Recurrent Infection and Externalities in Treatment

Flavio Toxvaerd*

Preliminary and Incomplete. Please do not circulate.

Abstract. This paper studies a model in which a population of individuals is exposed to an infectious disease. Individuals can choose to undergo privately costly treatment which, if successful, restores the individual's susceptibility. For extreme levels of disease prevalence, equilibrium play is uniquely determined and socially optimal. For intermediate levels of disease prevalence, multiple perfect foresight equilibrium paths coexist and can lead to different steady states. Furthermore, equilibrium play may be socially suboptimal.
JEL Classification: C73, I18.
Keywords: Economic epidemiology, susceptible-infected-susceptible models, treatment, externalities and complementarities, differential games.

1. Introduction and Literature Review

As a field of study, economic epidemiology has sought to understand the interplay between individuals' incentives, constraints and behavior (at the *micro* level) and the population-wide evolution of infectious diseases (i.e. the *macro* level). Gersovitz and Hammer (2004) offer the keen observation that

> "The economic approach to infectious disease is in its infancy, somewhat oddly because many economists have long had the intuition that epidemics and infectious diseases are quintessential manifestations of the principle of an externality, itself a central concept in economics [...]. Furthermore, epidemiology provides ready-made dynamic models of disease transmission and economics provides methods of valuing the costs and benefits of health interventions and methods of dynamic optimization to guide policy. Policy toward infections is of great importance. Yet only recently have economists begun to look at these questions in a formal way."

This quote nicely summarizes the research programme of economic epidemiology. The existing theoretical literature in the field falls into two broad, conceptually different, categories. First, there is a large and growing literature on the effects and desirability of preventive measures such as quarantines, prophylaxis, vaccines and reduced levels of promiscuity and rate of partner change, which include contributions from Sethi (1978), Geoffard and Philipson (1996, 1997), Kremer (1996), Auld (2003), Aadland and Finnoff (2007), Francis (2004), Gersovitz and Hammer (2004), Boulier et al. (2007), Brito et al. (1991), Barrett (2003), Reluga (2009) and Chen and Cottrell (2009). Second, there is a smaller literature that considers different models of treatment, including Sanders (1971), Sethi (1974), Sethi and Staats (1978) Goldman and Lightwood (1995, 2002), Rowthorn (2006) and Toxvaerd (2009a,b).

Methodologically different from the above literature are contributions that study disease evolution in explicitly network theoretic contexts (see Jackson, 2008 for a review). Last, there is an important body of work on the empirics of infectious disease, such as that on prevalence elasticity, partner choice etc. These include Ahituv et al. (1996), Philipson (2000), Auld (2006), Gibbison (2006), Dupas (2007) and Oster (2005, 2007).

*Faculty of Economics, University of Cambridge and CEPR. Address for correspondence: Faculty of Economics, University of Cambridge, Austin Robinson Building, Sidgwick Avenue, Cambridge CB3 9DD, United Kingdom. Phone: +44 (0) 1223 335259; Fax: +44 (0) 1223 335475; Email: fmot2@cam.ac.uk; Web: http://people.pwf.cam.ac.uk/fmot2/.

Especially worthy of mention are the contributions by Goldman and Lightwood (1995, 2002) and Gersovitz and Hammer (2004) who, to the best of my knowledge, were the first to explicitly consider decentralized decision making and to compare the outcomes with that chosen by a central planner. The present work seeks to contribute to this body of knowledge by taking the natural step of allowing strategically sophisticated decision makers to non-cooperatively choose exposure levels and thereby not only influence their own probability of becoming infected but also, through their interaction with other individuals, the overall evolution of the disease.

The paper is organized as follows. In Section 2, I review the classical susceptible-infected-susceptible model and introduce the economic extension. In Section 3, I treat the case of centralized decision making which serves as the benchmark for subsequent analysis. In Section 4, I analyze the case of decentralized decision making. Section 5 contains a comparison of the equilibrium outcomes of the two different setups. In the Appendix, I derive the main results of the two different settings via dynamic programming.

2. The Classical and Economic Models

The classical susceptible-infected-susceptible model is simple to describe.[1] Time is continuous and runs indefinitely. A population $\mathcal{P} = [0, 1]$ consists of a continuum of infinitely lived individuals who can at each instant t each be in one of two states, namely *susceptible* or *infected*. The set of infected individuals is denoted by $\mathcal{I}(t)$ and has measure $I(t)$, while the set of susceptible individuals is denoted by $\mathcal{S}(t)$ and has measure $S(t)$. Because the population size has been normalized to unity, these measures can be interpreted as fractions. Henceforth, $I(t)$ shall be referred to as *disease prevalence*.

At each instant, the population mixes homogeneously. This corresponds to pair-wise random matching where each individual has an equal chance of meeting any other individual, irrespective of the health status of the two matched individuals. Whereas a match between two infected individuals or two susceptible individuals does not create any new infection, a match between an infected and a susceptible individual may. The rate at which infection is transferred in such a match is denoted by $\beta > 0$. This parameter captures the infectivity of the disease. Coupled with the assumption of homogeneous mixing, this means that the rate at which susceptible individuals become infected is given by the simple expression $\beta I(t)S(t)$. This means that the rate of new infection, or *disease incidence*, is proportional to disease prevalence.[2] Note that while disease incidence is a flow value, disease prevalence is a stock value.

Infected individuals recover spontaneously at rate $\gamma \geq 0$. This means that the rate at which infected individuals become susceptible is given by $\gamma I(t)$. The dynamics of the model are thus described by the following system of differential equations:

$$\dot{S}(t) = I(t)\left[\gamma - \beta S(t)\right] \tag{1}$$

$$\dot{I}(t) = I(t)\left[\beta S(t) - \gamma\right] \tag{2}$$

$$I(t) = 1 - S(t), \quad I(0) = I_0 \tag{3}$$

Using the normalization, this system reduces to the following simple logistic growth equation:

$$\dot{I}(t) = I(t)\left[\beta(1 - I(t)) - \gamma\right], \quad I(0) = I_0 \tag{4}$$

[1] See Anderson and May (1991), Daley and Gani (2001) or Keeling and Rohani (2008) for good introductions and applications.

[2] The term $\beta I(t)S(t)$ should be thought of as the rate at which susceptible individuals have contact with other individuals, multiplied by the probability of the contact being with an infectious individual, multiplied by the probability that the infection is transmitted in such a contact. See e.g. Keeling and Rohani (2008) for a detailed derivation.

The steady states of this system are

$$I^* = 0, \quad I^* = \frac{\beta - \gamma}{\beta} \tag{5}$$

For $\beta > \gamma$, the stable steady state is endemic while for $\beta < \gamma$, the stable steady state involves eradication. In other words, if the rate at which individuals become infected surpasses the rate at which they recover, then some positive fraction of the population will always be infected. If recovery is not possible, the entire population ends up being infected. On the other hand, if individuals recover at a higher rate than the rate at which they become infected, then the disease eventually dies out. This completes the description of the classical SIS model.

2.1. Extension to an Economic Model. To turn the classical model into a fully fledged economic model, I will proceed by assigning payoffs to the different disease states and assume that time is discounted. Specifically, I will assume that individuals earn a flow payoff $\pi_{\mathcal{S}}$ per instant while susceptible and $\pi_{\mathcal{I}} < \pi_{\mathcal{S}}$ per instant while infected and that time is discounted at rate $\rho > 0$. For notational simplicity, let the quantity $\pi \equiv \pi_{\mathcal{S}} - \pi_{\mathcal{I}} > 0$ denote the health premium. Since protective behavior is disregarded in the present analysis, the rate of infection cannot be directly influenced. Instead, I assume that the rate at which agents recover (and become immune to further infection) can be influenced through costly treatment. In particular, for some treatment intensity $\tau(t) \in [0, 1]$, the rate at which the individual transitions from $\mathcal{I}(t)$ to $\mathcal{S}(t)$ is given by $\tau(t)\alpha$, where $\alpha > 0$ is interpreted as the efficiency of the treatment. The treatment costs $c > 0$ per instant.

Note that unlike the classical treatment, I assume that recovery can occur exclusively through treatment, i.e. the background rate of recovery γ is everywhere replaced by $\tau(t)\alpha$. This is done for tractability. Also note that there is no disease-induced deaths in the present analysis. Relaxing either assumption has interesting consequences that will be discussed in the Conclusion.

Before I analyze the economic model in more detail, it should be mentioned that the classical model presented here has a number of simplifying assumptions that are inherited by the economic version of the model. First, there is only one disease and one level (or severity) of infection.[3] In particular, this rules out the possibility of superinfection by different strains of the disease. Second, the moment an individual is infected coincides with the onset of symptoms such as the welfare loss brought about by infection (i.e. the *incubation period* has zero length), so no infected individual acts under the mistaken belief that he or she is susceptible. Last, once an individual becomes infected, he or she immediately becomes infectious to other individuals (i.e. the *latency period* has zero length).

3. Centralized Decision Making

In the centralized or controlled model, a social planner directly controls the aggregate treatment rate $r(t) \in [0, 1]$ in order to maximize aggregate, discounted expected welfare. The planner therefore solves the following problem:[4,5]

$$\max_{r(t)\in[0,1]} \int_0^\infty e^{-\rho t} I(t)[\pi_{\mathcal{I}} - \pi_{\mathcal{S}} - r(t)c]dt \tag{6}$$

$$s.t. \quad \dot{I}(t) = I(t)\left[\beta(1 - I(t)) - \alpha r(t)\right], \quad I(0) = I_0 \tag{7}$$

[3] Thus the sets $\mathcal{S}(t)$ and $\mathcal{I}(t)$ are disjoint and exhaust $\mathcal{P}$.

[4] A solution to the centralized problem exists, as does a solution to each agent's problem in the decentralized setting. See Seierstad and Sydsaeter (1987, Theorem 15, p. 237).

[5] The objective is a simplified version of the objective $\int_0^\infty e^{-\rho t}[I(t)\left(\pi_{\mathcal{I}} - r(t)c\right) + (1 - I(t))\pi_{\mathcal{S}}]dt$, which differs by the constant $\pi_{\mathcal{S}}$.

Letting $\lambda(t)$ denote the current-value costate variable (or multiplier), the current-value Hamiltonian for this optimal control problem is then given by[6,7,8]

$$H^C = I(t)\left[\pi_{\mathcal{I}} - \pi_{\mathcal{S}} - r(t)c\right] + \lambda(t)I(t)\left[\beta(1 - I(t)) - \alpha r(t)\right] \tag{8}$$

It should be noted that this Hamiltonian is convex in $I(t)$ for all $t \geq 0$ for $r(t)$ given. The necessary conditions for optimality (for $I(t) \neq 0$) are given by

$$c + \alpha\lambda(t) = 0 \tag{9}$$

The evolution of the multiplier is governed by the differential equation

$$\dot{\lambda}(t) = \lambda(t)\left[\rho + \alpha r(t) + \beta I(t) - \beta(1 - I(t))\right] + \left[\pi_{\mathcal{S}} - \pi_{\mathcal{I}} + r(t)c\right] \tag{10}$$

The necessary conditions for optimality means that the optimal policy is of the bang-bang type and given by

$$r(t) = 0 \quad for \quad -\alpha\lambda(t) < c \tag{11}$$

$$r(t) \in [0,1] \quad for \quad -\alpha\lambda(t) = c \tag{12}$$

$$r(t) = 1 \quad for \quad -\alpha\lambda(t) > c \tag{13}$$

In interpreting this result, it is useful to interpret the multiplier $\lambda(t)$ (which is negative). The costate variable quantifies the social cost of one additional infected individual and thus $-\alpha\lambda(t)$ is but the social value of treating one additional infected individual. With this in mind, the optimal policy then simply states that individuals should be fully treated as long as the social value of doing so outweighs the cost of treatment. If the social value is smaller than the cost, then no treatment is sought. Last, in the knife's edge case in which they coincide, any treatment intensity will do.

It is clear that in a stationary point, it must be that either of the following types of solutions obtain:

$$I(t) = 0 \tag{14}$$

$$I(t) = \frac{\beta - \alpha r(t)}{\beta} \tag{15}$$

The eradication steady state is stable if $\beta < \alpha r(t)$ while the endemic steady state is stable provided $\beta > \alpha r(t)$. If $\beta < \alpha r(t)$, the endemic solution becomes unstable but irrelevant and the relevant (and stable) solution is the one with eradication. If $r(t) = 0$, then $\beta < \alpha r(t) = 0$ cannot be satisfied. Thus $I(t) = 0$ is unstable and $I(t) = 1$ is stable. Last, if $r(t) = 1$, there are two cases to consider: (i) if $\beta < \alpha r(t) = \alpha$, then $I(t) = 0$ stable and $I(t) = (\beta - \alpha)/\beta < 0$ unstable but irrelevant; (ii) if $\beta > \alpha r(t) = \alpha$, then $I(t) = 0$ is unstable and $I(t) = (\beta - \alpha)/\beta \in (0,1)$ is stable. In conclusion, when $r(t) = 1$,

$$\max\left\{0, \frac{\beta - \alpha}{\beta}\right\} \tag{16}$$

[6] Note that throughout this treatment, standard optimal control notation is used and the dependence of the optimal policy on state variables is suppressed.

[7] Note that an admissible pair of functions $(I(t), r(t))$ is such that for all $t \geq 0$, $I(t)$ satisfies the logistic growth equation and $r(t) \in [0,1]$.

[8] Note that the control $r(t)$ must be piecewise continuous and the costate variable $\lambda(t)$ piecewise continuously differentiable.

is stable, while

$$\min\left\{0, \frac{\beta-\alpha}{\beta}\right\} \tag{17}$$

is unstable. For simplicity, assume that $\beta < \alpha$ so there is only one endemic steady state in centralized and decentralized models.[9]

Corresponding to the optimal policy, the dynamic system in $(I(t), \lambda(t))$ has two (boundary) steady states as follows:

$$I^* = 1, \quad \lambda^* = \frac{-\pi}{\rho+\beta} \tag{18}$$

$$I^* = 0, \quad \lambda^* = \frac{-(\pi+c)}{\rho+\alpha-\beta} \tag{19}$$

Note that the appropriate transversality condition is satisfied[10] and that these (boundary) steady state equilibria are both locally (and globally) saddle-path stable.[11]

The centralized problem also admits an interior solution in which prevalence is kept at an interior value by treating only a fraction of the infected population. This solution is given by

$$\hat{I}_C \equiv \frac{\alpha\pi - c(\rho-\beta)}{2\beta c} \tag{20}$$

$$\hat{r}_C \equiv \frac{c(\rho+\beta) - \alpha\pi}{2\alpha c} \tag{21}$$

It turns out that since the planner's Hamiltonian is strictly convex in $I(t)$ (for given $r(t)$), this solution is not an optimum.[12] In fact, this solution is dominated by one of the two extremal solutions. To complete the characterization, the optimal path towards steady state needs to be determined. In order to do this, note the following important result:

Lemma: The function H^C is supermodular and has increasing differences in $(r(t), I(t))$.

Proof: Since H^C is defined on the real line, supermodularity holds trivially. To show increasing differences, suppose that $-\alpha\mu(t) > c$, in which case the optimal policy is to fully treat, i.e. set $\tau(t) = 1$. It then follows that

$$\frac{d}{dt}\left[c + \alpha\mu(t)\right] = \alpha\mu(t)\left[\rho + \alpha + \beta I(t)\right] + \alpha\left[\pi + c\right] \tag{22}$$

This derivative is easily verified to be negative for $I(t) \leq \hat{I}_C$. Suppose now that $-\alpha\mu(t) < c$, so that not treating at all is optimal, i.e. such that $\tau(t) = 0$. Then

$$\frac{d}{dt}\left[c + \alpha\mu(t)\right] = \alpha\mu(t)\left[\rho + \beta I(t)\right] + \alpha\pi \tag{23}$$

This derivative is positive for $I(t) \geq \hat{I}_C$ ■

This result simply shows that if treatment is optimal at some disease prevalence, then treatment is optimal for all lower levels of infection too, which in turn are brought about by

[9] In classical model, the unique steady state is endemic if $\beta > \alpha$ and involves eradication if $\beta < \alpha$. For $\beta = \alpha$, the steady state is $I(t) = I(0) = I_0$ for all $t \geq 0$.

[10] Since this is an autonomous problem, the transversality condition is that $\lim_{t\to\infty} \lambda(t)e^{-\rho t} = 0$.

[11] The eigenvalues of the Jacobian at the steady state with $I^* = 1$ are $\rho + \beta > 0$ and $-\beta < 0$ respectively while at the steady state with $I^* = 0$ they are $\rho + \alpha - \beta > 0$ and $\beta - \alpha < 0$ respectively. For a proof of global saddle-path stability, see Appendix.

[12] Neither is it a global minimum. The Hessian is neither positive nor negative semi-definite and thus the interior solution is a saddle point. It is, however, the worst possible steady state.

the policy of full treatment. Similarly, if no treatment is optimal, then disease prevalence will increase and the policy of no treatment remains optimal in perpetuity.[13]

Last, increasing differences in $(r(t), I(t))$ can also be determined from inspecting the phase diagramme. This follows since along any optimal path,

$$\frac{d\lambda(t)}{dI(t)} = \frac{\dot{\lambda}(t)}{\dot{I}(t)} > 0 \tag{24}$$

Define the set

$$\Pi^C = \left\{ (I(t), \lambda(t)) \in [0,1] \times \mathbb{R} : \dot{I}(t)\dot{\lambda}(t) \geq 0 \right\} \tag{25}$$

State-costate pairs that do not belong to the set Π^C are strictly suboptimal and can thus be disregarded by the central planner. Paths that originate at points $(I(t), \lambda(t)) \notin \Pi^C$ are not most rapid approach paths (MRAPs) and hence cannot be optimal.

These properties of the planner's Hamiltonian can now be exploited to characterize the optimal policy more succinctly. First, introduce the following partial order on the set of policies:

Definition: Consider two policies r and r' for $t \geq 0$. Let $r \geq r'$ if $r(t) \geq r'(t)$ for all $I(t) \in [0,1]$ and $t \geq 0$.

Note that the policy belongs to the space $[0,1]^\infty$ which, with this partial (point-wise) order, is a lattice. The following useful properties of H^C can now be established:

Theorem: There is a largest and least optimal policy and each is monotone in the measure of infectives $I(t)$.

Proof: Follows directly from Topkis' Monotonicity Theorem ■

This result has important consequences. It implies that along an optimal path, there can be no switch between the regimes where $r(t) = 0$ and $r(t) = 1$ respectively. In other words, optimal policies are of the most rapid approach type in which steady states are reached as fast as possible. Furthermore, because extremal policies must be monotone in disease prevalence $I(t)$, the optimal policy can be described easily in terms of the state variable as follows:

Corollary: There exists some $k \in [0,1]$ such that the optimal policy is given by

$$r(t) = 0 \quad for \quad I(t) > k \tag{26}$$

$$r(t) \in [0,1] \quad for \quad I(t) = k \tag{27}$$

$$r(t) = 1 \quad for \quad I(t) < k \tag{28}$$

Since both steady states are locally stable, monotonicity of the optimal policy in $I(t)$ implies that there exists one (and only one) value of disease prevalence $k \in [0,1]$ with this property. In fact the critical level is given by $k = \hat{I}_C$. Thus for $I_0 > \hat{I}_C$, the optimal path leads to the endemic steady state $I^* = 1$ while for $I_0 < \hat{I}_C$, the optimal path leads to the eradication steady state $I^* = 0$.

The comparative statics of $\hat{I}_C$ with respect to the different parameters are collected in the following corollary:

[13] There is an alternative route to this result. Because this is an infinite horizon autonomous problem, the state variable $I(t)$ is monotone in time and the multiplier is a one-to-one function of $I(t)$, i.e. $\lambda(t) = \lambda(I(t))$. See Kamien and Schwartz (1995, p. 179), Rowthorn (2006) or Hartl (1987) for details. Since these two extremal steady states of the system are saddle-path stable, there is for each steady state is a unique optimal path towards it (in a neighborhood of it). Coupled with the bang-bang property of the optimal policy, monotonicity and continuity of the state variable imply the result.

Note that since $\lambda(t)$ is monotone in $I(t)$ on the set Π^C, monotonicity of the optimal policy in $I(t)$ follows immediately from the form of the bang-bang policy above.

Corollary: The critical level $\hat{I}_C$ is (i) increasing in α and π and decreasing in c and ρ and (ii) decreasing in β for $c < \alpha\pi/\rho$.

Note that the condition in part (ii) will be assumed in what follows in order to make corner solutions feasible in the decentralized setting.

The two regimes have been delineated in terms of constraints on the multipliers. These can be rewritten in terms of the basic parameters of the model. Specifically, in order for the extremal solutions to be steady states, it must be that

$$c \in \left[\frac{\alpha\pi}{\rho+\beta}, \frac{\alpha\pi}{\rho-\beta}\right] \tag{29}$$

For completeness, the path towards steady state is given by

$$I(t) = \frac{I_0\xi(t)}{e^{t\xi(t)}\xi(t) - \left(1-e^{t\xi(t)}\right)\beta I_0}, \quad \xi(t) \equiv \alpha r(t) - \beta \tag{30}$$

where $r(t)$ is appropriately chosen. Note that for $\xi(t) = -\beta$, the model reduces to that of a *simple epidemic*.

Last, the speed of convergence to the steady states is of interest. For a given steady state (I^*, r^*), the speed of convergence is given by

$$\sigma(I^*, r^*) \equiv -\left[\beta(1-2I^*) - \alpha r^*\right] \tag{31}$$

For the endemic steady state, the speed of convergence is

$$\sigma(1,0) = \beta > 0 \tag{32}$$

Note that this is independent of the treatment efficiency α since around this steady state, there is no treatment. For the eradication steady state, the speed of convergence is

$$\sigma(0,1) = \alpha - \beta > 0 \tag{33}$$

4. Decentralized Decision Making

In this section, I analyze the setting in which decision making if fully decentralized so that each infected individual $i \in [0,1]$, at each instant $t \geq 0$, decides non-cooperatively on its treatment intensity $\tau_i(t) \in [0,1]$. I start by considering the best response of an individual under the assumption that all other individuals' actions are held fixed.

For some fixed level of treatment τ and initial infection state Q_0, the individual's health state evolves according to a two-state continuous-time Markov process with the following transition rate matrix:

$$\begin{pmatrix} -\beta I(t) & \beta I(t) \\ \alpha\tau & -\alpha\tau \end{pmatrix}$$

Consider an individual agent's best response policy and assume that all agents have perfect foresight. Fix some path $I(t)$ and recall that the agent cannot influence aggregate variables. Let $Q(t)$ denote the probability that the agent is infected at time $t \geq 0$. The agent then solves the following problem:[14]

$$\max_{\tau_i(t)\in[0,1]} \int_0^\infty e^{-\rho t} Q(t)[\pi_{\mathcal{I}} - \pi_{\mathcal{S}} - \tau_i(t)c]dt \tag{34}$$

$$s.t. \quad \dot{Q}(t) = (1-Q(t))\,\beta I(t) - Q(t)\alpha\tau_i(t), \quad Q(0) = I_0 \tag{35}$$

[14] The objective is a simplified version of the objective $\int_0^\infty e^{-\rho t}[Q(t)\,(\pi_{\mathcal{I}} - \tau_i(t)c) + (1-Q(t))\pi_{\mathcal{S}}]dt$, which differs by the constant $\pi_{\mathcal{S}}$.

The individuals' problems are of course solved subject to the condition that the aggregate evolution of disease prevalence is consistent with individual decisions, i.e. subject to the constraint

$$\dot{I}(t) = I(t)\left[\beta(1 - I(t)) - \alpha \mathbf{r}(t)\right], \quad I(0) = I_0 \tag{36}$$

were the fraction seeking treatment at time $t \geq 0$ is given by

$$\mathbf{r}(t) \equiv \int_{i \in \mathcal{I}(t)} I(t)^{-1} \tau_i(t) di \tag{37}$$

The associated current value Hamiltonian for this problem is then[15]

$$H^D(\tau_i(t), \tau_{-i}(t)) = Q(t)[\pi_{\mathcal{I}} - \pi_{\mathcal{S}} - \tau_i(t)c] + \mu(t)\left[(1 - Q(t))\,\beta I(t) - Q(t)\alpha\tau_i(t)\right] \tag{38}$$

where $\mu(t)$ is the current value costate variable. In this definition, $H^D(\tau_i(t), \tau_{-i}(t))$ is the current value Hamiltonian at time $t \geq 0$ for agent $i \in [0, 1]$ setting treatment at $\tau_i(t)$, given that other players play according to $\tau_{-i}(t)$.

The current value Hamiltonian has the following useful interpretation. It is the current value of flow utility from treatment intensity $\tau_i(t)$ given state $Q(t)$. Increasing $\tau_i(t)$ reduces "current" utility since higher costs are incurred, but also adds value through the resulting reduction in the probability of subsequently being infected (an instant later). The shadow price $\mu(t)$ simply calculates this "future" value in current utility terms.

It is important to note that other agents' treatment decisions influences an individual's tradeoff through their impact on the evolution of $I(t)$, with higher $I(t)$ effectively eroding the value of treatment by increasing the rate at which reinfection occurs.

Next, I characterize the best response of an individual, noting that $\tau_i(t)$ must be piecewise continuous and while the costate variable $\mu(t)$ must be piecewise continuously differentiable.[16]

The necessary (and sufficient[17]) conditions for optimality (for $Q(t) \neq 0$) are given by

$$c + \alpha\mu(t) = 0 \tag{39}$$

This yields the following optimal bang-bang policy:[18]

$$\tau_i(t) = 0 \quad for \quad -\alpha\mu(t) < c \tag{40}$$

$$\tau_i(t) \in [0, 1] \quad for \quad -\alpha\mu(t) = c \tag{41}$$

$$\tau_i(t) = 1 \quad for \quad -\alpha\mu(t) > c \tag{42}$$

Wile this policy resembles that of the planer under centralized decision making, the costate variable multiplies a difference constraint in the individual's problem.

Next, the evolution of the multiplier is governed by the following differential equation

$$\dot{\mu}(t) = \mu(t)\left[\rho + \alpha\tau_i(t) + \beta I(t)\right] + \left[\pi_{\mathcal{S}} - \pi_{\mathcal{I}} + \tau_i(t)c\right] \tag{43}$$

From this law of motion, the dependence of the individual's problem on the evolution of disease prevalence is apparent. While the individual's "personal" state variable $Q(t)$ does not appear in this equation, it should be recalled that in equilibrium, it must be the case that $Q(t) = I(t)$ and the evolution of $Q(t)$ therefore mirrors the evolution of $I(t)$.

[15] In the agent's problem, an admissible pair of functions $(Q(t), \tau_i(t))$ is such that for all $t \geq 0$, $Q(t)$ satisfies the differential equation for the state variable $Q(t)$ and $\tau_i(t) \in [0, 1]$.

[16] Note that $-\alpha\mu(t)$ is to be interpreted as the expected, discounted benefit of treatment at time t under the assumptions that $I(t)$ is a constant for all $s \geq t$ and that the the individual will play optimally at all future dates $s \geq t$. Even if the actions at some point in time do not directly persist over time, the consequences (switches between disease states) do.

[17] Mangasarian's sufficiency conditions are easily verified for the agent's problem.

[18] Note that this policy is Markovian and stationary for $\mu(t)$ constant, which will be the case when $I(t)$ is constant. In other words, history dependence of an agent's optimal policy can only occur as a response to history dependence in other agents' policies.

4.1. Strategies and Equilibrium. Before presenting the characterization of the equilibrium set, I formally define the notions of strategy and equilibrium that will be studied.

Definition: A time t behavior strategy for player $i \in [0,1]$ is given by $\tau_i(t) \in [0,1]$.

Definition: A time t strategy for player $i \in [0,1]$ is given by $\overline{\tau}_i(t) = (\tau_i(s))_{s \geq t} \in [0,1]^{\infty}$.

Definition: A time t joint behavior strategy is given by $\boldsymbol{\tau}(t) = (\tau_i(t))_{i \in [0,1]} \in [0,1]$.

Definition: A time t joint strategy is given by $\overline{\boldsymbol{\tau}}(t) = (\overline{\tau}_i(t))_{i \in [0,1]} = (\boldsymbol{\tau}(s))_{s \geq t} \in [0,1]^{\infty}$.

Denote the set of possible time t histories by $\mathcal{H}_t$, with typical element $h_t \equiv \{I_0, \boldsymbol{\tau}(s)\}_{s \in [0t[}\}$. In all generality, a time t behavior strategy is a mapping $\tau_i : \mathcal{H}_t \to [0,1]$. It should be noted that these strategies are *conditional* on being infected, i.e. they prescribe the probability of treatment being sought by the agent *if he happens to be infected*. This convention obviates the need to specify state-dependent behavior strategies in which the strategy set for a susceptible individual is empty. Furthermore, these strategies, conditioning actions at each instant on past history, are so-called feedback (or closed loop) strategies. Accordingly, the equilibrium notion employed below is sequentially rational.

Next, in order to employ the tools of lattice programming to the game, I introduce a straightforward order on the strategy spaces of the individuals. Consider two strategies $\overline{\tau}_i(t)$ and $\overline{\tau}'_i(t)$.

Definition: Let $\overline{\tau}_i(t) \geq \overline{\tau}'_i(t)$ if $\tau_i(t) \geq \tau'_i(t)$ for all $I(t) \in [0,1]$ and $t \geq 0$.

Definition: Let $\overline{\boldsymbol{\tau}}(t) \geq \overline{\boldsymbol{\tau}}'(t)$ if $\overline{\tau}_i(t) \geq \overline{\tau}'_i(t)$ for all $i \in [0,1]$.

It is important to emphasize that with this partial order, both the individual strategy sets and the joint strategy set is a lattice.

Note that for $\overline{\boldsymbol{\tau}}(t) \geq \overline{\boldsymbol{\tau}}'(t)$, the path $(I(s))_{s \geq t}$ shifts downwards (weakly) for all $s \geq t$, i.e.

$$\left(I(s)'\right)_{s \geq t} \geq (I(s))_{s \geq t} \tag{44}$$

where $(I(s)')_{s \geq t}$ is the path corresponding to joint strategy $\overline{\boldsymbol{\tau}}'(t)$. This follows trivially since $\mathbf{r}(t) \geq \mathbf{r}(t)'$.

An equilibrium from $t = 0$ onwards is a strategy profile $\overline{\boldsymbol{\tau}}^*(t) = (\overline{\tau}^*_i(t))_{i \in [0,1]}$ such that no player $i \in [0,1]$ can benefit from choosing a strategy $\overline{\tau}_i(t) \neq \overline{\tau}^*_i(t)$ when all other players $j \in [0,1]$, with $j \neq i$ play $\overline{\tau}^*_{-i}(t)$. Formally, an equilibrium is defined as follows:

Definition: An equilibrium is a strategy profile $\overline{\boldsymbol{\tau}}^*(t) = (\overline{\tau}^*_i(t))_{i \in [0,1]}$ such that for all $i \in [0,1]$, all $\overline{\tau}_i(t)$ and all $t \geq 0$:

$$H^D(\tau^*_i(t), \tau^*_{-i}(t)) \geq H^D(\tau_i(t), \tau^*_{-i}(t)) \tag{45}$$

To characterize the equilibrium set, it shall prove useful to begin by characterizing an interior (mixed strategy) equilibrium. Using the interiority condition $-\alpha\mu(t) = c$ with the steady state equation $\dot{\mu}(t) = 0$ yields the following critical disease prevalence:

$$\hat{I}_D \equiv \frac{\alpha\pi - \rho c}{\beta c} \tag{46}$$

This is the level of infection that makes the agents willing to use interior behavioral strategies.[19] Using this value of disease prevalence in the steady state equation $\dot{I}(t) = 0$ (with $Q(t) = I(t)$) then yields the level of aggregate treatment necessary for steady state:

$$\widehat{\mathbf{r}}_D \equiv \frac{c(\rho + \beta) - \alpha\pi}{\alpha c} \tag{47}$$

[19] Note that $\hat{I}_D > 0$ for $c < \frac{\alpha\Delta\pi}{\rho}$ and $\hat{I}_D < 1$ for $c < \frac{\alpha\Delta\pi}{\rho+\beta}$. Both these inequalities have been assumed to hold.

At this is the level of aggregate treatment, infection remains at exactly the (constant) level required to make agents willing to employ interior behavioral strategies. It is interesting that for $c > \alpha\pi/(\rho+\beta)$, which is the condition ensuring that the endemic steady state exists in either setting, it is the case that $\widehat{I}_C > \widehat{I}_D$ and correspondingly, $\widehat{\mathbf{r}}_D > \widehat{r}_C$.

Next, note that except when $-\alpha\mu(t) = c$, the agents have corner best responses, i.e. when $-\alpha\mu(t) = c$, any $\tau_i(t) \in [0,1]$ yields the same payoff and hence the agents are willing to mix over their pure behavioral strategies. In any mixed strategy equilibrium profile $\boldsymbol{\tau}^*(t)$, it must therefore be the case that

$$\int_{i\in\mathcal{I}(t)} \tau_i(t)di = \left(\frac{c(\rho+\beta)-\alpha\pi}{\alpha c}\right)\left(\frac{\alpha\pi-\rho c}{\beta c}\right) \tag{48}$$

But this condition is necessary, not sufficient for $\boldsymbol{\tau}^*(t)$ to lead to a steady state (since it must hold identically for all $t \geq 0$). In order to satisfy the steady state condition for $I(t)$, the distribution of strategies across the set of infected individuals must be constant over time. But this can only be achieved if strategies are symmetric. Thus it must be the case that for all $i \in [0,1]$,

$$\tau_i(t) = \frac{c(\rho+\beta)-\alpha\pi}{\alpha c} \tag{49}$$

It should be mentioned that there are pure strategy equilibria that achieve the same disease prevalence $\hat{I}_D$ as in the mixed strategy equilibrium, but these have very unappealing properties.[20]

Note that for all other equilibria than the one maintaining disease prevalence $\hat{I}_D$, it is the case for all $t \geq 0$ that $\mathbf{r}(t) \in \{0,1\}$. That is, either all susceptibles seek full treatment or they all seek no treatment at all.

In the individual agent's problem, the Hamiltonian is linear in both state and control variable and thus the aggregate interior steady state cannot be discarded as is the case in the centralized model. Having said that, it should also be emphasized that the mixed strategy equilibrium may be very (socially) undesirable. This is because the centralized Hamiltonian is strictly convex and therefore *any* interior solution will be strictly worse than *one* of the boundary solutions.

I will proceed by characterizing the mixed strategy equilibrium. The comparative statics of $\hat{I}_D$ with respect to the different parameters are collected in the following corollary:

Corollary: The critical level $\hat{I}_D$ is (i) increasing in α and π and decreasing in c and ρ and (ii) decreasing in β for $c < \alpha\pi/\rho$.

The next step of the analysis is to show that the game played between the individuals under decentralized decision making is a supermodular game. I start by showing the following important result:

Lemma: The function H^D is supermodular and has increasing differences in $(r(t), I(t))$.

Proof: Since the function H^D is defined on the real line, supermodularity holds trivially. Suppose that $-\alpha\mu(t) > c$, which makes $\tau(t) = 1$ optimal. It then follows that

$$\frac{d}{dt}\left[c+\alpha\mu(t)\right] = a\mu(t)\left[\rho+\alpha+\beta I(t)\right] + \alpha\left[\pi+c\right] \tag{50}$$

[20] Apart from being asymmetric, the identity of the agents playing each pure strategy changes at each instant in such equilibria. This is because the composition of the sets $\mathcal{S}(t)$ and $\mathcal{I}(t)$ evolves over time. Thus even if all players adopt stationary pure strategies, the aggregate environment is no longer stationary, which upsets the interior pure strategy equilibria.

This expression is negative for $I(t) \leq \hat{I}_D$. Next, suppose that $-\alpha\mu(t) < c$, which makes $\tau(t) = 0$ optimal. In this case, it follows that

$$\frac{d}{dt}\left[c + \alpha\mu(t)\right] = a\mu(t)\left[\rho + \beta I(t)\right] + \alpha\pi \tag{51}$$

This expression is positive for $I(t) \geq \hat{I}_D$. This completes the proof ■

An alternative route to establish that the individual's objective function under decentralized decision making has increasing differences in $(I(t), \tau_i(t))$ and $(Q(t), \tau_i(t))$ is to start by noting that

$$\frac{d\mu(t)}{dI(t)} = \frac{\dot{\mu}(t)}{\dot{I}(t)}, \quad \frac{d\mu(t)}{dQ(t)} = \frac{\dot{\mu}(t)}{\dot{Q}(t)} \tag{52}$$

Next, following the definition under centralized decision making, define the set

$$\Pi^D = \left\{(Q(t), \mu(t)) \in [0,1] \times \mathbb{R} : \dot{Q}(t)\dot{\mu}(t) \geq 0\right\} \tag{53}$$

Clearly, state-costate pairs that do not belong to the set Π^D are strictly dominated strategies and can thus be disregarded.[21, 22] Next, observe that the influence of other players' strategies on a given player's payoffs are captured entirely by the value of $\mu_i(t)$ and one can therefore write $\mu_i(\boldsymbol{\tau}(t))$. It follows that

$$\frac{d\mu_i(t)}{d\boldsymbol{\tau}(t)} = \frac{d\mu_i(t)}{dI(t)}\frac{dI(t)}{d\boldsymbol{\tau}(t)} \leq 0 \tag{54}$$

In other words, higher strategies (more treatment) causes lower infection, which in turn causes higher marginal disutility of infection. In turn, this means that as $\boldsymbol{\tau}(t)$ increases, the marginal utility of treatment $-\alpha\mu(t)$ increases, thereby increasing the indifference point (the switching threshold) between non-treatment and full treatment. Thus there are strategic complementarities. This property allows me to use of the methodology of supermodular games to deduce the following important results:

Theorem: (i) optimal policies (best responses) exist for each agent $i \in \mathcal{P}$, (ii) each agent's strategy space is a lattice, (iii) each agent's objective function is supermodular in own action $\tau_i(t)$ and has increasing differences in $(\tau_i(t), \boldsymbol{\tau}(t))$ and $(\tau_i(t), I(t))$.

This result simply states that the game played under decentralized decision making fulfills all the formal requirements of a supermodular game. From Topkis' Monotonicity Theorem, it then follows that the usual properties of supermodular games obtain. Namely, I have the following:

Theorem: (i) there is a largest and least policy (best response), (ii) each best response is monotone in the measure of infectives $I(t)$, (iii) there exists a largest and a least equilibrium (iv) extremal equilibria are monotone in disease prevalence $I(t)$.

It is clear that for some constant path of $I(t)$, there can be no switch between the regimes where $\tau_i(t) = 0$ and $\tau_i(t) = 1$ respectively.

To find the aggregate steady states of the game, I start by considering the stationary point of an individual's problem. From the perspective of an individual agent, steady state obtains

[21] No agent will assign prositive probability to a strategy that is strictly dominated.

[22] Note that since $\mu(t)$ is monotone in $I(t)$ on the set Π^D, monotonicity of the best response in $I(t)$ follows immediately from the form of the bang-bang policy for the individual agent's problem.

when $\dot{Q}(t) = \dot{\mu}(t) = 0$, i.e. when

$$Q(t) = \frac{\beta I(t)}{\beta I(t) + \alpha \tau_i(t)} \tag{55}$$

$$\mu(t) = \frac{-(\pi + \tau_i(t)c)}{\rho + \tau_i(t)\alpha + \beta I(t)} \tag{56}$$

From these equations, it is clear that the individual's problem is in steady state only when the aggregate system is in steady state, i.e. when $\dot{I}(t) = 0$ (i.e. is in either of the boundary solutions or in some interior solution). Note that this dynamic system in $(Q(t), \mu(t))$ is saddle-path stable for any combination $(\tau_i(t), I(t))$.[23] Last, note that the appropriate transversality condition is satisfied.[24]

It turns out that the two boundary steady states of the setting with centralized decision making are also candidate steady states of the decentralized setting. In the endemic steady state, it must be that

$$Q(t) = 1, \quad \mu(t) = \frac{-\pi}{\rho + \beta} \tag{57}$$

In the eradication steady state, it must be that

$$Q(t) = 0, \quad \mu(t) = \frac{-(\pi + c)}{\rho + \alpha} \tag{58}$$

The constraints on the multipliers characterizing the two boundary solutions can be rewritten in terms of the parameters of the model as follows:

$$c \in \left[\frac{\alpha\pi}{\rho + \beta}, \frac{\alpha\pi}{\rho} \right] \tag{59}$$

Since the largest and smallest equilibrium strategy is monotone in disease prevalence $I(t)$, these can be usefully characterized as follows. Define the following threshold values of disease prevalence:

$$k' = \left\{ \max I(t) \in [0,1] : -\alpha\mu(t) \geq c \quad \text{with } (I(s))_{s \geq t} \,|_{\mathbf{r}(t)=0} \right\} \tag{60}$$

$$k'' = \left\{ \min I(t) \in [0,1] : -\alpha\mu(t) \leq c \quad \text{with } (I(s))_{s \geq t} \,|_{\mathbf{r}(t)=1} \right\} \tag{61}$$

Then clearly these thresholds are ranked and $k' \leq k''$. The smallest possible equilibrium strategy (given the partial order I have imposed on the strategy sets) is given by

$$\tau_i(t) = \begin{cases} 0 & \text{for all } I(t) \geq k' \\ 1 & \text{for all } I(t) < k' \end{cases} \tag{62}$$

The largest possible equilibrium strategy is similarly given by

$$\tau_i(t) = \begin{cases} 0 & \text{for all } I(t) \geq k'' \\ 1 & \text{for all } I(t) < k'' \end{cases} \tag{63}$$

The following result follows from the monotonicity of extremal equilibria:

Corollary: The extremal equilibria are Markovian and stationary.

[23] The relevant eigenvalues of the Jacobian are $\rho + \beta I(t) + \alpha\tau_i(t) > 0$ and $-\beta I(t) - \alpha\tau_i(t) < 0$ respectively.

[24] Since this is an autonomous problem, the transversality condition is that $\lim_{t\to\infty} \mu(t)e^{-\rho t} = 0$.

In games with strategic complementarities, there are typically multiple equilibria. In the context of the present game, multiple (monotone strategy) equilibria corresponds to the existence of multiple cutoff points $k \in [0,1]$ that constitute equilibrium strategies. More precisely, for any level of disease prevalence $I(t) \in [k', k'']$, there exist different paths of disease prevalence such that for one path $(I(s))_{s\geq t}$, the inequality $-\alpha\mu(t) + c \geq 0$ holds while for the other path $(I(s)')_{s\geq t}$, the inequality $-\alpha\mu(t) + c \leq 0$ holds. These equilibria are Pareto ranked. The following result shows this multiplicity formally:

Theorem: The game played under decentralized decision making has multiple equilibria.

Proof: To show that the equilibrium set is not a singleton, note that if $I(t)$ equals critical interior level $\hat{I}_D$, then each agent is indifferent and willing to mix (expecting constant $I(t) = \hat{I}_D$ for all t). But if they are willing to mix at $I(t) = \hat{I}_D$, then they will be willing to play (and continue to play) either of the extreme actions as long as all players do so (and continue to do so). This means that starting from $I(t) = \hat{I}_D$, there are two possible equilibrium paths.[25] But starting at $I(t) = \hat{I}_D$, if either of these equilibrium paths are expected to be followed by almost all individuals, then no individual agent will be willing to mix because he will strictly prefer one of the two extreme actions. This is because if $\tau_i(t) = 1$ is optimal for a constant path $(I(s))_{s\geq t}$ with $I(t) = I$ for all t, then it is also optimal for any decreasing path $(I(s))_{s\geq t}$ with $I(t) \leq I$ for all t. Similarly, if $\tau_i(t) = 0$ is optimal for a constant path $(I(s))_{s\geq t}$ with $I(t) = I$ for all t, then it is also optimal for any increasing path $(I(s))_{s\geq t}$ with $I(t) \geq I$ for all t. But then the indifference point must be shifted (up or down, depending on continuation play) and thus there are multiple equilibria in the sense that for a non-singleton set of initial conditions $I(0) = I_0$, there are multiple (distinct) equilibrium paths ■

Before I proceed to compare the decentralized and centralized outcomes, I present some straightforward comparative statics:

Proposition: Ceteris paribus, each agent's expected discounted utility is increasing in α and decreasing in β, c and π.

Proof: Note that the individual's Hamiltonian is, for all $t \geq 0$, jointly concave in $(Q(t), \tau_i(t))$ and in $(Q(t), x)$ for $x = \alpha, \beta, c, \pi_{\mathcal{I}}, \pi_{\mathcal{S}}$.[26] Given these properties, the comparative statics results follow by recalling that the properties of the value function are then inherited by the Hamiltonian[27] ■

5. Comparing the Centralized and Decentralized Setups

To compare the two settings, it's useful to compare the centralized problem's phase diagramme with that characterizing the aggregate outcome under decentralization. In the centralized setting, the construction is based on a straightforward application of the steady state conditions for the state and the costate variables. In constructing the aggregate phase diagramme for the decentralized setting, it is instructive to consider the individual agent's steady state condition. For the individual's problem to be at a rest point, it must be that

$$Q(t) = \frac{\beta I(t)}{\beta I(t) + \alpha\tau_i(t)} \tag{64}$$

For $-\alpha\mu(t) < c$, it is optimal not to seek any treatment and consequently the only steady state value is $Q(t) = 1$. In other words, when an agent doesn't seek any treatment, the probability of this agent being infected must tend to one. In turn, when $-\alpha\mu(t) \geq c$, full treatment

[25] This is equivalent to two distinct equilibrium outcomes in a simple 2x2 coordination game.

[26] The Hamiltonian is linear and thus concave $I(t)$ (for given $\tau_i(t)$) and jointly concave in $(I(t), x)$ for $x = \alpha, \beta, c, \pi_{\mathcal{I}}, \pi_{\mathcal{S}}$.

[27] See Seierstad and Sydsaeter (1987, p. 217) for details.

is optimal and hence the steady state becomes

$$Q(t) = \frac{\beta I(t)}{\beta I(t) + \alpha} \tag{65}$$

This expression depends explicitly on $I(t)$, which shows that whenever disease prevalence in constant, the individual agent's problem has a steady state corresponding to that level of prevalence. This means that by varying disease prevalence $I(t)$, one can trace out the points in $(I(t), -\alpha\mu(t))$ space at which the individual's problem is in steady state (for constant $I(t)$). This allows one to directly compare the centralized and the decentralized settings in the same phase diagramme.

Note that the scenario $-\alpha\mu(t) \geq c$ includes the special cases of full eradication $Q(t) = I(t) = 0$ and the mixed strategy equilibrium $Q(t) = I(t) = \hat{I}_D(t)$. This means that for $-\alpha\mu(t) \geq c$, all points on the optimal trajectory represent potential steady states. On the other hand, for $-\alpha\mu(t) < c$, the only point on the optimal trajectory that is a potential steady state corresponds to that of total infection, i.e. where $Q(t) = 1$.

When tracing out the curves $\dot{Q}(t) = 0$ and $\dot{I}(t) = 0$, they have the same overall appearance.[28] But these curves only bound the optimal trajectories, they do not directly represent them. It would therefore be premature to conclude that the optimal paths are therefore qualitatively similar under the two different settings. It turns out that there are important differences as has already been indicated by the discussion of the viability of the interior solutions.

The main difference is that the optimal policy in the centralized setting has a discontinuity (for an interior level of $I(t)$) corresponding to the switch between the optimal treat all and the treat no-one policies, namely at the critical level of disease prevalence $\widehat{I}_C$. By contrast, when tracing out the aggregate path under decentralized decision making, the aggregate path does include the corresponding point $\widehat{I}_D$. In other words, in the decentralized setting, the optimal policy is continuous in disease prevalence $I(t)$ (for some fixed profile of strategies).

While the planner's problem is convex in $I(t)$, the agents' problems are linear (and thus concave) in $Q(t)$ (and in $I(t)$). This means that in the decentralized problem, an interior (and thus a mixed strategy) equilibrium exists, whereas in the centralized problem no interior optimal solution exists.

Next, there are differences in the conditions that ensure existence of certain steady states. While the endemic (full infection) steady state is present under both models, the eradication steady state may fail to exist under decentralized decision making even when present in centralized model. This is because the disutility attached to infection by the individual and a social planner differ, a consequence of the externalities inherent in treatment activities. To appreciate this, observe that controlling for a given optimal path $(I(s))_{s \geq t}$, it is the case that

$$-\alpha\lambda(t) > -\alpha\mu(t) \quad for \quad I(t) \in [0, 1) \tag{66}$$

$$-\alpha\lambda(t) = -\alpha\mu(t) \quad for \quad I(t) = 1 \tag{67}$$

These expressions show that, for a fixed path $(I(s))_{s \geq t}$, the private value of treatment is lower than the social value of treatment, unless the system is in the endemic steady state (at which there are no susceptible individuals who can benefit from the positive externalities of treatment). This implies that even the most efficient decentralized equilibrium is smaller than the centralized equilibrium.

The sufficient conditions for the boundary steady states to exist are as follows: The steady state $I(t) = 1$ requires $c > c_1^C \equiv \alpha\pi/(\rho + \beta)$ in the centralized setting and $c > c_1^D \equiv \alpha\pi/(\rho + \beta)$ in the decentralized setting. In turn, the steady state $I(t) = 0$ requires $c < c_0^C \equiv \alpha\pi/(\rho - \beta)$ in

[28] This was also noted by Goldman and Lightwood (2002).

the centralized setting and $c < c_0^D \equiv \alpha\pi/\rho$ in the decentralized setting. It is easily verified that $c_1^C = c_1^D < c_0^D < c_0^C$ under the imposed parameter restrictions. In summary, both steady states exist in both settings if

$$c \in \left[\frac{\alpha\pi}{\rho+\beta}, \frac{\alpha\pi}{\rho-\beta}\right] \tag{68}$$

Under centralized decision making, both steady states exist if

$$c \in \left[\frac{\alpha\pi}{\rho}, \frac{\alpha\pi}{\rho-\beta}\right] \tag{69}$$

But in this range, only the full infection steady state exists under decentralized decision making.

A. Global Saddle-Path Stability

To show global saddle-path stability of the steady states, start by defining the functions

$$\begin{aligned} F(\lambda, I) &\equiv \lambda(t)\left[\rho - \beta(1-I(t)) + \alpha r(t) + \beta I(t)\right] \\ &\quad +(\pi + r(t)c) \qquad (70) \\ G(\lambda, I) &\equiv I(t)\left[\beta(1-I(t)) - \alpha r(t)\right] \qquad (71) \end{aligned}$$

These functions are continuous and C^1 in a neighborhood of either non-interior steady state (λ^*, I^*). Straightforward evaluation shows that

$$F_2'(\lambda^*, I^*)G_1'(\lambda^*, I^*) = 0, \quad F_1'(\lambda^*, I^*)G_2'(\lambda^*, I^*) < 0 \tag{72}$$

Also, the equation $F(\lambda, I) = 0$ has a unique solution $I = I(\lambda)$.

Next, consider the stability of the steady state with $r(t) = 0$ and $I(t) = 1$. It is straightforward to see that

$$\left|\frac{G(\lambda, I)}{F(\lambda, I)}\right| = \left|\frac{I(t)\beta(1-I(t))}{\lambda(t)\left[\rho - \beta + 2\beta I(t)\right] + \pi}\right| \tag{73}$$

is bounded for any relevant pair (λ, I). Next, consider the stability of the steady state with $r(t) = 1$ and $I(t) = 0$. Again, it is straightforward to see that

$$\left|\frac{G(\lambda, I)}{F(\lambda, I)}\right| = \left|\frac{I(t)\left[\beta(1-I(t)) - \alpha\right]}{\lambda(t)\left[\rho - \beta + \alpha + 2\beta I(t)\right] + c + \pi}\right| \tag{74}$$

is bounded for any relevant pair (λ, I). These results imply that the paths are globally saddle-path stable. See Seierstad and Sydsaeter (1987, Theorem 19, p. 256) for details.

B. Values Along Equilibrium Paths

In order to analyze payoffs in each disease state along the equilibrium paths, I derive the optimality conditions via dynamic programming. Consider an individual seeking treatment at rate $\tau_i(t) \in [0, 1]$. In an interval of time $[t, t + \Delta t]$, the rate at which successful treatment is achieved for this individual is given by $\alpha\tau_i(t)\Delta t$. The discrete-time approximation of the

equations describing the value functions for the two disease states are then given by

$$\begin{aligned} V_{\mathcal{S}}(I(t)) &= \pi_{\mathcal{S}}\Delta t + \frac{\beta I(t)\Delta t V_{\mathcal{I}}(I(t+\Delta t))}{1+\rho\Delta t} \\ &\quad + \frac{(1-\beta I(t)\Delta t)V_{\mathcal{S}}(I(t+\Delta t))}{1+\rho\Delta t} \end{aligned} \tag{75}$$

$$\begin{aligned} V_{\mathcal{I}}(I(t)) &= \max_{\tau_i(t)\in[0,1]} \{(\pi_{\mathcal{I}} - \tau_i(t)c)\,\Delta t \\ &\quad + \frac{\tau_i(t)\alpha\Delta t\,(V_{\mathcal{S}}(I(t+\Delta t)))}{1+\rho\Delta t} \\ &\quad + \frac{(1-\tau_i(t)\alpha\Delta t)V_{\mathcal{I}}(I(t+\Delta t))}{1+\rho\Delta t}\} \end{aligned} \tag{76}$$

Rearranging, dividing by Δt and taking the limit $\Delta t \to 0$ yields

$$\begin{aligned} \frac{dV_{\mathcal{I}}(I(t))}{dt} &= \max_{\tau_i(t)\in[0,1]} \{\pi_{\mathcal{I}} - \rho V_{\mathcal{I}}(I(t)) \\ &\quad + \tau_i(t)\alpha\left[V_{\mathcal{S}}(I(t)) - V_{\mathcal{I}}(I(t)) - \frac{c}{\alpha}\right]\} \end{aligned} \tag{77}$$

$$\frac{dV_{\mathcal{S}}(I(t))}{dt} = \pi_{\mathcal{S}} + \beta I(t)\left[V_{\mathcal{I}}(I(t)) - V_{\mathcal{S}}(I(t))\right] - \rho V_{\mathcal{S}}(I(t)) \tag{78}$$

From the coefficient of $\tau_i(t)$, the usual optimality condition is recovered, namely

$$c + \alpha\left[V_{\mathcal{I}}(I(t)) - V_{\mathcal{S}}(I(t))\right] = 0 \tag{79}$$

Next, note that by definition

$$\frac{dV_{\mathcal{I}}(I(t))}{dt} = \frac{dV_{\mathcal{I}}(I(t))}{dI(t)}\dot{I}(t), \quad \frac{dV_{\mathcal{S}}(I(t))}{dt} = \frac{dV_{\mathcal{S}}(I(t))}{dI(t)}\dot{I}(t) \tag{80}$$

For any constant function $I(t)$, the system of differential equations in value functions can be solved simultaneously to yield

$$\rho V_{\mathcal{I}}(I(t)) = \frac{\beta I(t)\,(\pi_{\mathcal{I}} - \tau_i(t)c) + \rho\pi_{\mathcal{I}} + \alpha\tau_i(t)\pi_{\mathcal{S}} - \rho\tau_i(t)c}{\rho + \beta I(t) + \alpha\tau_i(t)} \tag{81}$$

$$\rho V_{\mathcal{S}}(I(t)) = \frac{\pi_{\mathcal{S}}(\rho + \alpha\tau_i(t)) + \beta I(t)\,(\pi_{\mathcal{I}} - \tau_i(t)c)}{\rho + \beta I(t) + \alpha\tau_i(t)} \tag{82}$$

It is easily verified that the gross benefit of treatment equals

$$V_{\mathcal{I}}(I(t)) - V_{\mathcal{S}}(I(t)) = \frac{-(\pi + \tau_i(t)c)}{\beta I(t) + \rho + \alpha\tau_i(t)} = \mu(t) \tag{83}$$

which corresponds to the value derived using optimal control theory. For later reference, note that in the decentralized problem, each agent faces two discrete disease states, while in the centralized problem, the planner faces a "continuous disease state" (i.e. fraction of infected agents).

In the eradication steady state, the payoffs are given by

$$V_{\mathcal{S}}(0) = \frac{\pi_{\mathcal{S}}}{\rho}, \quad V_{\mathcal{I}}(0) = \frac{\pi_{\mathcal{S}}}{\rho} - \frac{\pi + c}{\rho + \alpha} < V_{\mathcal{S}}(0) \tag{84}$$

In the all-infected steady state, the payoffs are given by

$$V_{\mathcal{I}}(1) = \frac{\pi_{\mathcal{I}}}{\rho}, \quad V_{\mathcal{S}}(1) = \frac{\pi_{\mathcal{I}}}{\rho} + \frac{\pi}{\rho+\beta} > V_{\mathcal{I}}(1) \tag{85}$$

Turning to the centralized setup, the discrete-time approximation of the formula characterizing the evolution of social welfare is given by

$$V(I(t)) = \max_{r(t)\in[0,1]} \left\{ I(t)\left(-\pi - r(t)c\right)\Delta t + (1+\rho\Delta t)^{-1} V(I(t+\Delta t)) \right\} \tag{86}$$

Taking a first-order Taylor expansion yields

$$V(I(t+\Delta t)) \approx V(I(t)) + \frac{dV(I(t))}{dt}\Delta t \tag{87}$$

where higher order terms have been omitted. Rearranging, dividing by Δt and taking the limit $\Delta t \to 0$ yields

$$\rho V(I(t)) = \max_{r(t)\in[0,1]} \left\{ I(t)\left(-\pi - r(t)c\right) + \frac{dV(I(t))}{dt} \right\} \tag{88}$$

But by definition,

$$\frac{dV(I(t))}{dt} = \frac{dV(I(t))}{dI(t)}\dot{I}(t) \tag{89}$$

Hence the formula becomes

$$\rho V(I(t)) = \max_{r(t)\in[0,1]} \left\{ I(t)\left(-\pi - r(t)c\right) + \frac{dV(I(t))}{dI(t)} I(t)\left[\beta(1-I(t)) - \alpha r(t)\right] \right\} \tag{90}$$

The next step is to show that

$$\frac{dV(I(t))}{dI(t)} = \frac{-\pi - r(t)c}{\beta I(t) + \rho + \alpha r(t) - \beta(1-I(t))} = \lambda(t) \tag{91}$$

First, executing the maximization problem yields the equation

$$\begin{aligned} 0 \;=\; & \max_{r(t)\in[0,1]} \{ I(t)\left(-\pi - r(t)c\right) - \rho V(I(t)) \\ & + \frac{dV(I(t))}{dI(t)} I(t)\left[\beta(1-I(t)) - \alpha r(t)\right] \} \end{aligned} \tag{92}$$

This yields the familiar optimality condition

$$c + \alpha \frac{dV(I(t))}{dI(t)} = 0 \tag{93}$$

With an optimally chosen policy $r(t)$, it follows that

$$\begin{aligned} 0 \;=\; & I(t)\left(-\pi - r(t)c\right) - \rho V(I(t)) \\ & + \frac{dV(I(t))}{dI(t)} I(t)\left[\beta(1-I(t)) - \alpha r(t)\right] \end{aligned} \tag{94}$$

Differentiating this equation with respect to $I(t)$ and using the optimality condition then yields

$$\begin{aligned} 0 \;=\; & \left(-\pi - r(t)c\right) + \frac{d^2V(I(t))}{dI(t)^2}\dot{I}(t) \\ & - \frac{dV(I(t))}{dI(t)}\left(\rho - \beta(1-I(t)) + \alpha r(t) + \beta I(t)\right) \end{aligned} \tag{95}$$

Last, let

$$\frac{dV(I(t))}{dI(t)} \equiv \lambda(t) \tag{96}$$

and note that

$$\frac{d^2V(I(t))}{dI(t)^2}\dot{I}(t) = \frac{d\lambda(t)}{dI(t)}\dot{I}(t) = \dot{\lambda}(t) \tag{97}$$

This gives the final equation

$$\dot{\lambda}(t) = \lambda(t)\left[\rho - \beta(1 - I(t)) + \alpha r(t) + \beta I(t)\right] + (\pi + r(t)c) \tag{98}$$

Substitution then yields

$$\begin{aligned} \rho V(I(t)) \quad = \quad & \max_{r(t)\in[0,1]} \{(-\pi - r(t)c) \\ & \times \frac{\left[1 + I(t)\left(\beta I(t) + \rho + \alpha r(t) - \beta(1 - I(t))\right)\right]}{\rho + \beta I(t) + \alpha r(t) - \beta(1 - I(t))}\} \end{aligned} \tag{99}$$

In the eradication steady state, social welfare is given by

$$\rho V(0) = \frac{-\pi - c}{\rho + \alpha - \beta} \tag{100}$$

In the all-infected steady state, social welfare is given by

$$\rho V(1) = \frac{-\pi\left(1 + \beta + \rho\right)}{\rho + \beta} \tag{101}$$

REFERENCES

[1] AADLAND, D. AND D. FINNOFF (2007): Syphilis Cycles, *mimeo.*

[2] AHITUV, A., V. J. HOLZ AND T. PHILIPSON (1996): The Responsiveness of the Demand for Condoms to the Local Prevalence of AIDS, *Journal of Human Resources*, 31(4), 869-897.

[3] ANDERSON, R. M. AND R. M. MAY (1991): Infectious Diseases of Humans: Dynamics and Control, *Oxford University Press.*

[4] AULD, M. C. (2006): Estimating Behavioral Response to the AIDS Epidemic, *Contributions to Economic Analysis & Policy*, 5(1), Article 12.

[5] BARRETT, S. (2003): Global Disease Eradication, *Journal of the European Economic Association*, 1(2-3), 591-600.

[6] BOULIER, B. L., T. S. DATTA AND R. S. GOLDFARB (2007): Vaccination Externalities, *B.E. Journal of Economic Analysis & Policy*, 7(1), (Contributions), Article 23.

[7] BRITO, D. L., E. SHESHINSKI AND M. D. INTRILIGATOR (1991): Externalities and Compulsory Vaccinations, *Journal of Public Economics*, 45(1), 69-90.

[8] CHEN, F. AND A. COTTRELL (2009): Dynamic Equilibria in an Epidemic Model with Voluntary Vaccinations, *Journal of Biological Dynamics*, 3(4), 357-375.

[9] DALEY, D. J. AND J. GANI (2001): Epidemic Modelling: An Introduction, *Cambridge Studies in Mathematical Biology.*

[10] DUPAS, P. (2007): Relative Risks and the Market for Sex: Teenage Pregnancy, HIV, and Partner Selection in Kenya, *mimeo.*

[11] Francis, P. J. (2004): Optimal Tax/Subsidy Combinations for the Flu Season, *Journal of Economic Dynamics & Control*, 28(10), 2037 - 2054.

[12] Geoffard, P.-Y. and T. Philipson (1996): Rational Epidemics and Their Public Control, *International Economic Review*, 37(3), 603-624.

[13] Geoffard, P.-Y. and T. Philipson (1997): Disease Eradication: Private versus Public Vaccination, *American Economic Review*, 87(1), 222-230.

[14] Gersovitz, M. and J. S. Hammer (2004): The Economical Control of Infectious Diseases, *Economic Journal*, 114(492), 1-27.

[15] Goldman, S.M. and J. Lightwood (1995): The SIS Model of Infectious Disease with Treatment, *mimeo.*

[16] Goldman, S.M. and J. Lightwood (2002): Cost Optimization in the SIS Model of Infectious Disease with Treatment, *Topics in Economic Analysis and Policy*, 2(1), 1-22.

[17] Hartl, R. F. (1987): A Simple Proof of the Monotonicity of the State Trajectories in Autonomous Control Problems, *Journal of Economic Theory*, 41(1), 211-215.

[18] Hartl, R. F. and G. Feichtinger (1987): A New Sufficient Condition for Most Rapid Approach Paths, *Journal of Optimization Theory and Applications*, 54(2), 403-411.

[19] Jackson, M. O. (2008): Social and Economic Networks, *Princeton University Press.*

[20] Kremer, M. (1996): Integrating Behavioral Choice into Epidemiological Models of AIDS, *Quarterly Journal of Economics*, 111(2), 549-573.

[21] Oster, E. (2005): Sexually Transmitted Infections, Sexual Behavior and the HIV/AIDS Epidemic, *Quarterly Journal of Economics*, 120(2), 467-515.

[22] Oster, E. (2007): HIV and Sexual Behavior Change: Why not Africa?, *mimeo.*

[23] Philipson, T. J. (2000): Economic Epidemiology and Infectious Disease, in *Handbook of Health Economics*, volume 1B, Part 8; edited by J. Newhouse and T. Culyer. New York: North-Holland.

[24] Philipson, T., and R. A. Posner (1993): Private Choices and Public Health: An Economic Interpretation of the AIDS Epidemic. Cambridge, MA: Harvard University Press.

[25] Reluga, T. C. (2009): An SIS Epidemiology Game with Two Subpopulations, *Journal of Biological Dynamics*, 1751-3766.

[26] Rowthorn, R. (2006): The Optimal Treatment of Disease Under a Budget Constraint, in R. Halvorsen and D. Layton (eds), *Explorations in Environmental and Natural Resource Economics: Essays in Honor of Gardner M. Brown, Jr,* Edward Elgar.

[27] Sanders, J. L. (1971): Quantitative Guidelines for Communicable Disease Control Programs, *Biometrics*, 27(4), 883-893.

[28] Sethi, S. P. (1974): Quantitative Guidelines for Communicable Disease Control Program: A Complete Synthesis, *Biometrics*, 30(4), 681-691.

[29] SETHI, S. P. (1977): Nearest Feasible Paths in Optimal Control Problems: Theory, Examples, and Counterexamples, *Journal of Optimization Theory and Applications*, 23(4), 563-579.

[30] SETHI, S. P. (1978): Optimal Quarantine Programmes for Controlling an Epidemic Spread, *Journal of the Operational Research Society*, 29(3), 265-268.

[31] SETHI, S. P., P. W. STAATS (1978): Optimal Control of Some Simple Deterministic Epidemic Models, *Journal of the Operational Research Society*, 29(2), 129-136.

[32] TOXVAERD, F. (2009a): Recurrent Infection and Externalities in Prevention, *mimeo.*

[33] TOXVAERD, F. (2009b): Infection, Acquired Immunity and Externalities in Treatment, *mimeo.*