SIAM J. CONTROL OPTIM.
Vol. 60, No. 2, pp. S170–S195

INITIALIZATION AND CURING POLICIES FOR PÓLYA CONTAGION NETWORKS*

GREG HARRINGTON†, FADY ALAJAJI†, AND BAHMAN GHARESIFARD‡

Abstract. We investigate optimization policies for resource distribution in network epidemics using a model that derives from the classical Pólya process. This model, called the Pólya network contagion process, is based on a modified urn sampling scheme that accounts for both temporal and spatial contagion between neighboring nodes in a network. We study two infection mitigation problems—one which takes place upon initialization and one which occurs continually as the Pólya network process develops. We frame these problems as resource allocation problems with fixed budgets and analyze a suite of potential policies. Due to the complexity of these problems, we introduce effective proxy measures for the average infection rate in each case. We also prove that the two-sided infection-curing game on the so-called expected network exposure admits a Nash equilibrium. In both the curing and initialization scenarios, we introduce heuristic policies that primarily function on the basis of limiting the number of targeted nodes within a particular network setup. Simulations are run for mid-to-large–scale networks to compare performance of our heuristics to provably convergent gradient descent algorithms run on the simplified proxy measures.

Key words. network epidemics, interacting Pólya urn models, processes with reinforcement, temporal and spatial contagion, initialization and curing policies

AMS subject classifications. 92D30, 90B15, 60K35, 91A43

DOI. 10.1137/20M1358803

1. Introduction.

The study of epidemics on networks is an active research topic; see, e.g., [12, 33, 42] and the references therein. Real-life examples include the propagation of burst errors in a wireless communication channel [3], of a biological disease through a population [27], of malware in computer or smartphone systems [16], and the dissemination of rumors [40] and competing opinions [1] in social networks.

Controlling the spread of contagion in networks has been thoroughly investigated for various systems, including the well-known susceptible-infected-susceptible (SIS) and susceptible-infected-recovered (SIR) models [42] and extensions [34], models based on real epidemic data [43], and cascade models [12]. Other studies address link removal via immunization [29], optimization of curing resources [9], message-passing methods [26], optimal control under performance and resource usage trade-offs [25, 38], and competing or coevolving contagions within networks [24, 28, 30, 35, 44]. We herein study initialization and curing strategies and related game-theoretic problems for the (infinite-memory) Pólya network contagion model introduced in [21, 22].

The classical Pólya process [13, 36, 37] is a temporal contagion process that evolves via a sampling method from a single urn containing a finite number of red and black balls, representing units of "infection" and "healthiness," respectively. This model has been used in applications such as consensus dynamics [14] and generalized for various

*Received by the editors August 10, 2020; accepted for publication (in revised form) September 29, 2021; published electronically January 24, 2022. A preliminary version of section 4 was presented at the 16th Canadian Workshop on Information Theory as [19].

https://doi.org/10.1137/20M1358803

Funding: This work was funded in part by NSERC of Canada. The work of the third author was supported by NSERC of Canada and the Alexander von Humboldt Foundation.

†Department of Mathematics and Statistics, Queen's University, Kingston, ON K7L 3N6, Canada (12gth@queensu.ca, fa@queensu.ca).

‡Department of Electrical and Computer Engineering, University of California, Los Angeles, Los Angeles, CA 90095 USA (gharesifard@ucla.edu).

other purposes [10, 11]. The Pólya network contagion process [21, 22] generalizes this sampling scheme to a general network, where each node is equipped with its own urn, by allowing for spatial interactions between neighboring nodes. This is realized by forming a "super urn" for each node by combining its individual urn with the urns of its neighboring nodes [21, 22].

The resulting Pólya model for network contagion [21, 22] is similar to the SIS model [2, 12, 45]. The Pólya model can be operated in two modes—an infinite-memory mode where the reinforcing balls added at each time step remain permanently in the underlying urn of each node, and a finite-memory mode where the added balls at a given time step remain in the system only for a fixed number of future time steps (this mode results in a finite-order Markov network contagion process). While it was empirically observed in [22] that the Pólya network contagion model can mimic the behavior of the SIS process, particularly in the finite-memory (Markovian) mode, there are some notable differences between the two models. For example, the Pólya model benefits from exactly computable expressions for the joint and marginal probabilities of infection of its discrete-time contagion process, while the underlying (or exact) Markov process of the SIS model is typically represented and analyzed via a dynamical system using mean-field approximations; see, e.g., [2, 31, 41, 42]. Furthermore, the Pólya model captures the network contagion process at a microscopic level by associating an urn to each node, whose composition of red and black balls at each time step represents a granular profile for the degree of infection and healthiness of that node, which upon interaction with urns of nodes in proximity results in infected/healthy states with each super urn draw. This allows us to minutely and analytically capture the stochastic evolution of the network-wide contagion process. Finally, the underlying discrete-time Markov chain of the SIS model has an absorbing (all-healthy or disease-free) state [41], while the Pólya model does not in general[1] as its contagion process is symbiotically generated from its initial state (of red/black ball mixtures in the individual urns of the nodes) via the Pólya ball sampling mechanism.

In this paper, we examine resource allocation problems for the (infinite-memory) Pólya network contagion process. The first problem considers a one-time application of a control policy upon initialization of the network, while the second concerns the continual distribution of resources as the network contagion process develops. Both of these aim to manipulate the "average network-wide infection rate." Section 3 concerns the setup and analysis of the first such problem, while section 4 concerns the second. More specifically, the initialization problem concerns one party judiciously controlling the distribution of resources in an effort to minimize the average network-wide infection rate. This is a novel *preemptive* infection mitigation problem not typically considered in SIS models due to the Pólya model's inherent characteristics where contagion organically develops and propagates from initial network conditions. Resources in this case concern the distribution of red and black balls within a network prior to any draws taking place. We focus on the one-sided finite horizon case, wherein a player controls the allocation of black balls according to a fixed budget, with the goal of minimizing the average infection rate at some fixed point in the future. We provide a series of results showing that optimal policies for this problem satisfy two conditions: nested nodes will receive no resources, and symmetric red ball initializations will yield

[1]The Pólya model is more general than the exact SIS Markov process as it allows different asymptotic contagion behaviors. For example, in the finite-memory M mode with homogeneous parameters, the Pólya network process is an Mth order irreducible and aperiodic Markov chain with a unique stationary distribution that is not necessarily the all-healthy state with probability one.

symmetric black ball initializations (where the notion of "symmetry" is detailed using fundamental ideas from graph theory). Since determining the average infection rate at time $n \geq 1$ becomes increasingly complex as n grows, we use the average infection rate at $n = 1$ to simplify the problem, and we prove that under certain conditions this proxy measure admits an optimal policy. An algorithm for such an optimal policy using gradient descent is used for comparisons with our own heuristic policies. We detail three varieties of heuristic policies, each of which works on the basis of limiting our set of viable target nodes. The first set of policies targets the set of "inner nodes," the second works through a layered application of this inner node targeting technique, and the last works by iteratively targeting the most central nodes until full network coverage (either direct or indirect) is obtained. We provide two different algorithms for determining these target sets and variations based on previous heuristics (see [23]).

The second control problem, referred to as the "infection-curing problem" or the "Delta-curing problem," concerns *reactive* intervention policies deployed during the progress of contagion in the network. In this setup, which is an extension of the contagion-curing problem studied in [23], we define a two-player game on the average infection rate for a given network. Prior work on competitive dynamics over SIS and social networks include [35] and [30], respectively. To simplify the problem of finding optimal policies, we consider a game on the proxy measure of the expected network exposure, which we show to be convex in the curing parameters and concave in the infection parameters. We prove that under budget constraints, there exists a deterministic Nash equilibrium [32] for this game, which can be realized via gradient descent algorithms. We also develop heuristic policies for the Delta-curing problem using the same methods as in the initialization case. Finally, we provide simulations to demonstrate the performance of the proposed heuristics for the initialization and curing problems on both sparse and dense large-scale networks. These policies are compared to optimal policies for the one-step proxy measures.

The rest of the paper is organized as follows. In section 2, we describe in sufficient detail the Pólya network contagion process [22]. The initialization and infection-curing problems are theoretically investigated in sections 3 and 4, respectively. Heuristic optimization strategies for the above problems are developed in section 5, and simulation results are presented in section 6. Finally, conclusions are stated in section 7.

2. Pólya network contagion. We devote this section to some background on the classical Pólya urn process and the (infinite-memory) Pólya network contagion process under study. Throughout, results are derived using standard probability concepts, found in texts such as [4, 18]. Given the probability space $(\Omega, \mathcal{F}, P)$, consider the dimension-N stochastic process $\{\mathbf{Z}_n\}_{n=1}^{\infty}$, where $\mathbf{Z}_n = (Z_{1,n}, \ldots, Z_{N,n})$ is a random vector on Ω at time n. Here the process indices specify time indices, while the vector indices delineate spatial indices (for nodes in a network). For $i = 1, \ldots, N$, we write the n-tuple $(Z_{i,1}, \ldots, Z_{i,n})$ as Z_i^n. We represent by $\{\mathcal{F}_n\}_{n=1}^{\infty}$ the natural filtration on the process $\{\mathbf{Z}_n\}_{n=1}^{\infty}$ (i.e., $\mathcal{F}_n = \sigma(\{Z_j^n\}_{j=1}^N)$ is the σ-field generated by the stochastic process up to time n).

2.1. Classical Pólya process [13, 36, 37]. An urn initially contains $T = R+B$ balls of which $R \in \mathbb{Z}_{>0}$ are red and $B \in \mathbb{Z}_{>0}$ black. At time n, a ball is drawn from the urn and returned with $\Delta > 0$ balls of the same color. The random variable Z_n is used to indicate the color of the ball on the nth draw: if the draw is red, $Z_n = 1$; and if it is black, $Z_n = 0$. Setting $U_n := (R + \Delta \sum_{t=1}^{n} Z_t)/(T + n\Delta)$ as the proportion of red balls in the urn after the nth draw, the conditional probability of drawing a red

ball at time n, given the past draw history $(Z_1, \ldots, Z_{n-1})$, is given by

$$P(Z_n = 1 \mid Z_1, \ldots, Z_{n-1}) = \frac{R + \Delta \sum_{t=1}^{n-1} Z_t}{T + (n-1)\Delta} = U_{n-1}. \tag{2.1}$$

The process $\{Z_n\}_{n=1}^{\infty}$ is exchangeable (thus stationary) and nonergodic, with both U_n and $\frac{1}{n}\sum_{t=1}^{n} Z_t$ converging almost surely (a.s.) to a Beta distributed random variable with parameters $\frac{R}{\Delta}$ and $\frac{B}{\Delta}$ [15, 36]. We next describe the Pólya network contagion process [22, 23], which generates spatial contagion in addition to temporal contagion.

2.2. Pólya network contagion process [21, 22]. Let $\mathcal{G} = (V, \mathcal{E})$ be a graph, where $V = \{1, \ldots, N\}$ is the set of $N \in \mathbb{Z}_{>0}$ nodes and $\mathcal{E} \subseteq V \times V$ is the set of edges. Throughout, $\mathcal{G}$ is assumed to be undirected (i.e., $(i,j) \in \mathcal{E}$ if and only if $(j,i) \in \mathcal{E}$ for all $i, j \in V$) and connected (i.e., there exists a series of edges connecting any two nodes in $\mathcal{G}$). Let $\mathcal{N}_i = \{v \in V : (i,v) \in \mathcal{E}\}$ be the set of neighbors to node i, and let $\mathcal{N}_i' = \{i\} \cup \mathcal{N}_i$. Each node $i \in V$ is assigned an urn with $T_i = R_i + B_i$ balls of which $R_i \in \mathbb{Z}_{\geq 0}$ are red and $B_i \in \mathbb{Z}_{\geq 0}$ are black. We denote the black and red ball initializations by the vectors $\mathbf{B} := (B_1, \ldots, B_N)$ and $\mathbf{R} := (R_1, \ldots, R_N)$, respectively. We will refer to $\mathbf{B}$ as the curing initialization and $\mathbf{R}$ as the infection initialization.

As illustrated Figure 1, we form a super urn for node i by combining its individual urn with the urns of its neighboring nodes. We use $\bar{R}_i = \sum_{j \in \mathcal{N}_i'} R_j$ and $\bar{B}_i = \sum_{j \in \mathcal{N}_i'} B_j$ to denote the number of red and black balls, respectively, in node i's super urn, and we assume throughout that $\bar{R}_i + \bar{B}_i > 0$ for all $i \in V$ (to ensure that the super urn associated to each node is nonempty). The total number of balls in the ith super urn is given by $\bar{T}_i = \bar{R}_i + \bar{B}_i$. Similar to the classical Pólya process, a draw is conducted for each node at each time, and then a number of balls of the same color are added to that node's individual urn. In this case, however, the draw is conducted on the super urn of each node. As well, we may allow for the number of added balls to vary based on which node the draw was for, the color of the draw, and the time at which the draw occurred. Thus, for node i at time n, if a red ball is drawn, we add $\Delta_{r,i}(n)$ red balls to node i's individual urn, and if a black ball is drawn, we add $\Delta_{b,i}(n)$ black balls to node i's individual urn. We also assume that there exist $i \in V$ and n such that $\Delta_{b,i}(n) + \Delta_{r,i}(n) > 0$. We let $Z_{i,n}$ indicate the color of the nth draw for node i: if the draw is red, $Z_{i,n} = 1$; and if it is black, $Z_{i,n} = 0$.

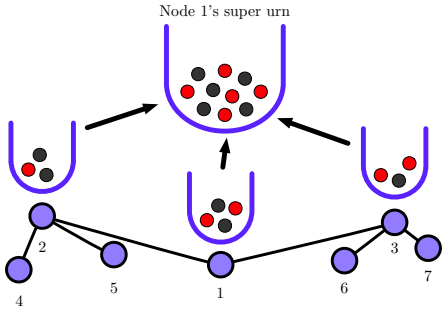

FIG. 1. *Illustration of a super urn in a network* [20].

The above parameters naturally translate into opinion dynamics applications involving two competing brands (red versus black). In this case, the initial red ball composition in node i represents the belief of the individual (residing in node i) about

the red brand at time $n = 0$. As each node interacts at time $n \geq 1$ with its neighbors (via the super urn), its belief is altered (via the reinforcement parameters $\Delta_{b,i}(n)$ and $\Delta_{r,i}(n)$) based on the opinions of the neighboring nodes. Furthermore, translating the above parameters to the framework of an epidemic in a population, black and red balls represent units of "healthiness" (such as white blood cells) and "infection" (such as bacteria or viruses), respectively [22]. In the super urn of a given node (such as a "person" or an "organism"), white blood cells (resp., bacteria) combine to improve (resp., impair) the overall health in the node's neighborhood. Drawing red at time n from the super urn means that the bacteria in the neighborhood were successful in replicating, making the individual more infected or less immune; otherwise, they were healthier since the white blood cells reproduced (as balls of the same color as the drawn balls are added). Thus, when $Z_{i,n} = 1$, we declare that node i is infected at time n, and if $Z_{i,n} = 0$, then it is healthy.

We refer to $\{\Delta_{b,i}(n)\}_{n=1}^{\infty}$ as the curing parameters (which in the biological epidemic setting can translate into medical therapeutics) and $\{\Delta_{r,i}(n)\}_{n=1}^{\infty}$ as the infection parameters for node i. From an urn-sampling point of view, these Δ's are nonnegative integers; however, we allow them to be nonnegative real numbers for mathematical convenience (e.g., when determining optimal control policies). The same assumption is used for the initialization parameters $\mathbf{B}$ and $\mathbf{R}$. The N-tuple of draw values at each time step yields the network contagion process $\{\mathbf{Z}_n\}_{n=1}^{\infty}$, where $\mathbf{Z}_n = (Z_{1,n}, \ldots, Z_{N,n})$. If desired, we can also separate out the individual draw process $\{Z_{i,n}\}_{n=1}^{\infty}$ for a given node i. We introduce some metrics to measure the infection spread in the network.

Similar to the classical Pólya process, we denote the proportion of red balls in node i's urn after the nth draw by $U_{i,n}$. For $T_i > 0$, letting

$$X_{i,n} = T_i + \sum_{t=1}^{n} \Delta_{r,i}(t) Z_{i,t} + \Delta_{b,i}(t)(1 - Z_{i,t}) \tag{2.2}$$

represent the total number of balls in node i's urn after the nth draw, we can write

$$U_{i,n} = \frac{R_i + \sum_{t=1}^{n} \Delta_{r,i}(t) Z_{i,t}}{X_{i,n}}. \tag{2.3}$$

Note that $U_{i,0} = \frac{R_i}{T_i}$ is the initial (at time $n = 0$) proportion of red balls in node i's individual urn. We denote the proportion of red balls in the super urn of node i at time n by $S_{i,n}$. This proportion is given by

$$S_{i,n} = \frac{\bar{R}_i + \sum_{j \in \mathcal{N}_i'} \sum_{t=1}^{n} \Delta_{r,j}(t) Z_{j,t}}{\sum_{j \in \mathcal{N}_i'} X_{j,n}} = \frac{\sum_{j \in \mathcal{N}_i'} U_{j,n} X_{j,n}}{\sum_{j \in \mathcal{N}_i'} X_{j,n}}, \tag{2.4}$$

where $S_{i,0} = \frac{\bar{R}_i}{\bar{T}_i}$ is the initial proportion of red balls in node i's super urn. We set

$$\tilde{S}_n := \frac{1}{N} \sum_{i=1}^{N} S_{i,n}$$

as the *network exposure*, which is a function of the underlying network contagion process $\{\mathbf{Z}_n\}_{n=1}^{\infty}$. For notational ease, we typically omit these arguments unless otherwise noted; e.g., for node i, we will write $S_{i,n}$ rather than $S_{i,n}(\{Z_j^n\}_{j=1}^{N})$.

The conditional probability of drawing a red ball for node i at time n, given all past draws $\{Z_j^{n-1}\}_{j=1}^N := \{(Z_{1,1},\dots,Z_{1,n-1}),\dots,(Z_{N,1},\dots,Z_{N,n-1})\}$, is given by

$$P\left(Z_{i,n}=1|\{Z_j^{n-1}\}_{j=1}^N\right)=\frac{\bar{R}_i+\sum_{j\in\mathcal{N}_i'}\sum_{t=1}^{n-1}\Delta_{r,j}(t)Z_{j,t}}{\sum_{j\in\mathcal{N}_i'}X_{j,n-1}}=S_{i,n-1}. \tag{2.5}$$

This is analogous to (2.1), except that it depends on $S_{i,n}$ rather than $U_{i,n}$. The n-fold joint probability of the network $\mathcal{G}$ is given for $a_i^n\in\{0,1\}^n$, $i\in V$, by

$$\begin{aligned}P_{\mathcal{G}}^{(n)}(a_1^n,\dots,a_N^n) &:= P(\{Z_i^n=a_i^n\}_{i=1}^N)=\prod_{t=1}^n P(\{Z_{i,t}=a_{i,t}\}_{i=1}^N|\{Z_i^{t-1}=a_i^{t-1}\}_{i=1}^N)\\ &=\prod_{t=1}^n\prod_{i=1}^N(S_{i,t-1})^{a_{i,t}}(1-S_{i,t-1})^{1-a_{i,t}}. \end{aligned}\tag{2.6}$$

Finally, we define the network *average infection rate at time n* as

$$\tilde{I}_n := \frac{1}{N}\sum_{i=1}^N P(Z_{i,n}=1), \tag{2.7}$$

which is a key metric to study the asymptotic behavior of the network contagion process. We can further break down (2.7) by noting that

$$\begin{aligned}P(Z_{i,n}=1) &= \sum_{\{a_j^{n-1}\}_{j=1}^N} P(Z_{i,n}=1|\{Z_j^{n-1}=a_j^{n-1}\}_{j=1}^N)P(\{Z_j^{n-1}=a_j^{n-1}\}_{j=1}^N)\\ &= \sum_{\{a_j^{n-1}\}_{j=1}^N} S_{i,n-1}(\{a_j^{n-1}\}_{j=1}^N)P_{\mathcal{G}}^{(n-1)}(\{a_j^{n-1}\}_{j=1}^N),\end{aligned}\tag{2.8}$$

where the summation is over all possible draw histories $a_j^{n-1}\in\{0,1\}^{n-1}$, $j\in V$.

Measures $\tilde{S}_n$ and $\tilde{I}_n$ are closely related and depend on the underlying network topology and initial ball distributions. As $\tilde{I}_n$ is difficult to analyze for all n, we will use $\tilde{I}_1$ and $\tilde{S}_n$ as proxy metrics to simplify the analysis or gain insight about $\tilde{I}_n$.

3. Initialization problems. When aiming to limit the spread of infection for the Pólya network contagion model, there are different ways to implement a curing policy. One method is to control the allocation of the curing parameters $\{\Delta_{b,i}(n)\}_{i=1}^N$ at each time n, as explored in section 4. We can, alternatively, control how the process is initialized as a preemptive measure, which we herein study.

The initialization parameters $\mathbf{B}\in\mathbb{R}_{\geq 0}^N$ and $\mathbf{R}\in\mathbb{R}_{\geq 0}^N$ can be tailored to alter the evolution of the Pólya network contagion process for a given network. The goal is to optimally allocate such resources. These optimizations may be performed for either or both sets of initialization parameters subject to a budget $\mathcal{B}\in\mathbb{R}_{\geq 0}$. We start with the one-sided initialization problem, wherein we aim to minimize $\tilde{I}_n$ over the black ball initialization parameters $\mathbf{B}$ subject to a budget $\mathcal{B}_b$. Throughout this section we assume the initial distribution of red balls $\mathbf{R}$ is known and fixed such that $\bar{R}_i>0$ for all $i\in V$.

PROBLEM 3.1 (one-sided finite horizon constrained budget initialization). *For a fixed time n, minimize the average infection rate $\tilde{I}_n$ subject to a budget $\mathcal{B}_b$ on the curing initialization $\mathbf{B}=(B_1,\dots,B_N)$; i.e., find*

$$\min_{\{B_i\}_{i=1}^N\in\mathbb{R}_{\geq 0}^N:\sum_{i=1}^N B_i\leq\mathcal{B}_b}\tilde{I}_n.$$

Although the above problem can be extended to the infinite horizon case, we limit ourselves to a finite horizon. Moreover, a *two-sided* initialization problem, wherein the objective is to minimize $\tilde{I}_n$ over both sets of initialization parameters $\mathbf{B}$ and $\mathbf{R}$, can be considered. However, we focus only on Problem 3.1, where for $n = 1$ we have

$$\tilde{I}_1 = \frac{1}{N}\sum_{i=1}^{N} P(Z_{i,1} = 1) = \frac{1}{N}\sum_{i=1}^{N} \frac{\bar{R}_i}{\bar{R}_i + \bar{B}_i}. \tag{3.1}$$

Note that $\tilde{I}_1$, which is well defined over $\{(\mathbf{B}, \mathbf{R}) \in \mathbb{R}^{2N}_{\geq 0} | \bar{B}_i + \bar{R}_i \neq 0 \, \forall i \in V\}$, is convex in $\mathbf{B}$ and concave in $\mathbf{R}$. We use (3.1) to provide some simplifications regarding solutions to Problem 3.1 for $n = 1$. We then extend these results for arbitrary n.

3.1. Outer node allocation. We first show that for Problem 3.1 with $n = 1$, nodes with a nested neighborhood (as specified in the lemma below) can be ignored. We refer to such nodes as "outer nodes" due to their topological location.

LEMMA 3.2. *For a general network $\mathcal{G} = (V, \mathcal{E})$ equipped with the Pólya network contagion model, if for any nodes $i, j \in V$ we have $\mathcal{N}_i' \subset \mathcal{N}_j'$, then any initial distribution that minimizes $\tilde{I}_1$ over the choice of $\mathbf{B}$ will satisfy $B_i = 0$. Likewise, any initial distribution that maximizes $\tilde{I}_1$ over the choice of $\mathbf{R}$ will satisfy $R_i = 0$.*

Proof. Consider a given initial distribution of black balls $\mathbf{B} = (B_1, \ldots, B_N)$ and red balls $\mathbf{R} = (R_1, \ldots, R_N)$. As written in (3.1), we have $\tilde{I}_1 = \frac{1}{N}\sum_{k=1}^{N} \frac{\bar{R}_k}{\bar{R}_k + \bar{B}_k}$. Without loss of generality, assume that $\mathcal{N}_1' \subset \mathcal{N}_2'$, and let the distribution $\mathbf{B}^* = (0, B_1 + B_2, B_3, \ldots, B_N)$ be an alternative curing initialization. For a given node $k \in V$, we have three cases to consider: (1) $k \in \mathcal{N}_1' \implies k \in \mathcal{N}_2'$; (2) $k \notin \mathcal{N}_1'$ and $k \in \mathcal{N}_2'$; (3) $k \notin \mathcal{N}_2' \implies k \notin \mathcal{N}_1'$. For cases (1) and (3), note that $\sum_{l \in \mathcal{N}_k'} B_l = \sum_{l \in \mathcal{N}_k'} B_l^*$. For case (2), we have $\sum_{l \in \mathcal{N}_k'} B_l \leq \sum_{l \in \mathcal{N}_k'} B_l^*$ with equality if and only if $B_1 = 0$. Denoting $\bar{B}_k^* = \sum_{l \in \mathcal{N}_k'} B_l^*$, we have

$$\frac{\bar{R}_k}{\bar{R}_k + \bar{B}_k} \geq \frac{\bar{R}_k}{\bar{R}_k + \bar{B}_k^*} \quad \forall k \in V.$$

Therefore, for $\tilde{I}_1$ as a function of the initial distribution of black balls we have that $\tilde{I}_1(\mathbf{B}^*) \leq \tilde{I}_1(\mathbf{B})$. The proof for the maximization case follows similarly. ∎

To extend Lemma 3.2 for general n, we derive the following intermediate results.

LEMMA 3.3. *Consider two sequences of draw values $\{a_j^n\}_{j=1}^N$ and $\{b_j^n\}_{j=1}^N$ that are equal, except for some $(k, s) \in V \times \{1, \ldots, n\}$, where $a_{k,s} = 1$ and $b_{k,s} = 0$. Then $P(Z_{i,n+1} = 1 | \{Z_j^n = a_j^n\}_{j=1}^N) \geq P(Z_{i,n+1} = 1 | \{Z_j^n = b_j^n\}_{j=1}^N)$ for all $i \in V$.*

Proof. Assume that $k \in \mathcal{N}_i'$ (note that if this is not the case, the two values we wish to compare are equal). Let

$$y = \sum_{j \in \mathcal{N}_i'} \sum_{t=1}^{n} a_{j,t} \Delta_{r,j}(t) \quad \text{and} \quad x = \sum_{j \in \mathcal{N}_i'} \sum_{t=1}^{n} (1 - a_{j,t}) \Delta_{b,j}(t).$$

Then letting $y^* = y - \Delta_{r,k}(s)$ and $x^* = x + \Delta_{b,k}(s)$, we obtain

$$\begin{aligned} P(Z_{i,n+1} = 1 | \{Z_j^n = a_j^n\}_{j=1}^N) &\geq \frac{\bar{R}_i + y}{\bar{R}_i + \bar{B}_i + y + x^*} \\ &\geq \frac{\bar{R}_i + y^*}{\bar{R}_i + \bar{B}_i + y^* + x^*} \\ &= P(Z_{i,n+1} = 1 | \{Z_j^n = b_j^n\}_{j=1}^N). \end{aligned}$$ ∎

Intuitively, this result states that the process is self-reinforcing; i.e., higher rates of infection increase the likelihood of the infection recurring. We will use this lemma in the proof of Lemma 3.4. Given the general form of $\tilde{I}_n$ given in (2.7), we next compare the performance of two curing initializations, $\mathbf{B}$ and $\mathbf{B}^*$, in limiting the spread of infection. In order to differentiate between the resultant probabilities, we will use * to denote any probability obtained using the curing initialization $\mathbf{B}^*$, e.g., $\tilde{I}_n^* = \tilde{I}_n(\mathbf{B}^*)$.

LEMMA 3.4. *Consider a general network $\mathcal{G} = (V, \mathcal{E})$ equipped with the Pólya network contagion model with two different curing initializations $\mathbf{B} = (B_1, \ldots, B_N)$ and $\mathbf{B}^* = (B_1^*, \ldots, B_N^*)$. If $S_{i,t}^*(a_1^t, \ldots, a_N^t) \leq S_{i,t}(a_1^t, \ldots, a_N^t)$ for all $(i,t) \in V \times \{1, \ldots, n\}$ and any sequence of draws $\boldsymbol{a}^n = (a_{1,s}, \ldots, a_{N,s})_{s=1}^n \in \{0,1\}^{N \times n}$, then $\tilde{I}_{n+1}^* \leq \tilde{I}_{n+1}$.*

Proof. Let $\{Y_{i,t} : t = 1, \ldots, n, i = 1 \ldots, N\}$ be independent random variables each uniformly distributed on the unit interval [0,1]. For all $i = 1, \ldots, N$, let $Z_{i,t}$ and $Z_{i,t}^*$ be inductively defined for $t = 1, \ldots, n$ by

$$Z_{i,t} = \begin{cases} 1 & \text{if } Y_{i,t} \leq S_{i,t-1}(\{Z_i^{t-1}\}_{i=1}^N), \\ 0 & \text{otherwise,} \end{cases} \tag{3.2}$$

and

$$Z_{i,t} = \begin{cases} 1 & \text{if } Y_{i,t} \leq S_{i,t-1}^*(\{Z_i^{*,t-1}\}_{i=1}^N), \\ 0 & \text{otherwise.} \end{cases} \tag{3.3}$$

Here we assume that $S_{i,0} \geq S_{i,0}^*$ for all $i = 1 \ldots, N$. Note that since

$$\begin{aligned} &P(Z_{i,t} = 1 | \{Z_i^{t-1}\}_{i=1}^N) = S_{i,t-1}(\{Z_i^{t-1}\}_{i=1}^N), \\ \text{and} \quad &P(Z_{i,t}^* = 1 | \{Z_i^{*,t-1}\}_{i=1}^N) = S_{i,t-1}^*(\{Z_i^{*,t-1}\}_{i=1}^N) \end{aligned}$$

for all $i = 1, \ldots, N$ and $t = 1, \ldots, n$, the processes have the required Pólya network process distribution with initializations $\mathbf{B}$ and $\mathbf{B}^*$, respectively.

We claim that if for some $t \in \{1, \ldots, n-1\}$ the sequence $\{Z_i^{t-1}\}_{i=1}^N$ dominates $\{Z_i^{*,t-1}\}_{i=1}^N$, in the sense that $Z_{i,s} \geq Z_{i,s}^*$ for all $i = 1, \ldots, N$ and $s = 1, \ldots, t-1$, then $\{Z_i^t\}_{i=1}^N$ dominates $\{Z_i^{*,t}\}_{i=1}^N$. To prove this claim, note that by repeated application of Lemma 3.3, we have that

$$S_{i,t-1}(\{Z_i^{t-1}\}_{i=1}^N) \geq S_{i,t-1}^*(\{Z_i^{*,t-1}\}_{i=1}^N)$$

for all $i = 1, \ldots, N$; specifically, since $Z_{i,s} \geq Z_{i,s}^*$ means that either $Z_{i,s} = Z_{i,s}^*$ or $Z_{i,s} = 1$ and $Z_{i,s}^* = 0$, Lemma 3.3 is applied $K = |\{(i,s) : Z_{i,s} \neq Z_{i,s}^*\}|$ times. This, (3.2), and (3.3) then give $Z_{i,t} \geq Z_{i,t}^*$, which in turn implies that $\{Z_i^t\}_{i=1}^N$ dominates $\{Z_i^{*,t}\}_{i=1}^N$, finishing the proof of the claim.

To proceed with the proof of the result, by using the conclusion of the claim for $t = 0$, which vacuously holds by the assumption that $S_{i,0} \geq S^*_{i,0}$ for all $i = 1, \ldots, N$, we obtain $Z_{i,1} \geq Z^*_{i,1}$, $i = 1, \ldots, N$. Applying now the result of the claim inductively, for $t = 2, \ldots, n$, we obtain that

$$Z_{i,t} \geq Z^*_{i,t}, \quad i = 1, \ldots, N, \quad t = 1, \ldots, n,$$

with probability one. Thus the event $\{Z^*_{i,t} = 1\}$ implies $\{Z_{i,t} = 1\}$; so $P(Z^*_{i,t} = 1) \leq P(Z_{i,t} = 1)$ for all $i = 1, \ldots, N$ and $t = 1, \ldots, n$. Hence $\tilde{I}^*_t \leq \tilde{I}^*_t$ for all $t = 1, \ldots, n$. ∎

We can now use this result to extend Lemma 3.2 to the case of a general time n.

THEOREM 3.5. *For a general network $\mathcal{G} = (V, \mathcal{E})$ equipped with the Pólya network contagion model, if for any nodes $i, j \in V$ we have $\mathcal{N}'_i \subset \mathcal{N}'_j$, then any curing initialization that minimizes $\tilde{I}_n$ will satisfy $B_i = 0$. Likewise, any infection initialization that maximizes $\tilde{I}_n$ will satisfy $R_i = 0$.*

Proof. Without loss of generality, assume $\mathcal{N}'_1 \subset \mathcal{N}'_2$. Consider the general form of $S_{i,n}(a^n_1, \ldots, a^n_N)$:

$$\frac{\bar{R}_i + \sum_{j \in \mathcal{N}'_i} \sum_{t=1}^{n} a_{j,t} \Delta_{r,j}(t)}{\bar{R}_i + \bar{B}_i + \sum_{j \in \mathcal{N}'_i} \sum_{t=1}^{n} (a_{j,t} \Delta_{r,j}(t) + (1 - a_{j,t}) \Delta_{b,j}(t))}. \tag{3.4}$$

Letting $B^*_1 = 0$, and $B^*_2 = B_1 + B_2$, as in the proof of Lemma 3.2, we have for any $i \in V$ that $\bar{B}^*_i \geq \bar{B}_i$. Substituting this into (3.4) yields that $S^*_{i,n} \leq S_{i,n}$ for any realization of the draw process, any $i \in V$, and any time t. Applying Lemma 3.4 yields the desired result. ∎

3.2. Symmetry. The next simplification we make regarding solutions to Problem 3.1 has to do with networks with "symmetry" properties, as described below. We start with some preliminaries, referring the reader to [7, 8, 17] for more details.

DEFINITION 3.6. *An automorphism of a graph $\mathcal{G} = (V, \mathcal{E})$ is a permutation τ of V, such that $(u, v) \in \mathcal{E}$ if and only if $(\tau(u), \tau(v)) \in \mathcal{E}$. The set of automorphisms of a given graph forms a group under the composition operation, called the automorphism group of $\mathcal{G}$ and denoted $Aut(\mathcal{G})$.*

DEFINITION 3.7. *For a graph $\mathcal{G} = (V, \mathcal{E})$, if $H \leq Aut(\mathcal{G})$ is a subgroup of automorphisms of $\mathcal{G}$, then $u, v \in V$ are similar under H if there exists an automorphism in H which maps u to v. Equivalence classes defined by similarity under H are called orbits of the graph $\mathcal{G}$ by H. The sets of orbits by H form a partition of the vertices of $\mathcal{G}$, called an orbit partition.*

If there exists an automorphism which maps between two different nodes, it means that those nodes are distinguishable only by their labels. For our purposes, "symmetry" between nodes is characterized by graph automorphisms, and an orbit partition allows us to divide a network based on these symmetries. Our goal is to prove that under certain conditions, sets of nodes that are similar under some $H \leq Aut(\mathcal{G})$ will be treated equivalently; i.e., they will all be given the same number of resources by an optimal policy. We start with a preliminary result.

LEMMA 3.8. *Consider a general network $\mathcal{G} = (V, \mathcal{E})$ with a given initialization $(\mathbf{B}, \mathbf{R}) = (B_1, \ldots, B_N, R_1, \ldots, R_N)$. Suppose that there exists a subset of nodes $V' \subset V$ such that $\bar{R}_i = \bar{R}_j$ for all $i, j \in V'$, and that we can find an alternate curing initialization $\mathbf{B}^* = (B^*_1, \ldots, B^*_N)$ such that $\bar{B}^*_i = \frac{1}{|V'|} \sum_{j \in V'} \bar{B}_j$ for all $i \in V'$. Then $\sum_{i \in V'} P^*(Z_{i,1} = 1) \leq \sum_{i \in V'} P(Z_{i,1} = 1)$.*

Proof. We prove this result by induction on $|V'|$. For the base case of $|V'| = 2$, consider two nodes $i, j \in V$ such that $\bar{R}_i = \bar{R}_j$. Let $\bar{B}_i^* = \bar{B}_j^* = \frac{1}{2}(\bar{B}_i + \bar{B}_j)$. We have

$$P(Z_{i,1} = 1) + P(Z_{j,1} = 1) = \frac{\bar{R}_i}{\bar{R}_i + \bar{B}_i} + \frac{\bar{R}_i}{\bar{R}_i + \bar{B}_j}, \tag{3.5}$$

$$P^*(Z_{i,1} = 1) + P^*(Z_{j,1} = 1) = \frac{2\bar{R}_i}{\bar{R}_i + \frac{1}{2}(\bar{B}_i + \bar{B}_j)}. \tag{3.6}$$

The difference between (3.5) and (3.6) yields

$$\begin{aligned}
&\frac{\bar{R}_i}{\bar{R}_i + \bar{B}_i} + \frac{\bar{R}_i}{\bar{R}_i + \bar{B}_j} - \frac{2\bar{R}_i}{\bar{R}_i + \frac{1}{2}(\bar{B}_i + \bar{B}_j)} \\
&\quad = \bar{R}_i \left[\frac{(2\bar{R}_i + \bar{B}_i + \bar{B}_j)(\bar{R}_i + \frac{1}{2}(\bar{B}_i + \bar{B}_j)) - 2(\bar{R}_i + \bar{B}_i)(\bar{R}_i + \bar{B}_j)}{(\bar{R}_i + \bar{B}_i)(\bar{R}_i + \bar{B}_j)(\bar{R}_i + \frac{1}{2}(\bar{B}_i + \bar{B}_j))} \right] \\
&\quad = \bar{R}_i \left[\frac{2\bar{R}_i^2 + 2\bar{R}_i(\bar{B}_i + \bar{B}_j) + \frac{1}{2}(\bar{B}_i + \bar{B}_j)^2 - 2\bar{R}_i^2 - 2\bar{R}_i(\bar{B}_i + \bar{B}_j) - 2\bar{B}_i\bar{B}_j}{(\bar{R}_i + \bar{B}_i)(\bar{R}_i + \bar{B}_j)(\bar{R}_i + \frac{1}{2}(\bar{B}_i + \bar{B}_j))} \right] \\
&\quad = \bar{R}_i \left[\frac{\frac{1}{2}(\bar{B}_i - \bar{B}_j)^2}{(\bar{R}_i + \bar{B}_i)(\bar{R}_i + \bar{B}_j)(\bar{R}_i + \frac{1}{2}(\bar{B}_i + \bar{B}_j))} \right] \geq 0.
\end{aligned}$$

Now, consider a subset of nodes $V' \subset V$ such that $\bar{R}_i = \bar{R}_j$ for all $i, j \in V'$, and suppose $|V'| = m+1$, where $m \geq 2$. Without loss of generality, let $V' = \{1, \ldots, m+1\}$. We assume the result holds for any subset of m of these nodes and prove it must hold for V'. We can write that

$$\sum_{i=1}^{m} P(Z_{i,1} = 1) = \sum_{i=1}^{m} \frac{\bar{R}_i}{\bar{R}_i + \bar{B}_i} \geq m \frac{\bar{R}_1}{\bar{R}_1 + \frac{1}{m}\sum_{i=1}^{m} \bar{B}_i}.$$

Let $\bar{B} = \frac{1}{m}\sum_{i=1}^{m} \bar{B}_i$, and let $\bar{B}_i^* = \frac{1}{m+1}(m\bar{B} + \bar{B}_{m+1})$ for all $i \in V'$. We note that $\bar{B}_i^* = \frac{1}{m+1}\sum_{j \in V'} \bar{B}_j$ for all $i \in V'$. We have

$$\sum_{i \in V'} P(Z_{i,1} = 1) \geq m \frac{\bar{R}_1}{\bar{R}_1 + \bar{B}} + \frac{\bar{R}_1}{\bar{R}_1 + \bar{B}_{m+1}}, \tag{3.7}$$

$$\sum_{i \in V'} P^*(Z_{i,1} = 1) = (m+1) \frac{\bar{R}_1}{\bar{R}_1 + \frac{1}{m+1}(m\bar{B} + \bar{B}_{m+1})}. \tag{3.8}$$

Taking the difference between the right-hand sides of (3.7) and (3.8), we have

$$\begin{aligned}
&\bar{R}_1 \left[\frac{m(\bar{R}_1 + \bar{B}_{m+1}) + (\bar{R}_1 + \bar{B})}{(\bar{R}_1 + \bar{B})(\bar{R}_1 + \bar{B}_{m+1})} - \frac{m+1}{\bar{R}_1 + \frac{1}{m+1}(m\bar{B} + \bar{B}_{m+1})} \right] \\
&= \bar{R}_1 \left[\frac{\frac{m}{m+1}(\bar{B}^2 + \bar{B}_{m+1}^2 - 2\bar{B}\bar{B}_{m+1})}{(\bar{R}_1 + \bar{B})(\bar{R}_1 + \bar{B}_{m+1})(\bar{R}_1 + \frac{1}{m+1}(m\bar{B} + \bar{B}_{m+1}))} \right] \\
&= \bar{R}_1 \left[\frac{\frac{m}{m+1}(\bar{B} - \bar{B}_{m+1})^2}{(\bar{R}_1 + \bar{B})(\bar{R}_1 + \bar{B}_{m+1})(\bar{R}_1 + \frac{1}{m+1}(m\bar{B} + \bar{B}_{m+1}))} \right] \geq 0.
\end{aligned}$$

Thus, we have shown that redistributing resources within the set V' leads to a reduction of the average infection rate within that set of nodes. ∎

We now use this result to prove that for networks with symmetry, solutions to Problem 3.1 will be symmetric in nature as well. Suppose for a graph $\mathcal{G} = (V, \mathcal{E})$ we have an automorphism τ of $\mathcal{G}$. Assuming $\mathcal{G}$ is finite, we have that $Aut(\mathcal{G})$ is a finite group, and thus we can use τ to generate a cyclic subgroup $\langle \tau \rangle$ of $Aut(\mathcal{G})$. The subgroup $\langle \tau \rangle$ is necessarily finite and of finite order. We also have that for any node i in V there exists some integer $k \in \mathbb{Z}_{>0}$ such that $i = \sigma^k(i)$. We will denote by k_i the smallest such integer for which this holds for a given $i \in V$.

THEOREM 3.9. *For a network $\mathcal{G} = (V, \mathcal{E})$ with a given initialization $(\mathbf{B}, \mathbf{R})$, consider an automorphism τ of $\mathcal{G}$. Let m be the order of the cyclic subgroup $\langle \tau \rangle$ of $Aut(\mathcal{G})$. Suppose $R_i = R_{\tau(i)}$ for all $i \in V$. Let $\mathbf{B}^*$ be an alternate curing initialization such that $B_i^* = \frac{1}{m} \sum_{j=1}^{m} B_{\tau^j(i)}$ for all $i \in V$. Then, $\tilde{I}_1^* \leq \tilde{I}_1$.*

Proof. Note that for a given $i \in V$, $\bar{R}_i = \sum_{j \in \mathcal{N}_i'} R_j$. We have that if $j \in \mathcal{N}_i'$, then $\tau^k(j) \in \mathcal{N}_{\tau^k(i)}'$ for any $k \in \mathbb{Z}_{>0}$. Since by assumption $R_j = R_{\tau(j)} = R_{\tau^k(j)}$ for all $j \in \mathcal{N}_i'$, we have that $\bar{R}_i = \bar{R}_{\tau^k(i)}$ for all $i \in V$ and $k \in \mathbb{Z}_{>0}$. From Lemma 3.8, we need only to prove that $\bar{B}_i^* = \frac{1}{k_i} \sum_{j=1}^{k_i} \bar{B}_{\tau^j(i)}$ for all $i \in V$. We start with

$$\bar{B}_i^* = \sum_{j \in \mathcal{N}_i'} B_j^* = \sum_{j \in \mathcal{N}_i'} \frac{1}{m} \sum_{l=1}^{m} B_{\tau^l(j)}.$$

We note that for a given positive integer p,

$$p \sum_{l=1}^{k_j} B_{\tau^l(j)} = \sum_{l=1}^{k_j} (B_{\tau^l(j)} + B_{\tau^{l+k_j}(j)} + \cdots + B_{\tau^{l+(p-1)k_j}(j)}) = \sum_{l=1}^{pk_j} B_{\tau^l(j)},$$

which follows from the definition of k_j. Since m is the order of $\langle \tau \rangle$, we have $\tau^m(j) = j$, and since k_j is the smallest such integer for which this holds, we must have $k_j | m$ (i.e., k_j divides m). Choosing $p = \frac{m}{k_j}$ and using the identity between the leftmost and rightmost terms in the above equation, we have that

$$\frac{1}{k_j} \sum_{l=1}^{k_j} B_{\tau^l(j)} = \frac{p}{pk_j} \sum_{l=1}^{k_j} B_{\tau^l(j)} = \frac{1}{pk_j} \sum_{l=1}^{pk_j} B_{\tau^l(j)} = \frac{1}{m} \sum_{l=1}^{m} B_{\tau^l(j)},$$

and thus

$$\begin{aligned} \bar{B}_i^* &= \sum_{j \in \mathcal{N}_i'} \frac{1}{m} \sum_{l=1}^{m} B_{\tau^l(j)} = \frac{1}{m} \sum_{l=1}^{m} \sum_{j \in \mathcal{N}_i'} B_{\tau^l(j)} = \frac{1}{m} \sum_{l=1}^{m} \sum_{j \in \mathcal{N}_{\tau^l(i)}'} B_j \\ &= \frac{1}{m} \sum_{l=1}^{m} \bar{B}_{\tau^l(i)} = \frac{1}{k_i} \sum_{l=1}^{k_i} \bar{B}_{\tau^l(i)}. \end{aligned}$$

The result then follows from Lemma 3.8. ∎

Another way to think about Theorem 3.9 is that the optimal allocation distributes resources equally within sets of the orbit partition by $\langle \tau \rangle$. We assume that infection resources are distributed evenly within this orbit partition, a consequence of assuming that $R_i = R_{\tau(i)}$ for all $i \in V$. We then prove that redistributing curing resources within each set of the partition, i.e., letting $B_i^* = \frac{1}{m} \sum_{j=1}^{m} B_{\tau^j(i)}$ for all $i \in V$, will result in a decrease in the overall average infection rate.

Algorithm 1. Constrained gradient descent on a simplex [5].

$f(x_1, \ldots, x_N) \leftarrow$ function to be minimized
Start at an arbitrary node:
$y_1 = (y_{1,1}, \ldots, y_{1,N}) = (\mathcal{B}_b, 0, \ldots, 0)$
for $k = 1 : stoptime$ **do**
  Find the direction of steepest descent:
  $i = \arg\min_{j \in V} \frac{\partial f}{\partial x_j}|_{y_k}$
  Move only in that direction:
  $\bar{y}_{k,i} = \mathcal{B}_b$, and $\bar{y}_{k,j} = 0$ for all $j \neq i$
  Select the step size using the limit minimization rule:
  $\alpha_k = \arg\min_{\alpha \in [0,1]} f(y_k + \alpha(\bar{y}_k - y_k))$
  Perform the gradient descent:
  $y_{k+1} = y_k + \alpha_k(\bar{y}_k - y_k)$
end for

We finish this part with a few remarks. First, due to the complicated nature of $\tilde{I}_n$, it is difficult to analytically prove this result for time $n > 1$. Also, while Theorem 3.9 can be applied to a given automorphism, one can push this idea further. Consider a network $\mathcal{G} = (V, \mathcal{E})$, and the orbit partition of $\mathcal{G}$ by $Aut(\mathcal{G})$, denoted $P = \{V_1, \ldots, V_q\}$, where $q \leq N$. For any given initialization, $(\mathbf{B}, \mathbf{R})$, for which $\bar{R}_i = \bar{R}_j$ if i and j are similar under $Aut(\mathcal{G})$, we postulate that we can reduce the average infection rate $\tilde{I}_n$ by redistributing curing resources evenly within each set of the partition. To elaborate, for each $i \in V$, if $i \in V_k$, we conjecture that taking $B_i^* = \frac{1}{|V_k|} \sum_{j \in V_k} B_j$ yields that $\tilde{I}_n(\mathbf{B}^*) \leq \tilde{I}_n(\mathbf{B})$ for any n. This conjecture is supported empirically, but for reasons of space, we omit including a sample simulation.

3.3. Gradient descent. As in [22], and further explored in section 4, one can implement a gradient descent algorithm [5] to find solutions to Problem 3.1. We focus on minimizing the proxy measure $\tilde{I}_1$ over the curing initialization $\mathbf{B}$ subject to a budget $\mathcal{B}_b$. It is easy to use gradient flow techniques to find an optimal policy, which will make full use of the budget $\mathcal{B}_b$, for this one-step problem.

PROPOSITION 3.10 (gradient descent conditions). *For a general network $\mathcal{G} = (V, \mathcal{E})$ equipped with the Pólya network contagion model under an infection initialization $\mathbf{R}$ satisfying $\bar{R}_i > 0$ for each $i \in V$, the average infection rate at time one, $\tilde{I}_1$, is convex with respect to the curing initialization parameters $\mathbf{B}$. Moreover, the feasible set $\mathcal{X} = \{\{B_i\}_{i=1}^N \in \mathbb{R}_{\geq 0}^N | \sum_{i=1}^N B_i = \mathcal{B}_b\}$ is convex and compact.*

Referring to (3.1), we note that $\tilde{I}_1$, as a function of the curing distribution, $\mathbf{B}$, is simply a summation of functions of the form $f(\mathbf{x}) = \frac{c}{c + a^T \mathbf{x}}$, where a^T is a row vector containing only ones or zeros, $\mathbf{x}$ is a column vector such that $x_i = B_i \ \forall i \in V$, and $c \in \mathbb{R}_{>0}$ is constant. Since any $f(\mathbf{x})$ of this form is convex in $\mathbf{x}$ for $\mathbf{x} \in \mathbb{R}_{\geq 0}^N$, we obtain that $\tilde{I}_1$ is convex in $\mathbf{B}$ (which follows because the feasible set $\mathcal{X}$ is convex and compact, the proof of which is straightforward). Of note, however, is the fact that we require $\bar{R}_i > 0$ for each $i \in V$ in order to ensure that the feasible set $\mathcal{X}$ remains valid. If this is satisfied, we can then use a constrained gradient descent method [5, Chapter 2], as described in Algorithm 1. This algorithm will converge to an optimal solution for the one-step optimization problem, but not necessarily for the n-step case.

4. Curing-infection problems. A reactive technique to mitigate network infection is to control the allocation of the curing parameters, $\{\Delta_{b,i}(n)\}_{i=1}^N$, subject to a

fixed budget, $\mathcal{B}$. As in the initialization setup, this can be done for a finite or infinite horizon. Previously, [23] examined the problem of minimizing the limiting average infection rate and introduced a number of heuristic optimization policies. We begin by refocusing the problem in a game-theoretic setup.

We consider a two-player game where player one minimizes the average infection rate $\tilde{I}_n$, while player two maximizes it. Player one controls the distribution of curing parameters, $\{\Delta_{b,i}(n)\}_{i=1}^N$, subject to budget $\mathcal{B}_b$, while player two controls the infection parameters, $\{\Delta_{r,i}(n)\}_{i=1}^N$, under budget $\mathcal{B}_r$. The budget in each case is fixed for all time steps, and the allocation of resources for a given time n is determined prior to any draws being made at that time. Thus, if either player allocates resources to a node for which the opposite color ball is drawn, those resources will go to waste.

Formally, player one's objective is to find

$$\min_{\{\{\Delta_{b,i}(k)\}_{i=1}^N:\sum_{i=1}^N \Delta_{b,i}(k)=\mathcal{B}_b\},k=1,\ldots,n} \tilde{I}_n, \tag{4.1}$$

assuming the minimum exists, while player two's objective is to find

$$\max_{\{\{\Delta_{r,i}(k)\}_{i=1}^N:\sum_{i=1}^N \Delta_{r,i}(k)=\mathcal{B}_r\},k=1,\ldots,n} \tilde{I}_n, \tag{4.2}$$

assuming the maximum exists. We refer to (4.1)–(4.2) as the Delta-curing problem.

Since finding an optimal control policy for either player is not tractable for a general network, we provide a simplified one-step optimization problem by adopting the expected network exposure, $E[\tilde{S}_n|\mathcal{F}_{n-1}]$, as a proxy metric to optimize instead of $\tilde{I}_n$. We thus consider a two-player zero-sum game, where player one minimizes $E[\tilde{S}_n|\mathcal{F}_{n-1}]$ over $\{\Delta_{b,i}(n)\}_{i=1}^N$ and player two maximizes the same value over $\{\Delta_{r,i}(n)\}_{i=1}^N$. We begin with an important result about the expected network exposure.

PROPOSITION 4.1 (convexity-concavity of network exposure). *For a general network $\mathcal{G} = (V, \mathcal{E})$ equipped with the Pólya network contagion model under arbitrary initial conditions, the expected network exposure $E[\tilde{S}_n|\mathcal{F}_{n-1}]$ is convex with respect to $\{\Delta_{b,i}(n)\}_{i=1}^N$ and concave with respect to $\{\Delta_{r,i}(n)\}_{i=1}^N$ for all n.*

Proof. Using (2.2) and (2.4), we consider $E[\tilde{S}_n|\mathcal{F}_{n-1}]$ as a function of the vectors

$$\mathbf{x} = (x_1, \ldots, x_N)^T = (\Delta_{b,1}(n), \ldots, \Delta_{b,N}(n))^T,$$
$$\mathbf{y} = (y_1, \ldots, y_N)^T = (\Delta_{r,1}(n), \ldots, \Delta_{r,N}(n))^T$$

by reformulating (2.4) as follows:

$$S_{i,n} = f_{i,n}(\mathbf{x}, \mathbf{y}, \mathbf{Z}_n) = \frac{c_i + \delta_i(\mathbf{y}, \mathbf{Z}_n)}{c_i + d_i + \sigma_i(\mathbf{x}, \mathbf{Z}_n) + \delta_i(\mathbf{y}, \mathbf{Z}_n)},$$

where

$$c_i = \bar{R}_i + \sum_{t=1}^{n-1} \sum_{j \in \mathcal{N}_i'} \Delta_{r,j}(t) Z_{j,t}, \quad d_i = \bar{B}_i + \sum_{t=1}^{n-1} \sum_{j \in \mathcal{N}_i'} \Delta_{b,j}(t)(1 - Z_{j,t}),$$
$$\delta_i(\mathbf{y}, \mathbf{Z}_n) = \sum_{j \in \mathcal{N}_i'} y_j Z_{j,n}, \quad \sigma_i(\mathbf{x}, \mathbf{Z}_n) = \sum_{j \in \mathcal{N}_i'} x_j (1 - Z_{j,n}).$$

Alternatively, we let A represent the adjacency matrix, including self-loops, for our given network $\mathcal{G}$, where $A_{ij} = 1$ if and only if $i \in \mathcal{N}_j'$, and $A_{ij} = 0$ otherwise. We

then construct an $N \times N$ square matrix, D, where $D_{ij} = A_{ij}(1 - Z_{j,n})$. Letting D_i represent the ith row of the matrix D, we have that $\sigma_i(x, \mathbf{Z}_n) = D_i\mathbf{x}$. Likewise, we can construct an $N \times N$ square matrix, C, where $C_{ij} = A_{ij}Z_{j,n}$. Letting C_i represent the ith row of the matrix C, we have that $\delta_i(y, \mathbf{Z}_n) = C_i\mathbf{y}$, and hence

$$f_{i,n}(\mathbf{x}, \mathbf{y}, \mathbf{Z}_n) = \frac{c_i + C_i\mathbf{y}}{c_i + d_i + C_i\mathbf{y} + D_i\mathbf{x}},$$

where we have dropped the dependencies of the matrices C and D on $\mathbf{Z}_n$ for simplicity. Taking the expectation of $S_{i,n}$ given the history of the contagion process up to time $n-1$, we have

$$\begin{aligned} E[S_{i,n}|\mathcal{F}_{n-1}] &= E[f_{i,n}(\mathbf{x}, \mathbf{y}, \mathbf{Z}_n)|\mathcal{F}_{n-1}] \\ &= \sum_{\mathbf{z}_n \in \{0,1\}^N} f_{i,n}(\mathbf{x}, \mathbf{y}, \mathbf{z}_n) P(\mathbf{Z}_n = \mathbf{z}_n | \{Z_j^{n-1}\}_{j=1}^N = \{z_j^{n-1}\}_{j=1}^N). \end{aligned} \tag{4.3}$$

$P(\mathbf{Z}_n = \mathbf{z}_n|\{Z_j^{n-1}\}_{j=1}^N = \{z_j^{n-1}\}_{j=1}^N)$ is independent of our choice of $\mathbf{x}$, and for any fixed realization $\mathbf{z}_n = (z_{1,n}, \ldots, z_{N,n}) \in \{0,1\}^N$, $f_{i,n}(\mathbf{x}, \mathbf{y}, \mathbf{z}_n)$ is convex in $\mathbf{x}$ over $\mathbb{R}^N_{\geq 0}$ (the proof is given in [23]). Hence, $E[S_{i,n}|\mathcal{F}_{n-1}]$ is convex in $\mathbf{x}$ over $\mathbb{R}^N_{\geq 0}$, and thus so too is $E[\tilde{S}_n|\mathcal{F}_{n-1}]$. Concavity in $\mathbf{y}$ follows from a symmetry argument since

$$1 - f_{i,n}(\mathbf{x}, \mathbf{y}, \mathbf{Z}_n) = \frac{d_i + D_i\mathbf{x}}{c_i + d_i + C_i\mathbf{y} + D_i\mathbf{x}}$$

is convex in $\mathbf{y}$, thus showing $f_{i,n}(\mathbf{x}, \mathbf{y}, \mathbf{Z}_n)$ is concave in $\mathbf{y}$. The rest of the proof follows as in the convexity argument for $\mathbf{x}$. ∎

The natural symmetry of this model makes this result intuitive and facilitates the game-theoretic setup. We have proven that $E[\tilde{S}_n|\mathcal{F}_{n-1}]$ is convex in $\mathbf{x}$ and concave in $\mathbf{y}$ over $\mathbb{R}^N_{\geq 0} \times \mathbb{R}^N_{\geq 0}$. In order to get a better understanding of this function, we consider the partial derivatives with respect to $\mathbf{x}$ and $\mathbf{y}$:

$$\begin{aligned} \nabla_{\mathbf{x}} f_{i,n}(\mathbf{x}, \mathbf{y}, \mathbf{z}_n) &= \frac{-\nabla_{\mathbf{x}} D_i\mathbf{x}(c_i + C_i\mathbf{y})}{(c_i + d_i + C_i\mathbf{y} + D_i\mathbf{x})^2}, \\ \nabla_{\mathbf{y}} f_{i,n}(\mathbf{x}, \mathbf{y}, \mathbf{z}_n) &= \frac{\nabla_{\mathbf{y}} C_i\mathbf{y}(d_i + D_i\mathbf{x})}{(c_i + d_i + C_i\mathbf{y} + D_i\mathbf{x})^2}, \\ \frac{\partial}{\partial x_j} D_i\mathbf{x} = A_{ij}(1 - z_{j,n}), &\text{ and } \frac{\partial}{\partial y_j} C_i\mathbf{y} = A_{ij}z_{j,n}. \end{aligned}$$

Since A_{ij} and $z_{j,n}$ can only take values in $\{0,1\}$, we have that $\frac{\partial}{\partial x_j} f_{i,n}(\mathbf{x}, \mathbf{y}, \mathbf{z}_n) \leq 0$ and $\frac{\partial}{\partial y_j} f_{i,n}(\mathbf{x}, \mathbf{y}, \mathbf{z}_n) \geq 0$. Using this fact in conjunction with (4.3), we obtain that over the space $\mathcal{X} \times \mathcal{Y} = \mathbb{R}^N_{\geq 0} \times \mathbb{R}^N_{\geq 0}$, $E[\tilde{S}_n|\mathcal{F}_{n-1}]$ has no saddle point. Given our fixed allocation budgets $\mathcal{B}_b$ and $\mathcal{B}_r$, we restrict ourselves to considering sets of the form $\mathcal{X} \times \mathcal{Y} = \{\{\Delta_{b,i}(n)\}_{i=1}^N \in \mathbb{R}^N_{\geq 0} | \sum_{i=1}^N \Delta_{b,i}(n) \leq \mathcal{B}_b\} \times \{\{\Delta_{r,i}(n)\}_{i=1}^N \in \mathbb{R}^N_{\geq 0} | \sum_{i=1}^N \Delta_{r,i}(n) \leq \mathcal{B}_r\}$. Returning to our game over the expected network exposure, we remark that for any given n, the sets $\mathcal{X} = \{\{\Delta_{b,i}(n)\}_{i=1}^N \in \mathbb{R}^N_{\geq 0} | \sum_{i=1}^N \Delta_{b,i}(n) \leq \mathcal{B}_b\}$ and $\mathcal{Y} = \{\{\Delta_{r,i}(n)\}_{i=1}^N \in \mathbb{R}^N_{\geq 0} | \sum_{i=1}^N \Delta_{r,i}(n) \leq \mathcal{B}_r\}$ are convex and compact. This gives rise to the following result.

THEOREM 4.2 (Nash equilibrium for network exposure). *For a general network $\mathcal{G} = (V, \mathcal{E})$ equipped with the Pólya network contagion model under arbitrary initial*

conditions and a given time n, consider a two-player zero-sum game where player one tries to minimize $E[\tilde{S}_n|\mathcal{F}_{n-1}]$ over $\{\Delta_{b,i}(n)\}_{i=1}^N$ and player two tries to maximize $E[\tilde{S}_n|\mathcal{F}_{n-1}]$ over $\{\Delta_{r,i}(n)\}_{i=1}^N$. Then, if we take our set of allowable policies to be of the form $\mathcal{X} \times \mathcal{Y} = \{\{\Delta_{b,i}(n)\}_{i=1}^N \in \mathbb{R}_{\geq 0}^N | \sum_{i=1}^N \Delta_{b,i}(n) \leq \mathcal{B}_b\} \times \{\{\Delta_{r,i}(n)\}_{i=1}^N \in \mathbb{R}_{\geq 0}^N | \sum_{i=1}^N \Delta_{r,i}(n) \leq \mathcal{B}_r\}$, the resulting game admits a Nash equilibrium. Moreover, the equilibrium policy will satisfy $\sum_{i=1}^N \Delta_{b,i}(n) = \mathcal{B}_b$ and $\sum_{i=1}^N \Delta_{r,i}(n) = \mathcal{B}_r$.

Proof. Since the function is convex-concave and over a compact set, the existence follows from the classical minimax theorem; see [6]. By the definition of $E[\tilde{S}_n|\mathcal{F}_{n-1}]$, and since the function has no saddle point in the interior of its domain, the optimal policy will utilize the full budget. ∎

The equilibrium policy from Theorem 4.2 can be determined numerically using gradient descent algorithms [5], as we are optimizing a convex-concave function over a simplex, but for large networks such algorithms can be computationally expensive, especially when considering the complexity of $E[\tilde{S}_n|\mathcal{F}_{n-1}]$, as seen in (4.3). We next investigate heuristic strategies that avoid this problem.

5. Heuristic strategies. Thus far in our analysis we have considered simplifying assumptions that can be made in order to minimize the average infection rate, $\tilde{I}_n$, for both the initialization and Delta-curing problems (studied in section 3 and section 4, respectively). The complex mathematical nature of these problems makes analytical determination of any optimal solutions difficult and makes computational determination of such solutions impractical, even when considering simplified versions of these problems. As such, we look to our analysis, as well as previous analyses in [20], [23], to give insight into basic heuristics that can avoid this complexity, while still providing effective solutions. Naturally, the effectiveness of such heuristics is measured via the original metric intended for these problems, which is the average network infection rate.

When developing heuristics in the initialization case, we make heavy use of Theorem 3.5, which proves that any optimal solution to Problem 3.1 would allocate no resources to outer nodes (i.e., nodes with a nested neighborhood). For the general Delta-curing problem, we have no such result, as it is possible to manufacture situations where the equilibrium policy as in Theorem 4.2 will assign resources to outer nodes. Such cases arise under specific circumstances, and it is uncommon for these nodes to receive significant resources. Thus, we introduce a similar set of resource allocation policies for both the initialization and Delta-curing problems of sections 3 and 4, respectively. While any of these resource allocation policies can be similarly applied to a player controlling the distribution of infection resources, we only consider these strategies from the perspective of the player controlling the curing resources.

5.1. Interior node targeting. Given a network, we determine which nodes are inner nodes, i.e., which nodes have a neighborhood that is not a strict subset of another node's neighborhood. We then distribute initialization resources uniformly among these nodes. Given curing initialization budget $\mathcal{B}_b$ and setting

$$V' = \{i \in V | \mathcal{N}_i' \not\subset \mathcal{N}_j' \, \forall j \in V\},$$

we allocate our resources according to the following scheme:

$$B_i = \begin{cases} \frac{1}{|V'|}\mathcal{B}_b & \text{if } i \in V', \\ 0 & \text{else.} \end{cases} \tag{5.1}$$

For the Delta-curing problem, a similar policy is employed as follows:

$$\Delta_{b,i}(n) = \begin{cases} \frac{1}{|V'|}\mathcal{B}_b & \text{if } i \in V', \\ 0 & \text{else.} \end{cases} \tag{5.2}$$

We can also combine these schemes with the heuristic strategies in [20, 23]. Among these strategies, the best one is based on a combination of centrality measures, node degree, and closeness centrality, as well as the infection level of each node's super urn. In the case of initialization policies, as they are implemented before the onset of contagion, the allocation strategy uses only the node degree and closeness centrality.

The degree of a node is simply measured as the size of a node's neighborhood, $|\mathcal{N}_i|$. The closeness centrality of a node, C_i, is a measure used to determine a node's topological position within a network, defined as

$$C_i = \frac{1}{\sum_{j \in V} d(i,j)}, \tag{5.3}$$

where $d(i,j)$ is the shortest path length from node i to node j. Thus the node with the highest closeness centrality has the shortest average path length to every node within the network. We can then distribute resources as follows:

$$B_i = \begin{cases} \frac{|\mathcal{N}_i| C_i}{\sum_{j \in V'} |\mathcal{N}_j| C_j}\mathcal{B}_b & \text{if } i \in V', \\ 0 & \text{else.} \end{cases} \tag{5.4}$$

In the Delta-curing setup, we also account for the evolution of the super urn proportion, $S_{i,n-1}$, of each node. As before, we let C_i represent the closeness centrality of a given node i, calculated by (5.3). We then allocate resources as follows:

$$\Delta_{b,i}(n) = \begin{cases} \frac{|\mathcal{N}_i| C_i S_{i,n-1}}{\sum_{j \in V'} |\mathcal{N}_j| C_j S_{j,n-1}}\mathcal{B}_b & \text{if } i \in V', \\ 0 & \text{else.} \end{cases} \tag{5.5}$$

Alternatively, similar strategies based on "betweenness centrality" can be used [23].

5.2. Minimized node targeting. We can improve upon our initialization and curing policies by minimizing the number of nodes we target at once. We introduce a method that targets nodes in such a manner that all nodes have resources within their immediate neighborhood $\mathcal{N}_i'$ (i.e., ensuring $\bar{B}_i > 0 \; \forall i \in V$), while aiming for the fewest possible nodes. The method is described in Algorithm 2.

The algorithm starts by locating outer nodes in the network and then targets nodes directly adjacent to these outer nodes. All nodes that are within the neighborhood of the targeted nodes are then removed from the set of nodes being evaluated. Within this new, smaller set of nodes, we once again locate outer nodes and repeat the process. In the case that there are no outer nodes within the set of nodes being evaluated, all remaining nodes are added to the set of targeted nodes. The purpose of this algorithm is to effectively employ Theorem 3.5 and ensure that our curing policy does not target outer nodes, while still requiring that $\bar{B}_i > 0$ for these nodes. Once the set of outer nodes is excluded from directly receiving resources, the algorithm then works inwards, attempting to minimize the number of targeted nodes.

Once the set of nodes to be targeted V' is determined, we can once again distribute resources in one of two manners as in the previous strategy set. We can distribute resources to these nodes uniformly according to (5.1), or we can distribute resources

Algorithm 2. Node targeting algorithm.

$V \leftarrow$ set of all nodes
$V' \leftarrow$ target set
$V_{test} \leftarrow$ nodes to test for this loop
Start with $V' = \emptyset$, $V_{test} = V$
while $V_{test} \neq \emptyset$ **do**
  Identify outer nodes:
  Set $V_{outer} = \{i \in V_{test} | \mathcal{N}_i' \subset \mathcal{N}_j', \text{ for some } j \in V_{test}\}$
  if $V_{outer} = \emptyset$ **then**
    $V' = V' \cup V_{test}$
    break
  end if
  Target nodes adjacent to outer nodes:
  Set $V_{inner} = V_{test} \backslash V_{outer}$
  Set $V_{added\_targets} = \{i \in V_{inner} | i \in \mathcal{N}_j' \text{ for some } j \in V_{outer}\}$
  Set $V' = V' \cup V_{added\_targets}$
  Update parameters for next loop:
  Set $V_{coverage} = \{i \in V_{test} | i \in \mathcal{N}_j' \text{ for some } j \in V_{added\_targets}\}$
  Set $V_{test} = V_{test} \backslash V_{coverage}$
end while

to more central nodes among those targeted using (5.4). Note that these heuristics are entirely based on network structure and by design will symmetrically allocate resources in accordance with Theorem 3.9. Also, by limiting the number of targeted nodes, the computational complexity is reduced.

We can apply the same techniques for the Delta-curing setup, where we consider a uniform allocation within the target set via (5.2), or use a strategy that accounts for the centrality and super urn proportion of the targeted nodes according to (5.5).

5.3. Dense networks. We now consider a limitation of the heuristics presented thus far. Our previous sets of solutions depend on the assumption that a network will possess a set of nodes whose respective neighborhoods are strict subsets of some other node's neighborhood. While this may be a useful tactic for sparse networks, it is a stringent requirement for dense networks. To account for this, we introduce a final node targeting strategy that does not have this limitation but works based on similar principles. Our previous methods utilized the idea of leveraging Theorem 3.5, in order to prevent suboptimal node targeting while still ensuring $\bar{B}_i > 0$, for every node $i \in V$. We introduce an alternative method for defining inner and outer nodes when the conditions of Theorem 3.5 are not met.

This technique, presented in Algorithm 3, begins by ranking the nodes by a centrality measure (closeness centrality for our purposes) from highest to lowest. Nodes are added one by one to the set of targeted nodes, V', in order of rank. Once the set of targeted nodes provides total coverage of the network (i.e., the neighborhoods of these nodes contain the entire network), the algorithm stops. For an additional reduction, we go through the list of targeted nodes in reverse order and remove any nodes not required to maintain full coverage of the network.

Once the target set has been determined, we can distribute resources either uniformly among these nodes via (5.1) or preferentially according to centrality as per (5.4). Likewise for the Delta-curing case, we can allocate resources uniformly among these nodes using (5.2) or based on centrality and super urn proportion of the targeted nodes via (5.5).

Algorithm 3. Node targeting algorithm for dense networks.

$V \leftarrow$ set of all nodes
$V' \leftarrow$ nodes to be targeted
$C_i \leftarrow$ centrality score of given node
Without loss of generality, assume $C_1 > C_2 > \cdots > C_N$, where N is the network size
Start with $V' = \emptyset$, $V_{coverage} = \emptyset$, n=0
while $V_{coverage} \neq V$ **do**
  Add most central node to target set:
  $V' = V' \cup \{n\}$
  Update loop parameters:
  Set $V_{coverage} = \{i \in V | i \in \mathcal{N}'_j, \text{ for some } j \in V'\}$
  n++
end while
Optional:
for $n = |V'| : 1$ **do**
  Let $V_{coverage} = \{i \in V | i \in \mathcal{N}'_j, \text{ for some } j \in V' \backslash \{n\}\}$
  if $V_{coverage} = V$ **then**
    Set $V' = V' \backslash \{n\}$
  end if
end for

Algorithm 4. Simulation setup.

$A \leftarrow$ adjacency matrix of the network
numTrials $\leftarrow$ trials to run for given policy
steps $\leftarrow$ number of time steps for each trial
initPolicy $\leftarrow$ initialization policy (uniform in no control case)
curePolicy $\leftarrow$ Delta-curing policy (uniform in no control case)
Assign $(\mathbf{B}, \mathbf{R})$ using *initPolicy* (under initialization budget)
for s = 1 : *numTrials* **do**
  $\vec{Z}_s \leftarrow$ RunTrial$(A, \mathbf{R}, \mathbf{B}, \mathcal{B}_r, \mathcal{B}_b,$ *steps*, *curePolicy*$)$
end for
RunTrial$A, \mathbf{R}, \mathbf{B}, \mathcal{B}_r, \mathcal{B}_b,$ *steps*, *curePolicy*
Initialize $S_{i,0}$ using R_i and B_i for all $i \in V$
for t = 1 : *steps* **do**
  Assign $\Delta_{b,i}(t)$, $\Delta_{r,i}(t)$ using *curePolicy* (under curing budget)
  Generate $\vec{Y} \sim$ **Uniform**$([0,1])^N$
  if $Y_i \leq S_{i,t-1}$ **then**
    $Z_{i,t} = 1$
  else
    $Z_{i,t} = 0$
  end if
  Update $S_{i,t}$ using $\Delta_{r,i}$ and $\Delta_{b,i}(t)$ for all $i \in V$
end for

6. Simulation results. We evaluate the performance of the heuristic policies detailed in section 5 and compare them to "optimal" strategies obtained through gradient descent (note that we say "optimal" because the gradient descent only converges to an optimal policy for the one-step proxy measures we introduced in both the initialization and Delta-curing cases). Algorithm 4 provides the general format of the simulations, while specifics of both the initialization and Delta-curing setups are provided in sections 6.2 and 6.3, respectively.

In order to ensure effectiveness of different strategies, we test these policies for

a few different networks, each of which is depicted in Figure 2. These networks include algorithmically generated 100 node Barabasi–Albert networks, depicted in Figures 2(a) and 2(b), with average densities of 1.98 and 18.9 connections per node, respectively. The third network, depicted in Figure 2(c), was generated by a tool [39] that crawled through posts in a Facebook group in order to establish connections between users based on interactions on user posts (via either commenting or "likes" on the post or comments). The resulting 1,363 node graph contains an average density of 3.56 connections per node and provides a topology for a real-world social network.

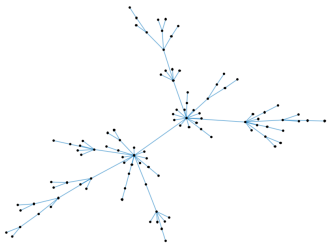

(a) Barabasi–Albert network with 100 nodes and 99 edges.

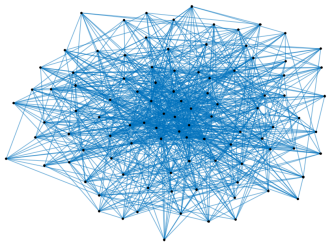

(b) Barabasi–Albert network with 100 nodes and 945 edges.

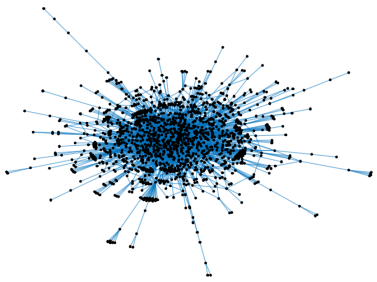

(c) Facebook network with 1,363 nodes and 2,425 edges.

FIG. 2. *Mid-to-large–scale networks used for simulation purposes. Adjacency matrices for each of these networks is available online at* https://bit.ly/2ygLEqg.

6.1. Targeting algorithms. Figure 3 depicts a comparison of the different targeting algorithms for the sparse network seen in Figure 2(a). For this specific network, the target set in Figure 3(b) is a subset of the set of inner nodes in Figure 3(a), which itself is a subset of the target set shown in Figure 3(c). While this is not a general rule, it is a common trend for very low-density networks. Algorithm 3 is not designed for use over sparse networks and does not successfully leverage Theorem 3.5. Thus, we know that distributing resources among this set of target nodes will result in inefficiencies. If we compare this to how Algorithm 3 applies to denser networks, as shown in Figure 3(d), we observe a significant reduction in the size of the set of target nodes. Algorithm 2, on the other hand, will be entirely ineffective at determining a subset of nodes to target for this same network, due to the lack of nested nodes.

6.2. Initialization trials. We start by comparing solutions to Problem 3.1 and evaluating their effectiveness for various network compositions. Table 1 provides an overview of the considered initialization strategies. More detailed explanations of

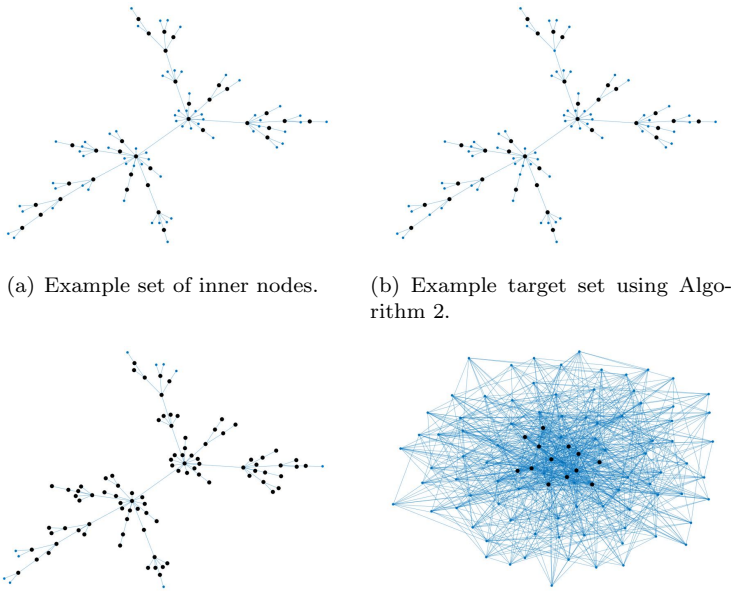

(a) Example set of inner nodes.

(b) Example target set using Algorithm 2.

(c) Example target set using Algorithm 3 for sparse networks.

(d) Example target set using Algorithm 3 for dense networks.

FIG. 3. *Comparison of target sets for various node targeting algorithms. In each case, the set of targeted nodes is enlarged and colored black. It is important to note that Figure* 3(c)*, as depicted here, does not apply the optional component of Algorithm* 3*. In fact, if we instead choose to apply that option, we end up getting the same target set as shown in Figure* 3(b)*.*

these strategies can be found in sections 3 and 5. It is important to note that we do not include any trials that use the optional component of Algorithm 3, simply because we have the same target set as Algorithm 2 when applied to the networks in Figures 2(a) and 2(c), and it does not provide any further reduction to the target set when applied to the network in Figure 2(b). The same reasoning applies to our curing strategies in Table 2.

We follow the same process for each of the three networks in Figure 2. We begin with equal initialization budgets for both red and black resources (i.e., take $\mathcal{B}_b = \mathcal{B}_r$) and let the infection initialization be uniform over the network (i.e., take $R_i = \frac{\mathcal{B}_r}{N}$ for all $i \in V$). We choose our curing initialization $\mathbf{B}$ in accordance with the desired strategy from Table 1. Then, we run the Pólya network contagion process with $\Delta_{r,i}(n) = \Delta_{b,i}(n) = \Delta$ fixed for each node $i \in V$ and each time n. We record the draw value, $Z_{i,n}$, for each node i at each time n. The draw values are averaged over the number of trials that are run and then averaged over the set of nodes in the network to obtain an empirical measure for the average infection rate $\tilde{I}_n$ at a given time n. The resulting output of these trials can be seen in Figure 4.

We start our comparison with a few observations. First, as expected, the solution determined using the gradient descent algorithm results in the lowest empirical average infection rate at time $n = 1$. For time $n > 1$, however, this is not always the case. Figure 4(c) is of particular interest, as we clearly see that the gradient descent strategy begins with a much lower infection rate but over time is surpassed by strategy (vi), which targets a reduced subset of nodes in accordance with Algorithm 2, with

TABLE 1
Initialization strategies.

(i)	Constrained gradient descent on a simplex: Find B_i using Algorithm 1 on the function $\tilde{I}_1$
(ii)	Uniformly allocate the budget to all nodes: $B_i = \frac{\mathcal{B}_b}{N}$
(iii)	Uniformly allocate the budget to inner nodes: $B_i = \begin{cases} \frac{1}{\lvert V' \rvert}\mathcal{B}_b & \text{if } i \in V', \\ 0 & \text{else,} \end{cases}$ where $V' = \{k \in V \lvert \mathcal{N}'_k \not\subset \mathcal{N}'_j \, \forall j \in V\}$
(iv)	Ratio of degree and closeness centrality (inner): $B_i = \begin{cases} \frac{\lvert\mathcal{N}_i\rvert C_i}{\sum_{j\in V'} \lvert\mathcal{N}_j\rvert C_j}\mathcal{B}_b & \text{if } i \in V', \\ 0 & \text{else} \end{cases}$
(v)	Allocate as in (iii) using V' from Algorithm 2
(vi)	Allocate as in (iv), using V' from Algorithm 2
(vii)	Allocate as in (iii) using V' from Algorithm 3 (without Optional)
(viii)	Allocate as in (iv), using V' from Algorithm 3 (without Optional)
(ix)	Ratio of degree and closeness centrality: $B_i = \frac{\lvert\mathcal{N}_i\rvert C_i}{\sum_{j\in V} \lvert\mathcal{N}_j\rvert C_j}\mathcal{B}_b$

resources being focused more around central nodes.

The fact that the best strategy at time $n = 1$ is not optimal as n increases supports our hypothesis that the optimal solution will change for increasing values of n and as $n \to \infty$. It is important to note, however, that our gradient algorithm merely converges to the optimal solution at time $n = 1$, and thus the solution we have from using this algorithm is not necessarily optimal, as the algorithm has not been run indefinitely. Indeed, we note that had we either chosen to use a finer filter when implementing the limit minimization rule, as described in Algorithm 1, or increased the stopping time, or even chosen a step size that allowed for faster convergence, we would have seen an improved performance using this strategy (at least at time $n = 1$).

Regardless, we consistently see that targeting a reduced subset of nodes, whether via Algorithm 2 or 3, results in performance that is close to or better than that obtained using gradient descent. These strategies, (vi) and (viii), always provide the best results of all the heuristic strategies being compared. This illustrates that we obtain near-optimal performance while only distributing resources among a greatly reduced subset of nodes. Performance of the other heuristic strategies varies greatly depending on the network, though we note that even just targeting the inner nodes using strategy (iii) sees a significant improvement over the no-control case (i.e., the results obtained using a uniform distribution of initialization resources).

We next examine whether these same results carry over to the curing case. We should note that the curing and initialization policies are by nature not identical. They do, however, mirror one another in such a way that we can at least provide a surface-level comparison between the two setups.

6.3. Delta-curing trials. We consider the one-sided curing problem in (4.1). As in the initialization case, we compare the effectiveness of the various curing policies presented in section 5. A brief overview of these policies is found in Table 2.

We begin with a uniform initialization, where each node is allocated a fixed number of red balls, R, and black balls, B, with $R = B$. We then run the Pólya network contagion process. Before each draw is performed, we assign a number of curing resources to each node $\Delta_{b,i}(n)$ in accordance with the desired policy from Table 2,

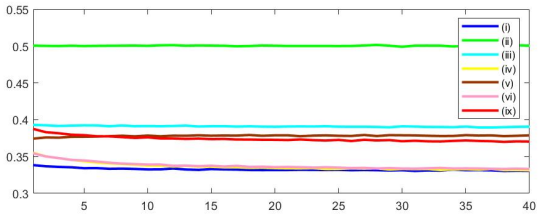

(a) Results of initialization trials for low density Barabasi–Albert network depicted in Figure 2(a).

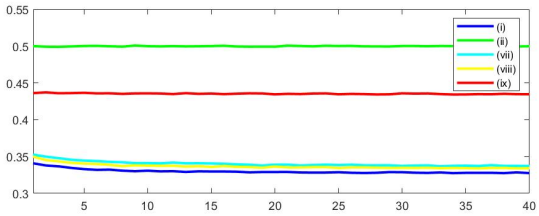

(b) Results of initialization trials for high density Barabasi–Albert network depicted in Figure 2(b).

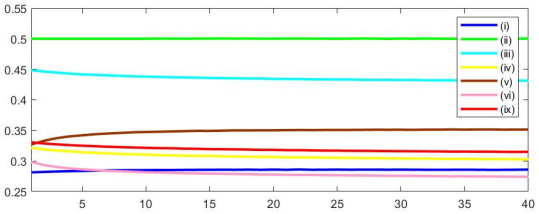

(c) Results of initialization trials for Facebook network depicted in Figure 2(c).

FIG. 4. *Plots of empirical average infection rate, $\tilde{I}_n$, versus time, n, for various initialization strategies across multiple networks. Lower values indicate lower levels of infection. In each case the network was given a fixed initialization budget of $\mathcal{B}_r = \mathcal{B}_b = 10N$, with $R_i = 10$ $\forall i \in V$, and the Pólya network contagion process was run over* 1000 *trials with a fixed* $\Delta = 5$.

while the infection resources are assigned uniformly, $\Delta_{r,i}(n) = \frac{\mathcal{B}_r}{N}$ (with a fixed budget $\mathcal{B}_b = \mathcal{B}_r$ for each time n). As in the initialization case, we record the draw value, $Z_{i,n}$, for each node $i \in V$ at each time n. The draw values are averaged over the number of trials and the set of nodes to obtain a measure for the empirical average infection rate, $\tilde{I}_n$. The results of these simulations are depicted in Figure 5.

We first remark that unlike the initialization case, targeting a reduced subset of nodes based on centrality does not always lead to the best performance. This is evidenced by both Figures 5(a) and 5(c), where strategy (vi) appears to perform worst of all the heuristics apart from the no-control case. In fact, among the three network setups, the heuristics vary wildly in performance, with the best performing heuristic from [20] performing extremely well in Figure 5(c), but performing significantly worse than the other techniques in Figure 5(b). The various strategies also perform to different degrees of success for different values of n (see strategy (vi) in Figure 5(c)).

Among the three networks, the strategies that perform the most consistently well are strategies that target only the inner set of nodes (strategies (iii) and (iv)

TABLE 2
Delta-curing strategies.

(i)	Constrained gradient descent on a simplex: Find $\Delta_{b,i}(n)$ using Algorithm 1 on the function $E[\tilde{S}_n \vert \mathcal{F}_{n-1}]$
(ii)	Uniformly allocate the budget to all nodes: $\Delta_{b,i}(n) = \frac{\mathcal{B}_b}{N}$
(iii)	Uniformly allocate the budget to inner nodes: $\Delta_{b,i}(n) = \begin{cases} \frac{1}{\vert V' \vert}\mathcal{B}_b & \text{if } i \in V', \\ 0 & \text{else,} \end{cases}$ where $V' = \{k \in V \vert \mathcal{N}'_k \not\subset \mathcal{N}'_j \, \forall j \in V\}$
(iv)	Ratio of degree, closeness centrality, and super urn proportion (inner): $\Delta_{b,i}(n) = \begin{cases} \frac{\vert\mathcal{N}_i\vert C_i S_{i,n-1}}{\sum_{j\in V'} \vert\mathcal{N}_j\vert C_j S_{j,n-1}}\mathcal{B}_b & \text{if } i \in V', \\ 0 & \text{else} \end{cases}$
(v)	Allocate as in (iii) using V' from Algorithm 2
(vi)	Allocate as in (iv), using V' from Algorithm 2
(vii)	Allocate as in (iii) using V' from Algorithm 3 (without Optional)
(viii)	Allocate as in (iv), using V' from Algorithm 3 (without Optional)
(ix)	Ratio of degree, closeness centrality, and super urn proportion: $B_i = \frac{\vert\mathcal{N}_i\vert C_i S_{i,n-1}}{\sum_{j\in V} \vert\mathcal{N}_j\vert C_j S_{j,n-1}}\mathcal{B}_b$

and the analogous strategies for denser networks, (vii) and (viii)). It is interesting that these strategies would perform most consistently for the Delta-curing problem but would perform far from optimally for the one-time initialization policy as seen in Figure 4. The main difference between these setups is that the solutions to the Delta-curing problem are applied continually as the infection evolves, while solutions to the initialization problem are applied once. Obviously, some of the benefits of targeting a greatly reduced set of nodes, as is seen in Figure 3(b), are lost when they do not properly account for changes in the urn compositions. Interestingly, for denser networks, targeting a greatly reduced set of nodes as in Figure 3(d) was still very effective, most likely as a direct result of the high network density.

In Figure 5 we did not show the performance using gradient descent due to the high computation burden for networks of such scale. We still provide a comparison on a smaller scale in order to see how far from "optimal" our various heuristics are. We provide such a comparison in Figure 6. While it is clear that the gradient descent algorithm does marginally improve the empirical average infection rate, we see that strategy (iv) still performs comparably, at a fraction of the computational cost.

7. Conclusion. We considered the (infinite-memory) Pólya network contagion model and formulated (preemptive and reactive) control policies to mitigate the spread of infection. We framed these problems in two lights: both as a one-time initialization problem and as a continual control problem that is an extension of contagion curing methods in prior work. We provided simplified metrics for the one-sided finite horizon version of these problems and proved the existence of optimal policies. We developed effective heuristics that functioned primarily by modifying the scope of solutions to only a limited number of nodes. We then compared these heuristics to the one-step optimal policies that were obtained through use of gradient descent. We established underlying results for the initialization problem proving that optimal initialization strategies would only allocate resources to inner nodes. We showed how network symmetry impacts the allocation of resources for optimal policies. We also demonstrated that the proposed heuristics perform comparably to the one-step optimal policy in certain situations. Finally, we framed the Delta-curing problem as

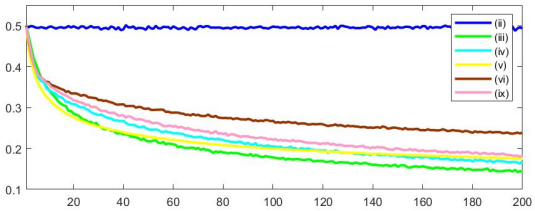

(a) Results of curing trials for low density Barabasi–Albert network depicted in Figure 2(a).

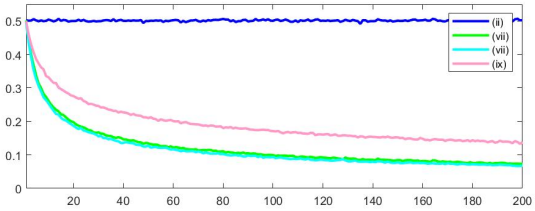

(b) Results of curing trials for high density Barabasi–Albert network depicted in Figure 2(b).

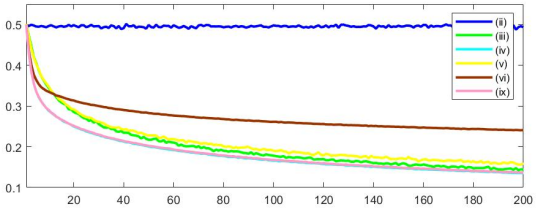

(c) Results of curing trials for Facebook network depicted in Figure 2(c).

FIG. 5. *Plots of empirical average infection rate, $\tilde{I}_n$, versus time, n, for various curing strategies across multiple networks. Lower values indicate lower levels of infection. In each case the network was initialized with $R_i = B_i = 10$ $\forall i \in V$, and the Pólya network contagion process was run over* 500 *trials with a fixed budget $\mathcal{B}_r = \mathcal{B}_b = 10N$, with $\Delta_r = 10$ fixed.*

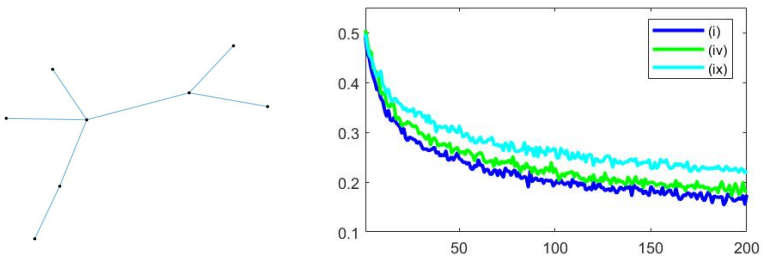

FIG. 6. *Empirical average infection rate, $\tilde{I}_n$, versus time, n, for an* 8 *node network. The performance of gradient descent over the expected network exposure, $E[\tilde{S}_n|\mathcal{F}_{n-1}]$, is compared to the two best performing heuristics from Figure* 5(a). *The network was initialized with $R_i = B_i = 10$ $\forall i \in V$, and the Pólya network contagion process was run over* 500 *trials with budget $\mathcal{B}_r = \mathcal{B}_b = 80$ and fixed $\Delta_r = 10$.*

a two-person zero-sum game on the expected network exposure, and we proved the existence of a Nash equilibrium. We borrowed ideas from the initialization problems in order to implement heuristics for the Delta-curing problem and demonstrated that these strategies can perform near-optimally at a fraction of the computational expense. Future directions include analogous optimization problems when underlying resource distribution is hidden from the observer, studying the initialization problem in a game-theoretic framework, and investigating the spread of the Pólya contagion process in networks with unknown topologies.

Acknowledgment. The authors sincerely thank Tamás Linder for considerably shortening the proof of Lemma 3.4.

REFERENCES

[1] E. Adar and L. A. Adamic, *Tracking information epidemics in blogspace*, in Proceedings of the IEEE/WIC/ACM International Conference on Web Intelligence, IEEE, Washington, DC, 2005, pp. 207–214.

[2] H. J. Ahn and B. Hassibi, *Global dynamics of epidemic spread over complex networks*, in Proceedings of the 52nd IEEE Conference on Decision and Control, IEEE, Washington, DC, 2013, pp. 4579–4585.

[3] F. Alajaji and T. Fuja, *A communication channel modeled on contagion*, IEEE Trans. Inform. Theory, 40 (1994), pp. 2035–2041.

[4] R. Ash and C. Doléans-Dade, *Probability and Measure Theory*, Academic Press, New York, 2000.

[5] D. P. Bertsekas, *Nonlinear Programming*, 2nd ed., Athena Scientific, Belmont, MA, 1999.

[6] D. P. Bertsekas, A. Nedić, and A. E. Ozdaglar, *Convex Analysis and Optimization*, Athena Scientific, Belmont, MA, 2003.

[7] N. Biggs, *Algebraic Graph Theory*, 2nd ed., Cambridge University Press, Cambridge, UK, 1994.

[8] J. A. Bondy and U. S. R. Murty, *Graph Theory with Applications*, Elsevier, New York, 1976.

[9] C. Borgs, J. Chayes, A. Ganesh, and A. Saberi, *How to distribute antidote to control epidemics*, Random Structures Algorithms, 37 (2010), pp. 204–222.

[10] M. Chen and M. Kuba, *On generalized Pólya urn models*, J. Appl. Probab., 50 (2013), pp. 1169–1186.

[11] F. Chung, S. Handjani, and D. Jungreis, *Generalizations of Pólya's urn problem*, Ann. Combin., 7 (2003), pp. 141–153.

[12] D. Easley and J. Kleinberg, *Networks, Crowds and Markets: Reasoning about a Highly Connected World*, Cambridge University Press, Cambridge, UK, 2010.

[13] F. Eggenberger and G. Pólya, *Über die Statistik verketteter Vorgänge*, Z. Angew. Math. Mech., 3 (1923), pp. 279–289.

[14] A. Fazeli and A. Jadbabaie, *On Consensus in a Correlated Model of Network Formation Based on a Pólya Urn Process*, ESE paper 614, University of Pennsylvania, Philadelphia, PA, 2011, https://repository.upenn.edu/ese_papers/614/.

[15] W. Feller, *An Introduction to Probability Theory and Its Applications, Vol.* II, 2nd ed., Wiley, New York, 1971.

[16] M. Garetto, W. Gong, and D. Towsley, *Modeling malware spreading dynamics*, in Proceedings of IEEE INFOCOM 2003: The Twenty-Second Annual Joint Conference of the IEEE Computer and Communications Societies, Vol. 3, IEEE, Washington, DC, 2003, pp. 1869–1879.

[17] C. D. Godsil and G. F. Royle, *Algebraic Graph Theory*, Grad. Texts in Math. 207, Springer, New York, 2001.

[18] G. R. Grimmett and D. R. Stirzaker, *Probability and Random Processes*, 3rd ed., Oxford University Press, Oxford, UK, 2001.

[19] G. Harrington, F. Alajaji, and B. Gharesifard, *Infection-curing games over Pólya contagion networks*, in Proceedings of the 16th Canadian Workshop on Information Theory (CWIT), Hamilton, ON, Canada, 2019, pp. 1–6.

[20] M. Hayhoe, *A Pólya Urn Stochastic Model for the Analysis and Control of Epidemics on Networks*, Master's thesis, Queen's University, Kingston, ON, Canada, 2017.

[21] M. Hayhoe, F. Alajaji, and B. Gharesifard, *A Pólya urn-based model for epidemics on*

networks, in Proceedings of the 2017 American Control Conference, Seattle, WA, 2017, pp. 358–363.

[22] M. Hayhoe, F. Alajaji, and B. Gharesifard, *A Pólya contagion model for networks*, IEEE Trans. Control Network Syst., 5 (2018), pp. 1998–2010.

[23] M. Hayhoe, F. Alajaji, and B. Gharesifard, *Curing epidemics on networks using a Pólya contagion model*, IEEE/ACM Trans. Networking, 27 (2019), pp. 2085–2097.

[24] A. R. Hota and S. Sundaram, *Game-theoretic protection against networked SIS epidemics by human decision-makers*, IFAC-PapersOnLine, 51 (2019), pp. 145–150.

[25] K. Kandhway and J. Kuri, *How to run a campaign: Optimal control of SIS and SIR information epidemics*, Appl. Math. Comput., 231 (2014), pp. 79–92.

[26] M. Khouzani, S. Sarkar, and E. Altman, *Optimal control of epidemic evolution*, in Proceedings of IEEE INFOCOM 2011: The Thirtieth Annual Joint Conference of the IEEE Computer and Communications Societies, IEEE, Washington, DC, 2011, pp. 1683–1691.

[27] L. Kim, M. Abramson, K. Drakopoulos, S. Kolitz, and A. Ozdaglar, *Estimating social network structure and propagation dynamics for an infectious disease*, in Proceedings of the International Conference on Social Computing, Behavioral-Cultural Modeling, and Prediction, Springer, New York, 2014, pp. 85–93.

[28] Q.-H. Liu, L.-F. Zhong, W. Wang, T. Zhou, and H. E. Stanley, *Interactive social contagions and co-infections on complex networks*, Chaos, 28 (2018), 013120.

[29] N. Madar, T. Kalisky, R. Cohen, D. Ben-Avraham, and S. Havlin, *Immunization and epidemic dynamics in complex networks*, European Phys. J. B, 38 (2004), pp. 269–276.

[30] W. Mei and F. Bullo, *Competitive propagation: Models, asymptotic behavior and quality-seeding games*, IEEE Trans. Network Sci. Engrg., 4 (2017), pp. 83–99.

[31] W. Mei, S. Mohagheghi, S. Zampieri, and F. Bullo, *On the dynamics of deterministic epidemic propagation over networks*, Annu. Rev. Control, 44 (2017), pp. 116–128.

[32] J. Nash, *Equilibrium points in n-person games*, Proc. Natl. Acad. Sci. USA, 36 (1950), pp. 48–49.

[33] C. Nowzari, V. M. Preciado, and G. J. Pappas, *Analysis and control of epidemics: A survey of spreading processes on complex networks*, IEEE Control Syst. Mag., 36 (2016), pp. 26–46.

[34] C. Nowzari, V. M. Preciado, and G. J. Pappas, *Optimal resource allocation for control of networked epidemic models*, IEEE Trans. Control Network Syst., 4 (2017), pp. 159–169.

[35] P. E. Paré, J. Liu, C. L. Beck, A. Nedić, and T. Başar, *Multi-competitive viruses over static and time-varying networks*, in Proceedings of the American Control Conference (ACC), Seattle, WA, 2017, pp. 1685–1690.

[36] G. Pólya, *Sur quelques points de la théorie des probabilités*, Ann. Inst. H. Poincaré, 1 (1930), pp. 117–161.

[37] G. Pólya and F. Eggenberger, *Sur l'interprétation de certaines courbes de fréquences*, Comptes Rendus C. R., 187 (1928), pp. 870–872.

[38] V. M. Preciado, M. Zargham, and D. Sun, *A convex framework to control spreading processes in directed networks*, in Proceedings of the 48th Annual Conference on Information Sciences and Systems, IEEE, Washington, DC, 2014, pp. 1–6.

[39] B. Rieder, *Studying Facebook via data extraction: The Netvizz application*, in Proceedings of the 5th ACM Web Science Conference, ACM, New York, 2013, pp. 346–355.

[40] E. M. Rogers, *Diffusion of Innovations*, 5th ed., Free Press, New York, 2003.

[41] N. A. Ruhi, H. J. Ahn, and B. Hassibi, *Analysis of Exact and Approximated Epidemic Models over Complex Networks*, preprint, https://arxiv.org/abs/1609.09565, 2016.

[42] P. Van Mieghem, J. Omic, and R. Kooij, *Virus spread in networks*, IEEE/ACM Trans. Netw., 17 (2009), pp. 1–14.

[43] L. Wang and J. T. Wu, *Characterizing the dynamics underlying global spread of epidemics*, Nature Commun., 9 (2018), 218.

[44] W. Wang, Q.-H. Liu, J. Liang, Y. Hu, and T. Zhou, *Coevolution spreading in complex networks*, Phys. Rep., 820 (2019), pp. 1–51.

[45] Y. Wang, D. Chakrabarti, C. Wang, and C. Faloutsos, *Epidemic spreading in real networks: An eigenvalue viewpoint*, in Proceedings of the 22nd IEEE International Symposium on Reliable Distributed Systems (Florence, Italy), IEEE, Washington, DC, 2003, pp. 25–34.