.e., deciding if the social optimum has cost at most B for a parameter B is in NP. When $C_j^t = 0$ for all $t \in \mathcal{T}$ and $j \in V$, this problem is equivalent to the GNS(d) game on the graph G, which is NP-hard [13]. ■

Approximating the social optimum. We now discuss an approximation algorithm for computing the social optimum. Our algorithm is based on the two-stage stochastic optimization approach. We first define some quantities which are needed for our algorithm. Let P_{ij}^T denote the set of all simple paths between nodes i and j having length at most T in G. We start with an integer programming formulation for the social optimum. We have the following variables

1) y_{0j} for each $j \in V$, which is an indicator for node j being vaccinated at time 0.
2) y_{sj} for each $j, s \in V$, which is an indicator for node j being vaccinated at time T, when the random source is s.
3) z_{ij} for each $i, j \in V$, which is an indicator for the event that there is no path $P \in P_{ij}^T$ consisting entirely of unvaccinated nodes.

Our integer programming $\mathcal{P}$ has the following structure:

$$\min f(\mathbf{y}, \mathbf{z}) = \sum_v C_v^0 y_{0v} + \frac{1}{n}\sum_{k,v} C_v^T y_{kv} + \frac{1}{n}\sum_{k,j} L_j(1 - z_{kj})$$

$$\text{s.t.} \sum_{v \in p} y_{0v} \ge z_{sj}, \forall p \in P_{sj}^T, \text{ with length} < T \quad (2)$$

$$\sum_{v \in p} y_{0v} + y_{sj} \ge z_{sj}, \forall p \in P_{sj}^T, \text{ with length } T \quad (3)$$

$$y_{0v}, y_{sv}, z_{ij} \in \{0,1\}, \quad \forall s, v, i, j \in V.$$

Lemma 4.2: Let $(\mathbf{y}, \mathbf{z})$ be the solution to the above integer program $\mathcal{P}$. Then, the strategy $Y(v, 0) = y_{0v}$ and $Y(v, s, T) = y_{sv}$ is an optimal solution to $(G, \mathcal{T}, \mathbf{C}, \mathbf{L})$.

Proof: Let $Y^{opt}(\cdot)$ be the socially optimal strategy to the given instance of the TemporalVaccination instance. If $Y^{opt}(j, \cdot, 0) = 0$, node j incurs a cost of L_j if the source is s and $d_{G[V - V_0(Y^{opt})]}(s, j) < T$, and it incurs a cost of C_j^T if $d_{G[V - V_0(Y^{opt})]}(s, j) = T$, since we assume that $C_j^t < L_j$ for all t, j. If $Y^{opt}(j, 0) = 1$, node j incurs cost C_j^0.

Let $Y(\cdot)$ be the strategy constructed from $(\mathbf{y}, \mathbf{z})$ in the lemma. We observe that $f(\mathbf{y}, \mathbf{z})$ equals $\text{cost}(Y)$, and $Y(\cdot)$ satisfies the same property as Y^{opt} above. For the nodes $j \in V_0(Y)$, we have $y_{0j} = 1$, which accounts for the cost $\sum_{j \in V_0(Y)} C_j^0$ of vaccination at time 0. Next, assume node j was not vaccinated at time t, i.e., $Y(j, \cdot, 0) = 0$. In the second stage, if j has distance less than T from the random source s in the residual graph, then there exists a path $p \in P_{sj}^T$ such that $y_{0v} = 0$ for all $v \in p$, including s and j. Then, constraint (2) in $\mathcal{P}$ causes $z_{sj} = 0$, so that node j incurs an infection cost of $L_j(1 - z_{sj}) = L_j$ in this case. On the other hand, if $d_{G[V - V_0(Y)]}(s, j) = T$, by constraint (3), we have $z_{sj} = 1$ if and only if $y_{sj} = 1$. This leads to a cost of $C_j^T y_{sj} = C_j^T$ for node j. Note that $y_{sj} = 0$ would have caused $z_{sj} = 0$, leading to an infection cost of $L_j(1 - z_{sj})$, which would be higher than C_j^T. Since s is the source with probability $\frac{1}{n}$, this cost component is scaled in $f(\mathbf{y}, \mathbf{z})$ by $\frac{1}{n}$. This implies $f(\mathbf{y}, \mathbf{z}) = \text{cost}(Y)$, and $Y(\cdot)$ satisfies the same property as Y^{opt}. ■

Lemma 4.3: For any T, the linear relaxation $\mathcal{P}_L$ can be solved in polynomial time.

Proof: When T is a constant, the linear program $\mathcal{P}_L$ is of polynomial size, since $|P_{sj}^T| = O(n^T)$. Therefore, in this case, we can directly solve the LP and find the fractional solution $(\mathbf{y}^1, \mathbf{z}^1)$. When T is not a constant, $\mathcal{P}_L$ has super-polynomially many constraints, and cannot be solved explicitly. We use the Ellipsoid method [21], which gives a polynomial time algorithm for finding a feasible solution x from a convex polytope $K \subseteq \mathbb{R}^N$. We refer to [21] for complete details, but note here that the key component needed for this method to work is to design a polynomial time "separation oracle", which, given a candidate solution

Algorithm 1: ApproxSocOpt

input : G, T, $\mathbf{C}$, $\mathbf{L}$
output: Strategy $Y(\cdot)$

Solve a linear relaxation $\mathcal{P}_L$ of the program $\mathcal{P}$, in which the constraints $y_{0v}, y_{sv} \in \{0,1\}$ and $z_{ij} \in \{0,1\}$ are replaced by

$$y_{0v}, y_{sv} \in [0,1], \qquad \forall s, v \in V$$
$$z_{ij} \in [0,1], \qquad \forall i, j \in V$$

If T is not a constant, $\mathcal{P}_L$ has exponentially many constraints, and we use the ellipsoid method to solve it, as discussed in Lemma 4.3.
Let $(\mathbf{y}^1, \mathbf{z}^1)$ denote the optimum fractional solution to $\mathcal{P}_L$.
We construct a new fractional solution $(\mathbf{y}^2, \mathbf{z}^2)$ in the following manner:
for *each* i, j **do**
  set $z^2_{ij} = 0$ if $z^1_{ij} \leq 1/2$, and $z^2_{ij} = 1$ otherwise
  for $j \in V$ **do**
    Set $y^2_{0j} = \min\{2y^1_{0j}, 1\}$ and $y^2_{sj} = \min\{2y^1_{sj}, 1\}$, for all $s \in V$.
for $j \in V$ **do**
  if $y^2_{0j} > 1/T$, we set $Y(\cdot, j, 0) = 1$, else we set $Y(\cdot, j, 0) = 0$.
  if $y^2_{sj} > 1/T$, we set $Y(s, j, T) = 1$, else we set $Y(s, j, T) = 0$.

$(\mathbf{y}, \mathbf{z})$, can decide if it is feasible, or finds a constraint that is infeasible.

We convert the $\mathcal{P}_L$ into a feasibility problem by "guessing" the cost of the objective, and adding a constraint that $f(\mathbf{y}, \mathbf{z}) \leq B$, where B is the estimate of the objective value. Such a separation oracle can be designed for the program $\mathcal{P}_L$ as follows:

1) For each pair s, j:
   a) Define weight $w_v = y_{0v}$ in the graph G
   b) Compute the shortest path distance $\text{dist}(s, j, T-1)$ from s to j restricted to paths with length at most $T-1$, based on the weights $\mathbf{w}$. If $\text{dist}(s, j, T-1) < z_{sj}$, we return $\sum_{v \in p} y_{0v} \geq z_{sj}$ as the violated constraint.
   c) Compute the shortest path distance $\text{dist}(s, j, T)$ from s to j restricted to paths with length at most T, based on the weights $\mathbf{w}$. If $\text{dist}(s, j, T) < z_{sj} - y_{sj}$, we return $\sum_{v \in p} y_{0v} + y_{sj} \geq z_{sj}$ as the violated constraint.
2) Finally, if $f(\mathbf{y}, \mathbf{z}) > B$, it is returned as a violated constraint.

These steps can be implemented in polynomial time. Therefore the separation oracle runs in polynomial time, so that the above algorithm returns a solution to program $\mathcal{P}_L$ in polynomial time, with cost at most the objective value, if it exists. We refer to [21] for complete details of the Ellipsoid method and its proof. ■

Lemma 4.4: Algorithm ApproxSocOpt gives a $2T$-approximation to the social optimum.

Proof: Let $Y^{opt}(\cdot)$ be the optimum solution. From Lemma 4.2, it follows that $f(\mathbf{y}^1, \mathbf{z}^1) \leq \text{cost}(Y^{opt})$.

We first argue that $(\mathbf{y}^2, \mathbf{z}^2)$ is feasible and $f(\mathbf{y}^2, \mathbf{z}^2) \leq 2f(\mathbf{y}^1, \mathbf{z}^1)$. By the construction, we have $1 - z^2_{ij} \leq 2(1 - z^1_{ij})$, whether $z^1_{ij} \leq 1/2$ or $z^1_{ij} > 1/2$. Constraints (2) and (3) corresponding to a pair s, j, continue to hold if $z^2_{sj} = 0$. The constraints corresponding to pairs s, j, for which $z^2_{sj} = 1$ also continue to hold, since $z^2_{sj} \leq 2z^1_{sj}$ in this case, and we have $y^2_{0j} = \min\{2y^1_{0j}, 1\}$ and $y^2_{sj} = \min\{2y^1_{sj}, 1\}$, for all $j, s \in V$.

For each pair s, j, if $z^2_{sj} = 1$, it follows that for every path $p \in P^T_{sj}$: (1) if $\text{len}(p) < T$, there must be some node $v \in p$ such that $y^2_{0v} \geq 1/T$, which implies $Y(\cdot, v, 0) = 1$; and (2) if $\text{len}(p) = T$, either there exists some node $v \in p$ such that $y^2_{0v} \geq 1/T$, or $y^2_{sj} \geq 1/T$. In the former case, we have $Y(\cdot, v, 0) = 1$, whereas in the latter case, we have $Y(s, j, T) = 1$.

Finally, by construction, we have $Y(\cdot, j, 0) \leq Ty^2_{0j}$ and $Y(s, j, T) \leq Ty^2_{sj}$. This implies that $\text{cost}(Y(\cdot)) \leq 2T \cdot \text{cost}(Y^{opt})$. ■

V. Experiments

We now study characteristics of Nash equilibria of the TemporalVaccination game in several social/communication networks and two random graph models, as summarized in Table II. In light of Lemma 3.2, we use a best response strategy to search for NE. We study the number and characteristics of nodes that get vaccinated at times 0 or T, and how this is affected by the relative vaccination costs at these times. Our main observations are summarized below.

TABLE II: Networks used in our experiments and their relevant properties: two real [15] and two synthetic graphs. We study the synthetic graphs with varying edge densities.

Network	Nodes (n)	Edges ($\lvert E \rvert$)
Ca-GRQC (co-authorship network)	4158	13422
AS20000102 (autonomous system network)	6474	12572
Barabasi-Albert	1000	varying
Erdős-Rényi	1000	varying

1. Number of vaccinated nodes at $t = 0$ and $t = T$
The number of nodes getting vaccinated initially, i.e., $|V_0(Y)|$, is very sensitive to $\frac{C}{L}$, and drops rapidly, as shown in Figure 2(c) and (d). Surprisingly, the number of nodes vaccinating at time T, i.e., $|V_T(Y)|$, however, is fairly stable across $\frac{C}{L}$ in both the networks. Further, there seems to be a cut-off point for C/L, where $|V_0(Y)|$ falls below $|V_T(Y)|$, which might be useful in policy design.

2. Performance of best response strategies
We find that the best response strategy generally converges to NE in linear number of rounds, as illustrated in Figure

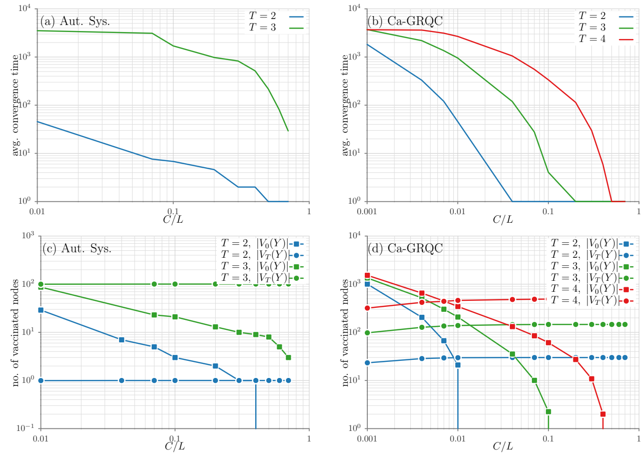

Fig. 2: Convergence time to NE using best response for (a) Autonomous System and (b) Co-authorship network. Variation of $|V_0(Y)|$ and $|V_T(Y)|$ with C/L, for different values of T in (c) Autonomous System and (d) Co-authorship network.

2. Moreover, the convergence time decreases rapidly with the C/L ratio, and increases with T. This suggests a high correlation between nodes the number of nodes vaccinating at time 0 and convergence time, which is an interesting topic for further investigation. The two real networks in Figure 2(a) and (b) exhibit very different behaviors, which might be the result of their structural differences. We also find that performance of best response is sensitive to the initial condition.

3. Effect of network density
Figure 3 (a) and (c) show the effect of network density, by varying the edge probability, p, for the Erdős-Rényi graph $G(n,p)$ and the number of edges per new node, m for the $BA(n,m)$ graph. We see that as the network becomes denser, more nodes vaccinate at time 0. Further, dense networks exhibit low diameter, and thus the game quickly approaches $GNS(\infty)$ of [13] for small increments in T.

4. Effect of heterogeneous costs
Figure 3 (b) considers the scenario when $C^T \neq C^0$. Note that for fixed C^0, as C^T increases, more nodes vaccinate at time 0.

5. Effect of T
From Figure 3 (d), as T, the minimum waiting time to next vaccination, increases, more nodes vaccinate at time 0. Further, this effect is more pronounced for denser networks. Qualitatively, this behavior is expected; as T increases, its T-ball becomes larger and the likelihood of a node getting infected before time T increases.

6. Correlation between degree and likelihood of vaccination at time 0
Figure 3 (e) shows the degree of nodes vaccinating at time 0 in NE plotted along with the degree of all nodes, for a Barabasi-Albert network with $N = 1000$, $m = 3$, $C/L = 0.1$, $T = 3$. We notice that almost all the nodes that get vaccinated at time 0 are among the top degree nodes. We observe a similar behavior in other networks, and these results are omitted because of space constraints.

VI. Related Work

There is a large literature on the use of non-cooperative game models for controlling the spread of epidemics and malware. We briefly summarize some of the main areas that are directly relevant to our paper.

A common approach in the mathematical epidemiology literature is based on differential equation models, e.g., [4], [10], [9], [6], [18], [19], [22]. These models are based on simplified assumptions about uniform mixing among the players, which allows for rigorous analysis. For instance, Bauch et al. [4] show that the NE can be completely characterized in terms of the reproductive number, which is the expected number of secondary infections caused by an infected individual. These models are deterministic and usually only consider vaccination strategies before the start of the epidemic. Reluga et al. [19] develop an approach that combines population games with Markov decision process, and consider decisions at different times.

Such differential equation models do not capture the complexity of interactions in real social contact networks. The work of Aspnes et al. [3] was among the first to study a network based formulation for vaccination games. They characterize NE in terms of the network structure and develop approximation algorithms for computing the social optimum. The utility function in their formulation requires the estimation of the probability that a node gets infected, which requires a lot of information. Kumar et al. [13] extend this formulation by restricting the amount of information needed by an individual. Our work builds on this formulation. Mean-field approximations have been used for detailed analysis in the SIS model, in which nodes switch from Susceptible to Infectious state, thereby capturing a more realistic epidemic model, e.g., [17], [22]. Saha et al. [20] consider a different formulation based on the spectral radius (the first eigenvalue of the network), in which the utility is based on whether or not the spectral radius is above a threshold or not— this is based on a characterization of the time to die out in the SIS model in terms of the spectral properties. However, all these approaches only consider vaccination decisions at the start of the epidemic, in a one-shot simultaneous game formulation. We also note that game-theoretical methods have been also used in other network security applications, e.g., [7], [8].

VII. Discussion and Conclusions

The TemporalVaccination game shows that vaccination decisions over time exhibit a very rich behavior. In both synthetic and real networks, and across a broad class of parameter regimes, we find that a significant fraction of nodes choose to get vaccinated later. The initially vaccinated fraction drops significantly as the ratio of vaccination and infection costs increases. Further, the timing of vaccination depends crucially on the network structure and information about the source and disease incidence. Therefore, the effect of delays in vaccination decisions has important implications for public health policy planning, and needs to be taken into careful consideration. In particular, vaccine availability needs to reflect the structure of equilibria. As a result, computing properties

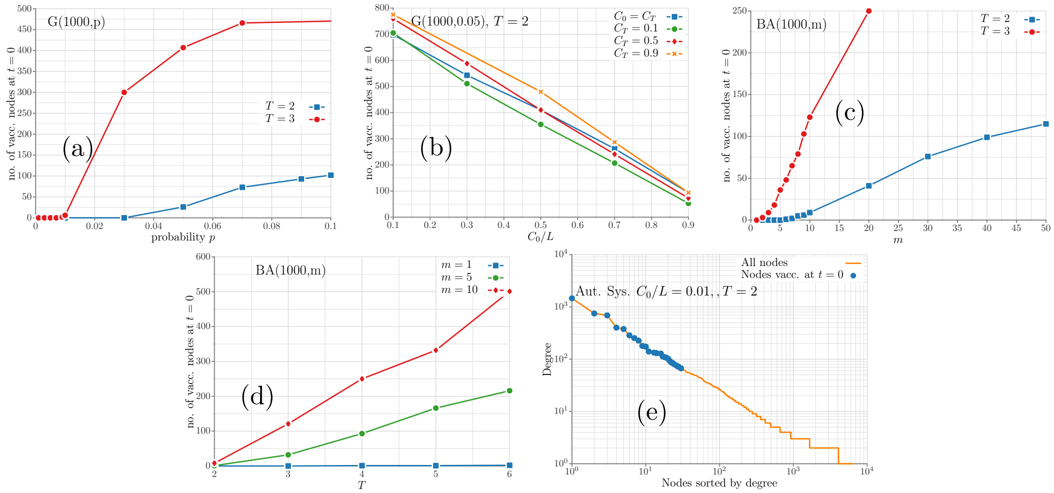

Fig. 3: The first four plots correspond to nodes vaccinating at $t = 0$ as a function of system parameters for the synthetic networks in Table II. The fifth plot shows the correlation between node degree and the chances of the node vaccinating at time 0.

of such games, including NE and social optimum, is an important issue. Admittedly, we have made a number of simplifying assumptions in our formulations. We only focus on a simple version of the **SI** model with infection probability 1, with perfect information about the source and other infections. Further, we assume vaccines have 100% efficacy and no delay. Relaxing all these assumptions are important directions for future work. We note that for many of our results, a vaccine delay of τ can be taken into account by considering an effective decision time of $T + \tau$, instead of T. Finally, we do not assume any resource constraints (e.g., a bound on the number of vaccines available at any time). In [1], we show that resource constraints change the structure of the game significantly.

Acknowledgments. This work has been partially supported by the following grants: DTRA Grant HDTRA1-11-1-0016, DTRA CNIMS Contract HDTRA1-11-D-0016-0010, NSF ICES CCF-1216000 and NSF NETSE Grant CNS-1011769.

References

[1] A. Adiga and A. Vullikanti. Temporal vaccination games under resource constraints. In *Proc. of AAAI*, 2016.

[2] J. Aspnes, K.L. Chang, and A. Yampolskiy. Inoculation strategies for victims of viruses and the sum-of-squares partition problem. *J. Comput. Syst. Sci.*, 2006.

[3] J. Aspnes, N. Rustagi, and J. Saia. Worm versus alert: Who wins in a battle for control of a large-scale network? *OPODIS*, 2007.

[4] C.T. Bauch and D.J.D. Earn. Vaccination and the theory of game. *PNAS*, 2004.

[5] Fan Chung and Linyuan Lu. The diameter of sparse random graphs. *Advances in Applied Mathematics*, 26(4):257–279, 2001.

[6] A. Galvani, T. Reluga, and G. Chapman. Long-standing influenza vaccination policy is in accord with individual self-interest but not with the utilitarian optimum. *PNAS*, 104(13):5692–5697, March 2007.

[7] J. Grossklags, N. Christin, and J. Chuang. Secure or insure? a game-theoretic analysis of information security games. In *World Wide Web Conference (WWW)*, 2008.

[8] M. Jain, V. Conitzer, and M. Tambe. Security scheduling for real-world networks. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2013.

[9] M. H. R. Khouzani, Eitan Altman, and Saswati Sarkar. Optimal quarantining of wireless malware through reception gain control. *IEEE Trans. Automat. Contr.*, 57(1):49–61, 2012.

[10] M. H. R. Khouzani, Saswati Sarkar, and Eitan Altman. Saddle-point strategies in malware attack. *IEEE Journal on Selected Areas in Communications*, 30(1):31–43, 2012.

[11] Donghyun Kim, Hao Guo, Yuchao Li, Wei Wang, Sung-Sik Kwon, and Alade O Tokuta. Social relation based long-term vaccine distribution planning to suppress pandemic. In *Proc. of CSoNet*, 2015.

[12] E. Koutsoupias and C. H. Papadimitriou. Worst-case equilibria. In *Proceedings of STACS*, 1999.

[13] V. S. Anil Kumar, R. Rajaraman, Z. Sun, and R. Sundaram. Existence theorems and approximation algorithms for generalized network security games. In *Proceedings of IEEE ICDCS*, 2010.

[14] M. Lelarge and J. Bolot. Economic incentives to increase security in the internet: The case for insurance. In *IEEE Infocom*, 2009.

[15] J. Leskovec. Snap: Stanford network analysis project, 2011.

[16] Kevin Leyton-Brown and Yoav Shoham. *Essentials of Game Theory: A Concise, Multidisciplinary Introduction.* Morgan and Claypool Publishers, 1st edition, 2008.

[17] J. Omic, A. Orda, and P. Van Mieghem. Protecting against network infections a game theoretic perspective. In *INFOCOM*, 2009.

[18] T. Reluga. Game theory of social distancing in response to an epidemic. *PLOS Computational Biology*, 6(5), 2010. e1000793.

[19] T. Reluga and A. Galvani. A general approach to population games with application to vaccination. *Mathematical Biosciences*, 2011.

[20] Sudip Saha, Abhijin Adiga, and Anil Kumar S. Vullikanti. Equilibria in epidemic containment games. In *The 28th AAAI Conference on Artificial Intelligence (AAAI)*, 2014.

[21] A. Schrijver. *Theory of Linear and Integer Programming.* Wiley, 1986.

[22] S. Trajanovski, Y. Hayel, E. Altman, H. Wang, and P. Van Mieghem. Decentralized protection strategies against sis epidemics in networks. *Control of Network Systems, IEEE Transactions on*, (99), 2015.