Forty-Ninth Annual Allerton Conference
Allerton House, UIUC, Illinois, USA
September 28 - 30, 2011

Market-Based Control of Epidemics

MHR. Khouzani, Santosh S. Venkatesh, and Saswati Sarkar
University of Pennsylvania

***Abstract*—We consider a single profit-maximizer healing or health-care provider that administers treatment to infective nodes during an epidemic outbreak in a homogeneous network. We derive the dynamic optimal policies of the provider for two cases: (a) when infective nodes gain permanent immunity upon treatment (SIRD model), and (b) when treated nodes are again susceptible to future re-infection (SIS model). We show that in the case of the SIRD model the optimal policy is to defer provision of healing until an infection threshold is reached and then to provide healing with maximum intensity. In the case of the SIS model the optimal policy is a Most-Rapid-Approach-Path (MRAP) to an intermediate density of infection.**[1]

I. Introduction

a) Motivation and Overture: Upon a new outbreak, counter-measures such as immunizing the susceptibles and healing the infectives can help contain the spread of the infection. In this paper, we investigate the cases in which a firm (a monopolist) with selfish economic incentive is providing treatment to infective nodes and healing them. In particular, the firm can dynamically decide when and with what intensity to provide healing in response to changes in the system, so as to maximize its own profit. At any given time, the provider faces the following trade-off: healing infectives yields an instantaneous reward, but it also reduces the density of infection, which in turn reduces the firm's aggregate future reward. The latter is due to the fact that each infective can potentially convert a susceptible to an infective upon future contact. Using tools from optimal control theory, we analytically derive the dynamic optimal policies of the provider.

We consider a network of *nodes*[2], i.e., humans, computers, communication devices, etc., a non-zero fraction of which are infected at time zero.[3] In one class of problems infective nodes are no longer prone to the infection upon receiving treatment. In this case, they are referred to as *recovered*. In general, the infective nodes, left untreated, may *die* at a certain rate as well. In such a case, a possible transformation for a node is from infective to the *dead* state. In the case of permanent immunity after treatment, the system is referred to as Susceptible-Infective-Recovered/Dead, or in short, an *SIRD* system.

Depending on the nature of the infection and the state of the art, it might be that the treatment does not grant the nodes immunity against future contraction. In this case, healing of an infective node changes its state to susceptible, as it will be vulnerable to re-infection. Such systems are referred to as Susceptible-Infective-Susceptible, or SIS systems, in brief.

b) Decision problems in an SIRD model: The firm may decide when to provide healing to the network. This binary decision (to provide or not to provide) can be relaxed to include the intensity (rate) of the provisioning. Intermediate values of provisioning can be interpreted as partial refusals of the infectives. The monopolist is rewarded per each node it treats. The objective of the monopolist is to maximize its overall profit over its interval of optimization.

We establish (in section III) that the optimal provision of healing is in the following simple class: no provision of healing till a particular time epoch, and subsequently provisioning with maximum intensity. In the language of optimal control theory, we prove that the optimal healing schedule, as a function of time, is *bang-bang* with at most one jump from zero to its maximum value. Of course, while a delay in countering epidemics may be profitable for the company it is damaging to society both directly for individuals and through negative externalities of the infection. The bang-bang result arises despite the relaxation to permit partial provisioning in principle. Hence, for an SIRD epidemic it is not optimal to ration treatment; it is all or nothing. These structural results are derived despite the lack of a closed form solution—a consequence of the non-linear nature of the epidemic evolution.

c) Decision problems in an SIS model: In an SIS setting, the difference is that there is no recovered state. As in the SIRD case, the monopolist is faced with a decision of when to start providing healing and with what intensity. Intermediate values of intensity imply granting partial access to infective nodes. We establish (in section VI) that in this case the optimal policy of the provider in general is to reach an intermediate density of infection as rapidly as possible and maintain the density of infection at that level. We calculate this optimum level and completely characterize the optimal control as a function of the current level of infection.

d) Related Work: Optimal dynamic defense policies against the outbreak of an epidemic in networks are investigated in [1]–[3] among others. However, in these papers the economic incentive of the defender is ignored, and the defender is assumed to be a social beneficent which dynamically manages the networks resources depending on the varying risk of infection. Recently, the role of economic incentives in security decisions in a network has seen a new wave of research interest([4]–[7] among others). However, unlike our work, these papers mostly deal with the economic incentive of individual nodes to seek prevention or treatment measures. Likewise, the economic incentive of the security providers is not included in their model. To the best of our knowledge, this is the first work that analytically investigates the dynamic decision of the network security provider as the infection propagates in the network.

[1]Support for this paper came from the following grants: CNS 1115547, CNS 0915697, CNS 0915203, CNS 0914955.

[2]Throughout the paper, there is a close correspondence between a human society and a communication network. These analogies are often spelled out.

[3]Time zero is the maximum of these two events: the epidemic is detected and the counter-measure is developed.

II. SIRD EPIDEMICS

Let $S(t)$, $I(t)$, $R(t)$ and $D(t)$ respectively represent the density of susceptible, infective, recovered and dead nodes per unit area in a vast area of consideration. At any given time $S(t) + I(t) + R(t) + D(t) =$ constant, where the constant is the (homogeneous) density of nodes per unit area which we may, without loss of generality, take to be 1. In light of this equality, we can represent the *state* of the system by the triplet $(S(t), I(t), D(t))$, as $R(t)$ can then be readily derived. At the time of the outbreak of the infection, that is at time zero, some nodes are infected: $0 < I(0) = I_0 \leq 1$. Without loss of generality, we may assume $R(0) = D(0) = 0$. Thus, $S(0) = 1 - I_0$.

We now model the dynamics of the propagation of the infection as a deterministic epidemic model. Epidemic modeling based on the classic Kermack-Mckendrick model [8] has extensively been used to analyze the spread of communicable diseases in a human society [9] as well as the propagation of malware in wired and wireless networks [10], [11]. These works show, through simulations and matching with actual data, that when the number of nodes in a network is large, the deterministic epidemic models can represent the dynamics of the spread of the infection with acceptable accuracy.

A susceptible is infected whenever it receives a copy of the source of infection from an infective. This requires the specific susceptible-infective pair to *contact* each other. In a human society, this can refer to the physical proximity of the pair to each other. Similarly, in a sensor/actuator network or a Delay Tolerant Network (DTN[4]), a contact occurs when a pair of susceptible-infective nodes roam into the communication range of each other. In communication networks such as 3G/4G cellular networks as well as p2p networks, a contact occurs when an infective sender generates a valid IP addresses of a susceptible node and can either infiltrate it or entice it to accept the communication.

The overall rate of such contacts per unit area at time t is thus proportional to the product $S(t)I(t)$. Under a *homogeneous mixing* assumption, the rate of contacts is linearly proportional to this product, i.e., is equal to $\beta_0 I(t)S(t)$, for some constant β_0. Implicit in the assumption of homogeneous mixing is a large network of indistinguishable parties in which pairs are equally likely to come into contact.

Suppose that infective nodes recover at rate γ when provided with healing service.[5] The monopolist (theoretically) can decide to provide service only to a (random) fraction of susceptible nodes seeking healing. Let the fraction of infective nodes that the monopolist provides healing to at time t be represented by $\tilde{v}(t)$. At any given time t, $0 \leq \tilde{v}(t) \leq 1$. The total rate of healing of the infective nodes per unit area is therefore $\tilde{v}(t)\gamma I(t)$. Combine terms and write $v(t) = \tilde{v}(t)\gamma$. Then $v(t)$ represents the control function of the monopolist. At any given time t, $v(t)$ must satisfy the *control constraint* $0 \leq v(t) \leq v_{\max}$ where $v_{\max} := \gamma$. The assumption of rationing healing by random access provisioning relaxes the problem statement to permit a partial entry into the market place from a purely binary decision of either to provide or to not provide healing. As we will shall see, the binary decision is retrieved in the optimal solution.

[4]A network of mobile wireless nodes where the communication range of the nodes is much smaller than the area they roam on.

[5]This can refer to provision of a medicine, distribution of a security patch, or skilled treatment.

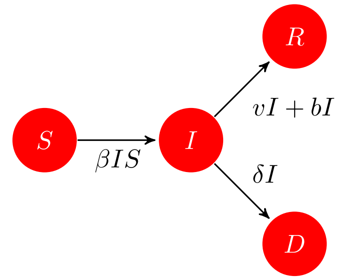

Fig. 1. State transitions in the deterministic model for the SIRD system and a monopolist. $v(t)$ is the control function of the monopolist.

In general, infective nodes left untreated may spontaneously recover or, conversely, die at rates b and δ, respectively. The corresponding rates can be taken as zero if such processes do not occur.

Our discussion motivates the following set of ODEs: the evolution of the system is represented by

$$\dot{S}(t) = -\beta_0 I(t)S(t), \tag{1a}$$
$$\dot{I}(t) = \beta_0 I(t)S(t) - v(t)I(t) - bI(t) - \delta I(t), \tag{1b}$$
$$\dot{D}(t) = \delta I(t), \tag{1c}$$

where $(S(\cdot), I(\cdot), D(\cdot))$ constitutes the system state and $v(\cdot)$ constitutes the control function of the monopolist. The system must satisfy the initial constraints

$$I(0) = I_0, \quad S(0) = 1 - I_0, \quad D(0) = 0, \tag{2}$$

as well as the evolutionary constraints

$$0 \leq S(t), I(t), D(t), \tag{3a}$$
$$S(t) + I(t) + D(t) \leq 1. \tag{3b}$$

Henceforth, wherever not ambiguous, we drop the dependence on t and make it implicit. Fig. 1 illustrates the transitions between different states of nodes and the notations used.

The monopolist is rewarded 1 (normalized) unit per healing an infective node. The overall reward of the monopolist over the interval $[0, T]$ is thus proportional to the following:[6]

$$J(v) = \int_0^T v(t)I(t)\,dt \tag{4}$$

The monopolist seeks to choose its control function $v(\cdot)$ so as to maximize J, subject to satisfying the state constraints in (3) and ensuring the following control constraint:

$$0 \leq v(t) \leq v_{\max}. \tag{5}$$

Any piecewise continuous function $v(t)$ that satisfies the above constraint is referred to as an *admissible* control.

III. SIRD: STRUCTURAL RESULTS

We start with a simple lemma that allows us to deal with optimal controls without state constraints. The proof is straightforward and is omitted for brevity.

[6]T is likely chosen long enough that by the end of it, the fraction of remaining infectives are insignificant.

Lemma 1: For all admissible controls $v(t)$, the state constraints of (3) are automatically satisfied. Moreover, for $I_0 > 0$, the inequalities $S(t), I(t) \geq 0$ are strictly satisfied for all $t \in [0, T]$.

In what follows, we use tools from control theory to derive the optimal healing policy. Define the *Hamiltonian* as:

$$\mathcal{H}(v, I, S, D, \lambda_S, \lambda_I, \lambda_D) := vI + (\lambda_I - \lambda_S)\beta IS - \lambda_I bI \\ -\lambda_I \delta I - \lambda_I vI.$$

Let $\psi(t) := (1 - \lambda_I)I$, which transforms the Hamiltonian to

$$\mathcal{H} = \psi v + (\lambda_I - \lambda_S)\beta IS - \lambda_I \delta I - \lambda_I bI. \quad (6)$$

Following Pontryagin's Maximum Principle [12, p.109], if $v(t)$ is an optimal solution and the states correspond to that optimal solution, then there exist continuous and piecewise continuously differentiable functions $\lambda_S(\cdot)$, $\lambda_I(\cdot)$, $\lambda_D(\cdot)$, referred to as *adjoint* or *co-state* functions, such that for all $t \in [0, T]$,

$$v(t) \in \arg \max_{0 \leq \underline{v} \leq v_{\max}} \mathcal{H}(\underline{v}, I, S, D, \lambda_S, \lambda_I, \lambda_D), \quad (7)$$

and at every t where $v(\cdot)$ is continuous,

$$\dot{\lambda}_S = -\frac{\partial \mathcal{H}}{\partial S} = -(\lambda_I - \lambda_S)\beta I,$$
$$\dot{\lambda}_I = -\frac{\partial \mathcal{H}}{\partial I} = -\frac{\psi}{I}v - (\lambda_I - \lambda_S)\beta S + \lambda_I \delta + \lambda_I b,$$

with the *transversality* condition

$$\lambda_S(T) = 0, \quad \lambda_I(T) = 0.$$

Referring to (6), (7) leads to the maximization of ψv, which yields the following condition for an optimal $v(t)$:[7]

$$v(t) = \begin{cases} v_{\max} & \psi(t) > 0 \\ 0 & \psi(t) < 0 \end{cases} \quad (8)$$

Since ψ is a continuous differentiable function of the states and the co-states, and since states and co-states are continuous and piecewise differentiable functions of time, we infer that ψ is a continuous and piecewise differentiable function of time. From (8), it is clear that v may only be discontinuous at times at which $\psi = 0$. For all other times, the co-state functions are differentiable and we can write:

$$\dot{\psi} = -\dot{\lambda}_I I + \frac{\psi}{I}\dot{I}$$

Let τ be a time at which $\psi = 0$. Hence

$$\dot{\psi}|_{\tau^+} = \dot{\psi}|_{\tau^-} = -(-(\lambda_I - \lambda_S)\beta IS + \lambda_I \delta I + \lambda_I bI)$$

Note that in our problem, neither the Hamiltonian nor the control constraints has explicit dependence on t. Such systems are referred to as *autonomous* systems, and it is shown that their Hamiltonian is constant [13, P.236]. Thus, for all $t \in [0, T]$, $\mathcal{H}(t) = \mathcal{H}(T)$, and we can write:

$$\dot{\psi}|_{\tau^+} = \dot{\psi}|_{\tau^-} = \mathcal{H}(\tau) = \mathcal{H}(T) = I(T) > 0.$$

[7]Note specifically that so far we are agnostic about the occasions at which $\psi = 0$. As we will argue, ψ cannot be equal to zero over an interval of non-zero length.

This shows that ψ cannot be equal to zero over an interval of non-zero length, since in that case $\dot{\psi}(t)$ should be equal to zero over such an interval too. Furthermore, it shows that $\psi(t)$ can cross zero at most once (and it would be from below zero to above zero, as time progresses). This is because otherwise, the right-hand side and the left-hand side limits of the time derivatives at consecutive zero-crossing epochs should have opposite signs (or at least one of these limits must be zero), which is impossible. Now, $\psi(T) = (1 - \lambda_I(T))I(T) = I(T) > 0$ and following its continuity in time, $\psi(t) > 0$ on an interval of non-zero length ending at T. Therefore, referring to (8), we have the following preposition about the structure of an optimal healing policy:

Proposition 1: There exists $0 \leq t_0 < T$ such that the optimal control takes the form

$$v(t) = \begin{cases} 0 & \text{if } 0 \leq t < t_0, \\ v_{\max} & \text{if } t_0 < t \leq T. \end{cases}$$

This still does not refute the possibility of $t_0 = 0$, but e.g., if $\beta S_0 > \delta + b$, which implies $\dot{I}(0) > 0$, that is the epidemic left untreated will initially grow, and if $v_{\max}$ is high, meaning that the healing is very effective, then the monopolist will not provide the healing option until a threshold is reached. Our numerical examples show that for a wide range of parameters, t_0 is strictly greater than zero.

The intuition behind theorem 1 is that the healing provider instead of throttling the market prematurely by removing the customer-generators too quickly, waits until the epidemic flourishes and ripens and then harvests it.

Remark 1: (8) has an intuitive economical explanation based on the shadow reward interpretation of the co-state functions: at any given time t, $v(t)I(t)\,dt$ is the instantaneous reward that the monopolist receives by healing the infectives, while $\lambda_I(t)v(t)I(t)\,dt$ is the reduction in the aggregate future reward as a result of reducing the density of infection. (8) states withhold healing if the reduction of the future reward of infectives outweighs the instantaneous reward of treating the infectives, and provide (maximum) healing otherwise.

Remark 2: Note that the optimal policy to defer the provision of healing treatment can be observed despite deaths of the infectives (for $\delta > 0$), unless the death rate is so large that it leads to loss in the potential customers. That is, unless δ is so large as to make $\delta + b > \beta S_0$, which means the epidemic is naturally damping, we should expect deferral of the provisioning.

Another action that helps contain the epidemic is immunization of susceptibles. Similar to the case of healing, we can define a control function for a monopolist which provides immunization as well. Let this control function be denoted by $u(t)$, with the constraint of $0 \leq u(t) \leq u_{\max}$. This transforms the expression for $\dot{S}$ in (1a) to $\dot{S} = -\beta IS - uS$.

The optimal policy of a monopolist that provides only immunization is trivial. In this case, there is no trade-off in favor of deferring the provision of the vaccine: as time progresses the susceptibles, as potential customers, are being lost to the pool of infectives. Formally, following lemma 1, and the equation for $\dot{S}$ in (1a), S is a strictly decreasing function of time. The objective of the monopoly in this case is $J = \int_0^T u(t)S(t)\,dt$. Hence, its optimal policy is $u(t) = u_{\max}$ for all $t \in [0, T]$.

In general, the monopolist can provide both immunization and healing services. Let the (normalized) reward per immunization of a susceptible be 1 and per healing an infective be ρ. Thus, the aggregate reward per unit area of the monopolist is $J = \int_0^T (u(t)S(t) + \rho v(t)I(t))\, dt$ Although we do not establish the structural properties of the optimal policies for this case, numerical examples show that when $\rho > 1$, i.e., when the monopolist is rewarded more per unit service of healing than immunization, it may not only defer the provision of healing, but the provision of immunization may be deferred as well (fig. 2). The intuition is straightforward: if serving the infective yields more reward, both deferral of immunizing the susceptibles and will result in better maturing the market toward more customers and especially more lucrative customers.

Note also that as subfig. 2(a) and 2(b) show, in general there is no order between the deferral of the immunization and healing in the optimal control either: depending on the parameters of the system, immunization or healing may be deferred longer.

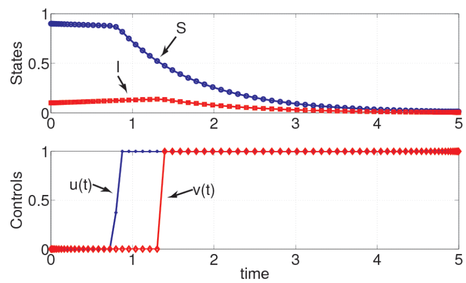

(a) $\beta = 0.3$

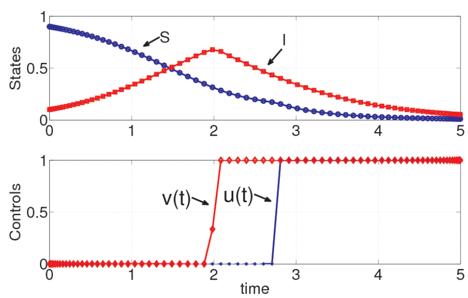

(b) $\beta = 1.5$

Fig. 2. Example of the optimal controls and the state evolution of a monopolist in the SIRD case when the monopolist can control both immunization and healing. The parameters in common are $T = 5$, $\rho = 1.5$, $I_0 = 0.1$, $\delta = 0$.

IV. SIS Epidemics

Consider now the cases in which the healing will remove the infection but leaves the infectives exposed, i.e., susceptible to the infection, maybe because of the nature of the infection which is highly evolving, or because there is no other means of treatment yet known. This model is known as the SIS model for epidemics. Suppose further that there is no mortality as a result of infection. We only investigate the case of a monopolistic provider for the SIS model. We also allow intermediate values of control, with the interpretation of limiting the accessibility/provision of the service. The system dynamics will change to the following (also refer to fig. 3):

$$\dot{S} = -\beta IS + vI, \qquad \dot{I} = \beta IS - vI$$

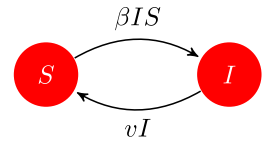

Fig. 3. State transitions in the deterministic model for the SIS system. $v(t)$ is the control functions of the monopolist.

Since S, I are now the only possible states and we have $S + I = 1$, we can write:

$$\dot{I} = \beta I(1 - I) - vI \tag{10}$$

with the state constraint: $0 \leq I(t) \leq 1$.

Lemma 2: For all admissible controls $v(\cdot)$, the state constraint of $0 \leq I(t) \leq 1$ is automatically satisfied. For $0 < I_0 < 0$, this state constraint is strictly satisfied.
The proof of this lemma, as for the case of SIRD in lemma 1 is straightforward and is omitted for brevity.

The profit function (objective functional to be maximized over $0 \leq v(t) \leq v_{\max}$):

$$J = \int_0^\infty e^{-rt} v(t)I(t)\, dt \tag{11}$$

The Hamiltonian in this case is:

$$\mathcal{H} := vI + \lambda\beta I(1 - I) - \lambda vI$$

The Hamiltonian is linear in control variable v. Pontryagin Maximum principle gives:

$$v(t) = \begin{cases} v_{\max} & \frac{\partial\mathcal{H}}{\partial v} > 0 \\ \text{singular} & \frac{\partial\mathcal{H}}{\partial v} = 0 \\ 0 & \frac{\partial\mathcal{H}}{\partial v} < 0 \end{cases}$$

where

$$\dot{\lambda} = r\lambda - \frac{\partial\mathcal{H}}{\partial I} = r\lambda - \beta\lambda + 2\beta\lambda I - v + \lambda v$$

[12, P. 132] Hence PMP cannot help us determine what the optimal control is in the case $\frac{\partial\mathcal{H}}{\partial v} = 0$ over some interval of non-zero length. The solution to the optimal control problem over such intervals is called the *singular* control. So unless we rule out the existence of such intervals, for example by proving that these are isolated time epochs as we did for the SIRD system, Pontryagin's Maximum Principle fails to identify the optimal solution.

In the sequel, we adopt a different approach, which is particularly developed to find a class of controllers, referred to as Most Rapid Approach Paths (MRAPs). We first briefly overview the statement of the results that we will apply.

V. MRAP Overview

Consider the optimal control problem:

$$\max_{u(t)} \int_0^\infty G(u(t), x(t))e^{-rt}\, dt, \tag{12}$$

$$\dot{x}(t) = f(x(t), u(t)), \quad x(0) = x_0, \tag{13}$$

$$u \in U(x), \quad U(x) \text{ compact and connected.} \tag{14}$$

Note that the last condition in our problem is satisfied as it is the box constraint $0 \leq u \leq u_{\max}$. We also have a state constraint, but according to lemma 2, they are automatically satisfied.

Definition 1 ([12] - P. 135): **Most Rapid Approach Path (MRAP)** A most rapid approach path (MRAP in short, or nearest feasible path, as [14] calls it) x^* to a given state (in general, trajectory) $\tilde{x}$ has the property:

$$|x^*(t) - \tilde{x}(t)| \leq |x(t) - \tilde{x}(t)|, \quad \forall t \in [0, \infty)$$

for all admissible state trajectories $x(\cdot)$.

More formally ([15]), define:

$$\overline{z}_t(x_0) := \sup\{z | \exists \text{ admissible path including the points } x(0) = x_0 \text{ and } x(t) = z\}$$

$$\underline{z}_t(x_0) := \inf\{z | \exists \text{ admissible path including the points } x(0) = x_0 \text{ and } x(t) = z\}$$

The MRAP from x_0 to $\hat{x}$ is the path which follows $\overline{z}(x_0)$ if $\hat{x} > x_0$ or follows $\underline{z}(x_0)$ if $\hat{x} < x_0$ until $\hat{x}$ is reached (if it ever is). MRAP's are always feasible by the following lemma (Lemma 1 in [15]):

Lemma 3: The path $\overline{z}_t(x_0)$ is an admissible path which can be generated by choosing the control from:

$$\overline{U}(y) = \{u | f(y, u) \geq f(y, \tilde{u}), \quad \text{all } \tilde{u} \in U(y)\}$$

that is, choosing a u that maximizes the time derivative of x. Similarly, the path $\underline{z}_t(x_0)$ can be generated by choosing the control from:

$$\underline{U}(y) = \{u | f(y, u) \leq f(y, \tilde{u}), \quad \text{all } \tilde{u} \in U(y)\}$$

that is, choosing a u that minimizes the time derivative of x.

Here, we quote a theorem (which we will refer to as the MRAP theorem) that determines a class of optimal control problem for which an MRAP to a derived trajectory is proven optimal ([12, Theorem 3.48 - P.136], also [16, Theorem 2.1, P.2(404)]). Consider the optimal control problem:

$$\max_{\dot{x}(t)} \int_0^\infty e^{-rt}[M(x(t)) + N(x(t))\dot{x}]\, dt \tag{15a}$$

$$\text{s.t.} \quad x(0) = x_0, \quad \dot{x}(t) \in \Omega(x(t)), \quad \underline{x} \leq x \leq \overline{x} \tag{15b}$$

where M and N are continuously differentiable and Ω depends continuously on x. Define

$$\iota(x) := rN(x) + \frac{d}{dx}M(x) \tag{16}$$

and consider the equation:

$$\iota(\hat{x}) = rN(\hat{x}) + \frac{d}{dx}M(\hat{x}) = 0 \tag{17}$$

The MRAP Theorem: Suppose that (17) has a unique solution $\hat{x}$ and that $\hat{x}$ is admissible, i.e.,

$$0 \in \Omega(\hat{x}), \quad \text{and} \quad \underline{x} \leq \hat{x} \leq \overline{x} \tag{18a}$$

Assume further that

$$\iota(x) \begin{Bmatrix} > \\ < \end{Bmatrix} 0 \quad \text{for} \quad \begin{cases} \underline{x} \leq x < \hat{x} \\ \hat{x} < x \leq \overline{x} \end{cases} \tag{18b}$$

and, finally, that for any admissible trajectory $x(\cdot)$,

$$\lim_{t\to\infty} e^{-rt} \int_{x(t)}^{\hat{x}(t)} N(\xi)\, d\xi \geq 0. \tag{18c}$$

Under these conditions if there exists an MRAP $x^*(t)$ from x_0 to $\hat{x}$ then it is the optimal solution.

If, on the other hand, there exists no admissible solution $\hat{x}$ of (17), and if

$$\iota(x) \begin{Bmatrix} > \\ < \end{Bmatrix} 0 \quad \text{for all} \quad [\underline{x}, \overline{x}] \tag{19}$$

then the MRAP to $\begin{Bmatrix} \overline{x} \\ \underline{x} \end{Bmatrix}$ is optimal.

VI. SIS: Structural results

First, let us re-cast our problem, i.e., finding the optimal healing by a monopolist provider in a SIS epidemic, in a way amenable to the MRAP theorem. (10) can be rewritten as $\left(\beta I(1-I) - \dot{I}\right) = vI$. Let $\Omega(I) := \{\beta I(1-I) - vI \;:\; 0 \leq v \leq 1\}$, which continuously depends on I. Hence, the model in (11) can be written in the following equivalent form:

$$\max_{\dot{I}(\cdot)} \int_0^\infty e^{-rt} \left(\beta I(1-I) - \dot{I}\right) dt \tag{20a}$$

$$\text{s.t.} \quad I(0) = I_0, \quad \dot{I} \in \Omega(I(t)), \quad I \in [\underline{I}, \overline{I}] \tag{20b}$$

We can now apply the MRAP Theorem to obtain the following structural result.

Proposition 2: For $r > 0$ and $\beta + r \leq 2v_{\max}$, depending on $\beta \begin{Bmatrix} \geq \\ < \end{Bmatrix} r$, MRAP to $\begin{Bmatrix} \frac{\beta - r}{2\beta} \\ 0 \end{Bmatrix}$ is an optimal solution.

Proof: comparing (20a) with (15a), we obtain $N(x) = -1$ and $M(x) = \beta x(1-x)$. Thus, referring to (16), we have:

$$\iota(I) = -r + \beta - 2\beta I \tag{21}$$

We first consider the case of $\beta \geq r$.

$$\iota(\hat{I}) := -r + \beta - 2\beta\hat{I} = 0 \Rightarrow \hat{I} = \frac{\beta - r}{2\beta} \tag{22}$$

From (17):

- $\iota(\hat{I}) = 0$ has the unique admissible solution $\frac{\beta - r}{2\beta}$. It is admissible because (conditions in (18a)):
  - $0 \in \Omega(\hat{I})$, which is achieved by $v = \frac{\beta + r}{2}$ and since we assumed $\beta + r \leq 2v_{\max}$, this is admissible.
  - $0 \leq \hat{I} \leq 1$ (in fact, $0 \leq \hat{I} < \frac{1}{2}$).
- Furthermore, $\iota(I) \begin{Bmatrix} > \\ < \end{Bmatrix} 0$ for $\begin{cases} 0 \leq I < \hat{I} \\ \hat{I} < I \leq 1 \end{cases}$, which holds readily from (22).
- For any admissible trajectory $I(\cdot)$, with $N = -1$, the right-hand-side of condition (18c) becomes

$$\lim_{t\to\infty} e^{-rt} \int_{I(t)}^{\hat{I}} -\, d\xi = \lim_{t\to\infty} e^{-rt} \underbrace{[I(t) - \hat{I}]}_{\text{bounded}} = 0$$

  hence is trivially nonnegative.
- An MRAP for our problem exists for any $0 \leq I_0 \leq 1$, as by lemma 3.

Now consider the case in which $\beta < r$, then the last part of the MRAP Theorem applies: for $\beta < r$, $\iota(I)$ in (21) is negative for all $0 \leq I \leq 1$ (compare with (19)). Thus MRAP to 0 is optimal. ■

Following lemma 3, we have the following proposition that specifies the optimal healing policy of a monopolist in a SIS system.

Proposition 3: Write $\hat{I} := \frac{\beta - r}{2\beta}$. If $\beta + r \leq 2v_{\max}$ in the SIS system then the optimal control of the monopolist is given by

$$v(t) = \begin{cases} v_{\max} & I(t) > \hat{I}, \\ 0 & I(t) < \hat{I}, \\ \frac{\beta + r}{2} & I(t) = \hat{I}. \end{cases}$$

Proof: According to theorem 2, the optimal policy of the provider is when $I(t)$ follows an MRAP from I_0 to $\hat{I}$ until it reaches it (if at all), and subsequently maintaining $I(t)$ at $\hat{I}$. That is, if $I_0 < \hat{I}$, then while $I(t) < \hat{I}$, the control that makes the $\dot{I}$ most positive should be selected (lemma 3), so that $I(t)$ approaches $\hat{I}$ the quickest. $\dot{I} = \beta I(1-I) - vI$ is maximized when $v(t) = 0$. If on the other hand $I_0 > \hat{I}$, then again according to lemma 3 the MRAP is achieved when $\dot{I} = \beta I(1-I) - vI$ is made the smallest, which is realized for $v(t) = v_{\max}$. Now, when $\hat{I}$ is reached, then $\dot{I} = \beta\hat{I}(1-\hat{I}) - v\hat{I}$ must vanish so that the density of infection is maintained at $\hat{I}$. $\hat{I} = 0$ is never reached (lemma 2), and for $\hat{I} > 0$, the value of v which makes $\dot{I} = 0$ for $I(t) = \hat{I}$ is thus $v = \beta(1-\hat{I}) = \frac{\beta + r}{2\beta}$. ■

Remark 3: Considering the cases of $\beta + r \leq 2v_{\max}$, the theorem is essentially saying if you are more impatient than the rate of propagation ($\beta < r$), then start harvesting the market at the beginning with maximum intensity. Otherwise if $\beta > r$, reach an intermediate infection level $\hat{I} = \frac{\beta - r}{2\beta}$, which by the way is always less than $1/2$, as fast as possible by picking the extreme of the controls (depending on whether the market starts at higher than $\hat{I}$ or lower), and then maintain the level of infection at $\hat{I}$ by switching to a suitable intermediate control which zeros $\dot{I}$ and stay there forever.

Remark 4: Once again, as in the SIRD case, the optimal healing policy for the monopolist is not quite so ideal from the point of view of the network. Indeed, the best outcome for the network is when the monopolist provides healing with maximum intensity from time zero and maintained throughout, and the epidemic is pushed towards extinction. Here, for example, when β and $v_{\max}$ are large and S_0 and r are small, or in other words, when the firm is not too impatient (not a too large r), the healing is effective, and the spread of the infection is fast and the initial level is infection is low, one should expect that the provision of healing is deferred. More troubling is that the firm will then maintain an optimum level of infection at slightly less than half of the population forever.

The intuition behind theorem 3 is that the firm waits (depending on its patience) till the market ripens and then instead of throttling the market, maintains it. The instantaneous profit of the firm is maximized when the infection level is at half of the population, as it yields the maximum rate of conversion of susceptibles to infectives and vice versa.

VII. Concluding Discussion and Future Work

We derived the optimal policy of a monopolist security service provider for both cases of permanent immunity (SIRD model) and temporary immunity (SIS model). We showed that in both cases the optimal policies by the provider leads to inauspicious outcomes for the network. The bounded time interval for the SIRD can the discounted infinite horizon for the SIS model was chosen for simplicity of exposition of the results. Generalization to each other can be derived.

A next natural question is whether presence of competition between the providers can alleviate the situation and leads to provisioning of security services upon outbreak of an infection and towards its extinction with maximum intensity. Our first steps in this direction shows that although competition may lessen the severity of the social damage as compared to the monopolist case, it does not completely rectify the problem. In example models of competition, we show (appendix-A) that in the open-loop N.E. of the resulting (non-cooperative) dynamic game, the strategy of each competitor in general is of the same type as for the monopoly case: deferring the provision of healing until a threshold is reached. Currently, we are expanding different scenarios of competition and investigating open-loop and closed loop N.E. strategies of such dynamic games.

This work underscores the necessity for stringent policy making and enforcement, or alternatively availability of a viable beneficent service provider for countering epidemics.

References

[1] M. Khouzani, S. Sarkar, and E. Altman, "Dispatch then Stop: Optimal Dissemination of Security Patches in Mobile Wireless Networks," in *49th IEEE CDC*, 2010.

[2] M. Bloem, T. Alpcan, and T. Basar, "Optimal and robust epidemic response for multiple networks," *Control Engineering Practice*, vol. 17, no. 5, pp. 525–533, 2009.

[3] H. Behncke, "Optimal control of deterministic epidemics," *Optimal control applications and methods*, vol. 21, no. 6, pp. 269–285, 2000.

[4] A. Bensoussan, M. Kantarcioglu, and S. Hoe, "A game-theoretical approach for finding optimal strategies in a botnet defense model," *Decision and Game Theory for Security*, pp. 135–148, 2010.

[5] M. Lelarge and J. Bolot, "A local mean field analysis of security investments in networks," in *Proceedings of the 3rd international workshop on Economics of networked systems*, pp. 25–30, ACM, 2008.

[6] M. Lelarge, "Economics of malware: epidemic risks model, network externalities and incentives," in *Allerton 2009*, pp. 1353–1360, IEEE.

[7] G. Theodorakopoulos, J. Baras, and J. Le Boudec, "Dynamic network security deployment under partial information," in *Allerton 2008*, pp. 261–267, IEEE.

[8] W. Kermack and A. McKendrick, "A contribution to the mathematical theory of epidemics: I," *Proc. Roy. Soc. A,*, 1927.

[9] R. Anderson, "Discussion: the Kermack-McKendrick epidemic threshold theorem," *Bulletin of math. biology*, vol. 53, no. 1, pp. 1–32, 1991.

[10] C. Zou, W. Gong, and D. Towsley, "Code red worm propagation modeling and analysis," in *Proceedings of the 9th ACM conference on Computer and communications security*, pp. 138–147, ACM, 2002.

[11] W. Shengjun and C. Junhua, "Modeling the spread of worm epidemics in wireless sensor networks," in *WiCom'09*, pp. 1–4, IEEE, 2009.

[12] D. Grass, J. Caulkins, G. Feichtinger, G. Tragler, and D. Behrens, *Optimal control of nonlinear processes: with applications in drugs, corruption, and terror*. Springer Verlag, 2008.

[13] D. Kirk, *Optimal Control Theory: An Introduction*. Prentice Hall, 1970.

[14] S. Sethi, "Nearest feasible paths in optimal control problems: Theory, examples, and counterexamples," *Journal of Optimization Theory and Applications*, vol. 23, no. 4, pp. 563–579, 1977.

[15] M. Spence and D. Starrett, "Most rapid approach paths in accumulation problems," *International Economic Review*, vol. 16, no. 2, pp. 388–403, 1975.

[16] R. Hartl and G. Feichtinger, "A new sufficient condition for most rapid approach paths," *Journal of Optimization Theory and Applications*, vol. 54, no. 2, pp. 403–411, 1987.

Appendix A
An example of competition

In a SIRD setting, suppose that the dynamics of the system in presence of M competing providers of healing is governed by the following

$$\dot{S} = -\beta IS \tag{23a}$$

$$\dot{I} = \beta IS - \delta I - bI - \sum_{i=1}^{M} v_i I \tag{23b}$$

$$\dot{D} = dI \tag{23c}$$

with the action constraints

$$v_i(t) \leq v^i_{\max}, \quad \forall i = 1, \dots, M, \tag{24}$$

and each competitor (say player i) is interested in maximizing its own aggregate reward,

$$J_i = \int_0^T \rho_i v_i(t) I(t)\, dt.$$

by dynamically choosing its strategy $v_i(t)$. This represents the case where the healing takes a service time and the aggregate capacity of each of the providers is limited and patients are not instantaneously initiated receiving treatments, and instead have to be wait-listed (this is compatible with the assumption of having a continuum of nodes and limited number of providers). In this case, provision by different competitors translates to an infective as pooling of the resources and the expected waiting time of infectives proportionally reduces. The reward of the i'th player is hence:

$$J_i = \int_0^T \rho_i v_i(t) I(t)\, dt.$$

Each player maximizes his objective while considering his competitors' actions. Before we can treat this problem as a (non-cooperative) dynamic game, we need to specify the available information to each player. Here we consider *open-loop* strategies, that is the action (strategy) of each player v_i is only a function of time. That is to say, competitors choose their strategy once in the beginning at time zero, knowing that other players do the same. A generalization of this assumption, which we leave untreated for now, is when the strategy of each player can also depend on the instantaneous state (hence a closed-loop policy.)

Let $v_{-i} := (v_1, \dots, v_{i-1}, v_{i+1}, \dots, v_M)$. A Nash Equilibrium is simply defined as an M-tuple $(v_1^*, \dots, v_M^*)$ of strategies where:

$$J_i(v^*_{-i}, v^*_i) \geq J_i(v^*_{-i}, v_i), \quad \forall i \in \{1, \dots, M\}$$

Proposition 4: In any open-loop N.E. strategy $(v_1, \dots, v_M$, for each of the players we have $v_i(t) = 0$ for $0 \leq t^i_0$ and $v_i(t) = v^i_{\max}$ for $t^i_0 \leq T$.

Proof: Let player $i-$th's Hamiltonian be:

$$\mathcal{H}_i := \rho_i v_i I + (\lambda^i_I - \lambda^i_S)\beta IS - \lambda^i_I \delta I - \lambda^i_I bI - \lambda^i_I \sum_{j=1}^{M} v_j I.$$

Define $\psi_i(t) := (\rho_i - \lambda^i_I)I$. This transforms the i-th Hamiltonian to

$$\mathcal{H}_i := \psi_i v_i + (\lambda^i_I - \lambda^i_S)\beta IS - \lambda^i_I \delta I - \lambda^i_I bI - \lambda^i_I \sum_{j\in\bar{i}} v_j I,$$

where $\bar{i} := \{1 \dots i-1, i+1, \dots M\}$. Then $(v_1(t), \dots, v^M(t))$ is an open-loop N.E. if

$$\dot{\lambda}^i_S = -\frac{\partial H_i}{\partial S} = -(\lambda^i_I - \lambda^i_S)\beta I$$
$$\dot{\lambda}^i_I = -\frac{\partial H_i}{\partial I} = -\frac{\psi_i}{I} v - (\lambda^i_I - \lambda^i_S)\beta S + \lambda^i_I \delta + \lambda^i_I b + \lambda^i_I \sum_{j\in\bar{i}} v_j$$
$$\lambda^i_S(T) = 0, \quad \lambda^i_I(T) = 0.$$

and

$$v_i \in \arg\max \mathcal{H}_i, \quad \text{subject to (24)}$$

The latter leads to

$$v_i = \begin{cases} v^i_{\max} & \psi_i > 0 \\ 0 & \psi_i < 0 \end{cases}$$

Note that for all i, $\psi_i(T) = \rho_i I(T) > 0$, and thus $v_i(t) = 1$ on a time interval leading to T. Now, $\dot{\psi}_i = -\dot{\lambda}^i_I I + \frac{\psi_i}{I}\dot{I}$. Let τ be a time at which $\psi_i = 0$. Hence

$$\dot{\psi}_i|_{\tau^+} = \dot{\psi}_i|_{\tau^-} = -[-(\lambda^i_I - \lambda^i_S)\beta IS + \lambda^i_I \delta I + \lambda^i_I b + \lambda^i_I \sum_{j\in\bar{i}} v_j]$$

The system is autonomous, thus the Hamiltonian is constant, and hence

$$\dot{\psi}_i|_{\tau^+} = \dot{\psi}_i|_{\tau^-} = H_i(\tau) = H_i(T) = I(T) > 0.$$

Hence, the N.E. strategy of each of the players is of $v_i(t) = 0$ for $0 \leq t^i_0$ and $v_i(t) = v^i_{\max}$ for $t^i_0 \leq T$. ■

Remark 5: The social inefficiency of these outcomes can be arbitrarily large. By social inefficiency, we mean the ratio of the social cost if this N.E. is played over the social cost if the socially optimal decision of entering from time zero was applied. To see this, note that the aggregate social cost is representable as

$$\int_0^T [f(I(t)) + g(D(t))]\, dt$$

where $f(I)$ captures the (direct and indirect) rate of social cost due to presence of infective nodes and $g(D)$ represents the rate of social cost per dead nodes. Now, depending on the infection, f and g can be large.

Remark 6: If there is a provider which is committed to enter from time zero and stays throughout and has resources to prevent the initial increase of the density of infectives, then all of the players will do the same and we reach the socially optimal outcome. However, this requires that player not be a rational profit-maximizer.[8]

[8]In the authors' opinion, this might suggest the necessity for an effective public option.