The Optimal Control of Infectious Diseases via Prevention and Treatment*

Robert Rowthorn† and Flavio Toxvaerd‡

Abstract. This paper characterizes the optimal control of a recurrent infectious disease through the use of prevention and treatment. We find that under centralized decision making, treatment induces positive destabilizing feedback effects, while prevention induces negative stabilizing feedback effects. While optimal treatment pushes prevalence towards the extremes, optimal prevention pushes it towards interior solutions. As a result, the dynamic system may admit multiple steady states and the optimal policy may be history dependent. We find that steady state prevalence levels in decentralized equilibrium must be equal to or higher than the socially optimal levels. The differences between the equilibrium outcome and the social optimum derive from the existence of a pure externality effect and a separate risk effect due to individuals being small. Last, we derive two separate corrective subsidy schemes that decentralize the socially optimal outcome, namely subsidies to prevention and treatment and a tax on the infected.

JEL Classification: C73, I18.

Keywords: Economic epidemiology, treatment, prevention, optimal and equilibrium policy mix, hysteresis, non-convex systems.

1. Introduction

Infectious diseases are a leading cause of morbidity and mortality in both developing and developed countries and impose a major strain on public budgets.[1] In parallel with rapid advances in the biomedical field, there is an ongoing effort to improve disease control through a better use of existing techniques and resources.

A recurrent issue in the debate on infectious diseases is the relative importance of prevention and treatment (Russell, 1986, Krauthammer, 2009, Kremer and Snyder, 2013). Although they are distinct forms of intervention, targeting different individuals, prevention and treatment cannot be evaluated in isolation since their effects interact. Prevention reduces the rate at which individuals become infected, thereby reducing the future need for treatment. In contrast, treatment reduces the proportion of the population who are infected, thereby reducing the risk of infection and the future need for protection. An optimal policy will typically combine both forms of intervention in proportions that vary

*We thank William Brock, Partha Dasgupta, Mark Gersovitz, Chryssi Giannitsarou, Steven Goldman, Philipp Kircher, Saeed Moghayer, Selma Telalagic, Florian Wagener, Jörgen Weibull, Peter White, Tasos Xepapadeas and seminar participants at the University of Cambridge, Columbia University, Wake Forest University, Johns Hopkins University, SUNY at Stony Brook, University of Glasgow, Royal Holloway University of London, University of Southampton, Institute for Advanced Study Toulouse and the Health Protection Agency's Modelling and Economics Unit for constructive feedback. We are also grateful for the constructive suggestions of the editor and the anonymous referees.

†Faculty of Economics, University of Cambridge and King's College, Cambridge.

‡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/.

[1]See Piot et al. (2008).

through time. Thus, rather than asking whether prevention or treatment is best for managing infectious diseases, the aim should be to determine the appropriate combination of these two interventions at any given time. This is in keeping with the approach advocated by Rose (1985, 1992), one of the central thinkers in the field of public health.

To take full account of the interactions between prevention and treatment, requires a unified intertemporal model. To this end, we study a simple susceptible-infected-susceptible (SIS) model, in which individuals are either infected or susceptible. No-one is immune. Examples of diseases that have been modeled within the SIS framework, and for which there is both prevention and treatment available, include gonorrhea (Hethcote and Yorke, 1984), malaria (Anderson et al. 2012) and syphilis (Cannefax, 1965).[2] We assume that new infections can be avoided through costly prevention and that individuals can be cured through costly treatment.[3]

Our analysis is conducted in three steps. First, we analyze *centralized decision making*, in which a benevolent social planner implements the command optimum. We derive the optimal policy, steady states and transition paths. Next, we analyze *uncontrolled decentralized decision making*, i.e. the market solution that obtains when forward-looking individuals behave non-cooperatively without any inducements from the planner. We pay special attention to the differences between the resulting equilibrium outcomes and those preferred by the social planner. Last, we analyze *controlled decentralized decision making*, i.e. the equilibrium solution that obtains when the social planner offers taxes and/or subsidies with a view to aligning public and private incentives.[4]

A major advantage of considering treatment and prevention within a unified framework, is that it helps organize and clarify results that are known from single-instrument models. Thus we can both analyze the interaction of multiple policies and obtain existing models as special cases. This makes it easier to trace different effects to specific policy instruments. Despite superficial similarities, prevention and treatment turn out to be profoundly different in their effects and desirability at different levels of disease prevalence (i.e. the fraction of the population that is infected). For example, the marginal benefit of treatment is a decreasing function of disease prevalence. The more prevalent a disease is, the greater is the risk of re-infection following cure, and hence the lower the benefit from treatment. In contrast, the marginal benefit of prevention is an increasing function of disease prevalence. The more prevalent a disease is, the greater is the risk of infection and the greater is the benefit from protection. In the terminology of Brock and Starrett (2003), with treatment there is a *destabilizing positive feedback effect* while with prevention, there is a *stabilizing negative feedback effect*. This is one of the central findings of our analysis and drives a number of other interesting results. For example, the positive destabilizing effect of treatment generates multiple potential steady states, while the negative stabilizing effect of prevention implies that it is suboptimal to eradicate the disease through prevention.

In general, for extreme levels of disease prevalence, treatment and prevention will

[2] According to Cannefax (1965), '*the cycle of cure, re-infection, cure, re-infection etc. occurs so frequently in given individuals that the term "ping-pong" syphilis [...] was coined to describe this frequent clinical observation.*'

[3] We focus on temporary measures that must be sustained through time in order to remain effective. In particular, we exclude vaccinations which confer prolonged (or permanent) immunity.

[4] We use the terms *controlled* and *uncontrolled* settings in the sense of Arrow and Kurz (1969). This should not be confused with the use of *control variables* in solving optimal control problems.

tend to be strong substitutes and used in very asymmetric proportions, whereas for intermediate prevalence levels they are weaker substitutes, such that it may be optimal to use them in conjunction. Along optimal paths, treatment and prevention are always at their maximum or minimum possible levels, whereas this is not true once a steady state is reached.

Next, we find that under uncontrolled, decentralized decision making, disease prevalence may be socially suboptimal. This will occur when the central planner chooses a higher level of protection than individuals would choose if left to their own devices. While steady state treatment levels under centralized and decentralized decision making may coincide, the corresponding steady state prevention levels only coincide when they are optimally equal to zero. Whenever prevention is actively used in steady state, its level under decentralization is suboptimally low, thereby distorting disease prevalence upwards. On the transition paths, centralized and decentralized treatment and prevention levels may coincide, even if they do not coincide once steady state is reached.

In comparing the uncontrolled equilibrium outcomes with those under social planning, we show that the differences in valuations between the planner and the individuals can be usefully decomposed into a conventional external effect and a risk effect. The former effect derives from the fact that individuals take no account of the harm their actions impose on others, while the latter derives from the fact that numerically insignificant individuals take aggregate disease prevalence as given, while the planner can directly control it.

Last, we consider the equilibrium outcomes under controlled decentralized decision making. We derive two different incentive schemes that decentralize the command optimum. In one scheme, the planner offers individuals (state dependent) subsidies to prevention and treatment. When faced with these subsidies, the equilibrium outcome under decentralized decision making exactly mimics that chosen by the social planner under centralized decision making. In the other scheme, the planner imposes a tax on infected individuals (or equivalently, gives a bonus to healthy individuals). This scheme also decentralizes the first-best solution. We also offer some discussion of the possibly perverse effects of non-optimal corrective measures such as simple fixed subsidies to prevention and treatment and the potential for simplified optimal incentive schemes.

1.1. Related Literatures. The literature on economic epidemiology is varied and growing. There are several good surveys, such as Philipson (2000), Gersovitz and Hammer (2003) and Klein et al. (2007). Of direct relevance to the present work is research that deals with prevention and treatment, separately or in conjunction.

The earliest contributions, by Sanders (1971), Sethi (1974) and Sethi and Staats (1978), consider treatment in different versions of the SIS model from a planner's perspective. Goldman and Lightwood (1995) consider treatment in the SIS model under learning, while Goldman and Lightwood (2002) also study treatment in the controlled SIS model, but considers different cost structures than the earlier literature. Goldman and Lightwood's (2002) analysis focuses mainly on necessary conditions for optimality and provide an informal analysis using phase diagrams. Both Sanders (1971) and Sethi (1974) assume that individual treatment cost is a sharply decreasing function of prevalence, whereas Goldman and Lightwood (2002) consider linear or increasing costs. Non-linear cost structures are most appropriate when considering the running of a health authority (i.e. the problem of a social planner) rather than to analyze the incentives

of individuals. Rowthorn (2006) and Anderson et al. (2012) extend the analysis of the controlled SIS model to settings with budget and wealth constraints. Toxvaerd (2009a) considers decentralization to strategic decision makers and the possibility of multiple equilibria (rather than merely multiple steady states), while Toxvaerd (2009b) considers the effects of treatment when recovery confers immunity to further infection.

The literature on prevention is more varied than that on treatment. Sethi (1978) considers quarantines, while Geoffard and Philipson (1996) and Aadland et al. (2010) consider non-vaccine prevention in the SI and SIS models respectively. Reluga (2009) analyzes prevention by strategic individuals in linked subpopulations, while Reluga (2010) considers prevention through social distancing**.** Toxvaerd (2010) analyzes continual prevention in the SIS model and decentralization of optimal policy to strategic decision makers. There are also important literatures on vaccination and on abstinence, exemplified by Brito et al. (1991) and Kremer (1996), respectively. The issues dealt with in those papers are somewhat orthogonal to the present work and are reviewed in more detail in Toxvaerd (2010). Greenwood et al. (2009) consider a search-theoretic matching model of the SI variety and analyze the incentives to form long and short term partnerships.

There are a few papers that explicitly consider multiple instruments. Most related to our work is that of Gersovitz and Hammer (2004) who, like us, consider prevention and treatment in an SIS framework. In contrast to us, they bypass the issue of multiplicity by assuming that there is a unique steady state. Furthermore, they assume that the steady state is interior in both control variables. As we will show, these assumption have radical consequences for both the analysis and the conclusions derived from it. In a short note, Zaman et al. (2007) consider vaccination and treatment in an SIR setting and simulate optimal paths. A similar exercise is done in Almeder et al. (2007) for an HIV type disease. Goyal and Vigier (2010) consider a static two stage model with vaccination and abstinence. Dodd et al. (2010) consider multiple concurrent interventions and discuss when there are likely to be synergies between these in the sense that raising the level of one instrument increases the benefit to increasing the level of other instruments. Blayneh et al. (2009) consider multiple interventions in a setting with a vector-borne disease, as do Agusto et al. (2012). Feichtinger (1984) and Behrens et al. (2000) analyze models that are structurally similar to ours, but which deal with non-disease applications. Apart from Gersovitz and Hammer (2004), these papers are similar to ours only in spirit and their analyses are not directly comparable to the one we carry out. Our work is also related to the empirical work by Cohen et al. (2011), who consider the (possibly perverse) effects of subsidies to malaria treatment and diagnostic tests. We will discuss this contribution further in the context of the implementation of socially optimal outcomes through different incentive schemes.

Last, our paper contributes to an important literature on equilibrium multiplicity and history dependence in systems with non-convexities (in both economics and ecology) as surveyed in Dasgupta and Mäler (2003), Brock and Starrett (2003), Wagener (2003), Mäler et al. (2003), Deissenberg et al. (2004) and Horan et al. (2011).

The remainder of the paper is structured as follows. In Section 2, we outline the classical susceptible-infected-susceptible model. In Section 3, we introduce the economic version of the model and partially characterize the optimal policies. In Section 4, we characterize the steady states of the system and the optimal paths formally. In Section 5, we describe the equilibria and dynamics of the model and interpret the central features

driving the results. In Section 6, we analyze the equilibria under decentralized decision making and compare these to the command optimum. In Section 7, we offer the effect decomposition result and show how to decentralize the command optimum via taxes and subsidies. Section 8 concludes. Most proofs are found in appendices and supplementary material on different aspects of the planner's solution is available in an online appendix.

2. The Classical SIS Model

We start by expounding the classical epidemiological version of the susceptible-infected-susceptible model in some detail. This will not only aid in understanding the economic model that follows, but also highlight the contrast in predictions based on the separate modeling approaches.

The classical SIS model is simple to describe.[5] 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 $t \geq 0$ 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 is normalized to unity, these measures can be interpreted as fractions. Henceforth, $I(t)$ will be referred to as *disease prevalence*.

At each instant, the population mixes homogeneously. This corresponds to pairwise 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 transmitted 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 $\beta I(t)S(t)$. Thus the rate of new infection, or *disease incidence*, is proportional to disease prevalence.[6] Note that while disease incidence is a flow, disease prevalence is a stock.

Finally, infected individuals recover spontaneously at rate $\gamma \geq 0$. This means that the aggregate rate at which infected individuals become susceptible is given by $\gamma I(t)$.

The dynamics of the model are described by the following system of differential equations:

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

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

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

This system reduces to the following simple logistic growth equation:

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

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

[6] 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 $\hat{I} = 0$ and $\hat{I} = (\beta - \gamma)/\beta$. For $\beta > \gamma$, the stable steady state is such that the disease is endemic while for $\beta < \gamma$, the relevant and stable steady state is such that the disease is eradicated. 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. Last, note that disease prevalence in the endemic steady state is increasing in infectivity and decreasing in the rate of recovery.

At the aggregate level, there is no uncertainty and thus the probability that a randomly chosen individual is infected must coincide with the fraction of infected individuals. From the perspective of an infected individual, the transition to susceptibility is governed by a Poisson process with rate γ, which is memoryless. Similarly, for a fixed level of aggregate infection $I(t)$, the transition to infectivity for a susceptible individual is governed by a Poisson process with rate $\beta I(t)$. Thus transition probabilities are memoryless, a fact that greatly simplifies the analysis that follows.

For simplicity, we will assume throughout that both the incubation period and the latency period have zero length. Furthermore, there is no uncertainty about individuals' health status. This means that individuals in each category, i.e. infected and susceptible, can be perfectly identified and thus targeted for treatment and prevention respectively.

3. The Economic Model and Optimal Policies

In the classical version of the model, there is no behavior or decision making and thus the model is lacking as a vehicle for analyzing human populations. To study an economic version of the model, assume that each individual earns flow payoffs that depend on the state of their health. For simplicity, assume that an individual earns flow payoff $\omega_{\mathcal{S}}$ while susceptible and $\omega_{\mathcal{I}} < \omega_{\mathcal{S}}$ while infected. Let $\omega \equiv \omega_{\mathcal{S}} - \omega_{\mathcal{I}} > 0$ be the health premium. The future is discounted at rate $\rho > 0$. The basic epidemiological parameters $\beta > 0$ (infectiousness) and $\gamma > 0$ (background rate of spontaneous recovery) are retained from the classical model.

The two policy instruments at the decision maker's disposal are *prevention* and *treatment*. These instruments influence the flows from $\mathcal{S}(t)$ to $\mathcal{I}(t)$ and from $\mathcal{I}(t)$ to $\mathcal{S}(t)$ respectively. Specifically, a planner can set some level of prevention $\pi(t) \in [0, 1]$ at time $t \geq 0$, which translates into effective disease incidence $(1 - \pi(t))\beta I(t)S(t)$. The factor $(1 - \pi(t))$ can be thought of as the proportion of susceptible individuals who are exposed at time $t \geq 0$. Turning to treatment, the planner can set the level of treatment $\tau(t) \in [0, 1]$ at time $t \geq 0$, which translates to an effective recovery rate $(\tau(t)\alpha + \gamma)$. Here, $\alpha > 0$ is the efficiency of treatment in inducing recovery. Last, the marginal costs of protection and treatment (i.e. per individual per instant) are $c_P \geq 0$ and $c_T \geq 0$ respectively.

We now consider the optimal control of the SIS system from the perspective of a benevolent social planner. The planner's objective is assumed to be a straightforward sum of the individuals' infinite horizon, discounted expected utilities. The planner's problem is therefore to solve the following program:

$$\max_{\tau(t),\pi(t)\in[0,1]} \int_0^\infty e^{-\rho t} \left[I(t) \left(\omega_I - c_T \tau(t) \right) + (1 - I(t))(\omega_{\mathcal{S}} - c_P \pi(t)) \right] dt \tag{5}$$

$$s.t. \quad \dot{I}(t) = (1 - \pi(t))\beta I(t)(1 - I(t)) - (\gamma + \alpha\tau(t))\, I(t), \quad I(0) = I_0 \tag{6}$$

where $\pi(t), \tau(t)$ are piecewise continuous.

The optimal value function for this program is denoted by $V(I_0)$, where dependence on the parameters has been suppressed for ease of notation. It can be shown that this problem admits an optimal solution under mild conditions. A proof is contained in an online appendix available from the authors. Note that the problem to be solved is autonomous, i.e. time enters in the integrand only through the discount term $e^{-\rho t}$. This has the important implication that $I(t)$ is monotonic along an optimum path.

Throughout this paper, we will maintain the following:

Assumption (i) $\omega - c_P > 0$ and (ii) $\beta - \gamma - \alpha > 0$.

The inequality (i) implies that a susceptible individual would always choose full protection, if the only alternative was to become instantly and permanently infected. The inequality (ii) implies that a policy without prevention, but with maximal treatment, cannot eradicate infection even asymptotically.[7] These are required to make the tradeoffs in the model interesting.

The current-value Hamiltonian for this problem is given by

$$\begin{aligned} H^C \equiv & -\omega I(t) - c_P \pi(t)(1 - I(t)) - c_T \tau(t) I(t) \\ & + \lambda(t) \left[(1 - \pi(t))\beta I(t)(1 - I(t)) - (\gamma + \alpha\tau(t))\, I(t) \right] \end{aligned} \tag{7}$$

where $\lambda(t)$ is the current-value costate variable (or shadow price).[8] The above Hamiltonian is linear in both control variables, which has important implications for the characterization of optimal policies.

For the solution to be optimal, $\lambda(t)$ must satisfy the following differential equation:

$$\dot{\lambda}(t) = \lambda(t) \left[\rho + \gamma + \alpha\tau(t) + (1 - \pi(t))\beta(2I(t) - 1) \right] + \left[\omega + \tau(t) c_T - \pi(t) c_P \right] \tag{8}$$

Moreover, the instruments $(\tau(t), \pi(t))$ must maximize the current-value Hamiltonian (7).

[7] To interpret the former assumption, suppose that there is no spontaneous recovery or treatment ($\alpha = \gamma = 0$) and that unprotected individuals are immediately infected ($\beta \to \infty$). This is a worst case scenario that makes prevention as useful an instrument as possible. In this setting, the effective choice is between perpetual protection at per instant cost c_P or perpetual infection at per instant cost ω. Assumption (i) ensures that under this scenario, prevention will be socially useful. To interpret the latter assumption, suppose that there is no prevention but full treatment. If (i) is violated, prevention is never used in steady state. Assumption (ii) then simply means that the rate of infection is higher than the effective rate of recovery. If (ii) is violated, then full treatment will eradicate the disease.

[8] Note that a term $\omega_{\mathcal{S}}$ has been dropped from the current-value Hamiltonian because its presence does not affect the optimal solution.

This yields the following Hamiltonian conditions for optimality:

$$\tau(t) = 0 \quad if \quad \alpha\lambda(t) > -c_T \tag{9}$$
$$\tau(t) \in [0,1] \quad if \quad \alpha\lambda(t) = -c_T \tag{10}$$
$$\tau(t) = 1 \quad if \quad \alpha\lambda(t) < -c_T \tag{11}$$

and

$$\pi(t) = 0 \quad if \quad \beta\lambda(t)I(t) > -c_P \tag{12}$$
$$\pi(t) \in [0,1] \quad if \quad \beta\lambda(t)I(t) = -c_P \tag{13}$$
$$\pi(t) = 1 \quad if \quad \beta\lambda(t)I(t) < -c_P \tag{14}$$

An additional necessary condition for optimality in this setting is that[9]

$$lim_{t\to\infty} e^{-\rho t} H^C(t) = 0 \tag{15}$$

The above Hamiltonian conditions imply that if the marginal benefit of increasing an instrument (i.e. treatment or prevention) exceeds the marginal cost of doing so, then it is optimal to increase the level of the instrument. Similarly, if the marginal cost exceeds the marginal benefit, then it is optimal to decrease the level of the instrument. To see this, recall that $\lambda(t) < 0$ is the (negative) social utility associated with a marginal increase in disease prevalence. With this in mind, it is straightforward to interpret the optimal policies in terms of the marginal costs and benefits of intervention. In the case of treatment, the marginal benefit of intervention is given by $-\alpha\lambda(t)$, which follows from the fact that α is the rate at which increased treatment induces recovery (i.e. it is the efficiency of treatment) and each recovery benefits society at level $-\lambda(t)$. In the case of preventive effort, the marginal benefit of intervention is given by $-\beta I(t)\lambda(t)$. This follows since $\beta I(t)$ is the rate at which unprotected susceptible individuals become infected and each infected individual costs society $\lambda(t)$.

Before continuing with the detailed analysis, a few comments on our modeling choices are in order. First, in virtually all existing models of treatment, e.g. Sanders (1971), Sethi (1974), Goldman and Lightwood (1996, 2002), Rowthorn (2006), Anderson et al. (2012) and Agusto et al. (2012), therapy is either modeled as a discrete (i.e. zero-one) choice, or as a continuous choice with constant returns to scale. If one allows for randomization, these modeling choices are of course equivalent. More to the point, the discrete nature of treatment (or the equivalent alternative assumptions) are descriptively accurate representations of how infectious diseases are treated in practice. As pointed out by McKinnon and Davis (2004), both so-called 'time-dependent' and 'concentration-dependent' treatments mandate that sufficiently high concentrations of the medicine in the blood is reached and sustained for a minimum period of time. But this does not mean that the treatment is scalable in practice. If the dose is below the prescribed target, then the infection will not be effectively eliminated and resistance may be induced. Similarly, if the dose is above the target, it can create adverse effects like abdominal pain, diarrhea and organ failure, without increasing the speed or probability of recovery.[10] In terms of

[9] $H(t) \equiv e^{-\rho t} H^C(t)$ is the conventional discounted Hamiltonian. It is shown in Michel (1982) that a necessary condition for optimality is that $\lim_{t\to\infty} H(t) = 0$.

[10] In addition, the World Health Organization (2004) recommends that sexually transmitted diseases

modeling of treatment, the exceptions are Gersovitz and Hammer (2004) and Gersovitz (2010), who assume that treatment is continuous but subject to decreasing returns to scale. These assumptions mean that their treatment variable cannot be interpreted as a randomization over discrete choices and must be interpreted literally. As noted above, such an assumption seems to be problematic on a purely descriptive level, but as will become clear in what follows, will also have important consequences for the characterization of optimal policy.

Second, we have for simplicity assumed that prevention is perfect in the sense that with full prevention, the infection probability is exactly zero. Our main results carry over to a setting with imperfect protection, but the analysis is considerably less transparent. The model with imperfect prevention is set out in the online appendix.

4. Optimal Paths and Steady States

We now proceed with a detailed formal analysis of the optimal paths and the steady states of the system, through a number of propositions. In the next section, we will offer a more informal discussion of these results. The dynamic system defined by equations (6) and (8) is in a steady state when all variables are constant, i.e. when $I(t) = \hat{I}, \lambda(t) = \hat{\lambda}, \tau(t) = \hat{\tau}$ and $\pi(t) = \hat{\pi}$. A steady state is said to be *feasible* if it also satisfies the Hamiltonian conditions (9) to (15) and $\hat{I} \in [0, 1]$.

Proposition 1. *Any optimal path converges to a feasible steady state in which* $\hat{I} > 0$ *and* $\hat{\pi} < 1$. *This implies that eradication is never optimal, even asymptotically.*

Proof: See Appendix B for the proof that under our parameter assumptions, eradication is never optimal. The proof follows from the observation that any path that leads towards eradication does not satisfy the transversality condition (15) ■

The intuition for this result is that when prevention is kept at a level that forces the disease towards eradication, the marginal value of prevention becomes negligible and prevention is no longer cost effective. Since the problem is autonomous and there is only one state variable $I(t)$, this variable is monotonic along an optimal path. Such a path must therefore converge to some $\hat{I} > 0$. At this point it must be the case that $\hat{\pi} < 1$.

Proposition 2. *The dynamic system defined by equations (6) and (8) has six potentially feasible steady states with* $\hat{\pi} < 1$. *These are characterized as follows: Solution* A *:* $\hat{\tau} = 0$ *and* $\hat{\pi} \in (0, 1)$. *Solution* B*:* $\hat{\tau} = 1$ *and* $\hat{\pi} \in (0, 1)$. *Solution* C*:* $\hat{\tau} \in (0, 1)$ *and* $\hat{\pi} \in (0, 1)$. *Solution* A_0*:* $\hat{\tau} = 0$ *and* $\hat{\pi} = 0$. *Solution* B_0*:* $\hat{\tau} = 1$ *and* $\hat{\pi} = 0$. *Solution* C_0*:* $\hat{\tau} \in (0, 1)$ *and* $\hat{\pi} = 0$. *For any given set of parameter values, it is not possible for both* A *and* A_0 *or for both* B *and* B_0 *or for both* C *and* C_0 *to be simultaneously feasible.*

The six potentially feasible steady state values are listed in Appendix A. The final part of this proposition is established by comparing the parameter restrictions under which these various steady states satisfy the Hamiltonian conditions (9) to (15) and $\hat{I} \in [0, 1]$.

Proposition 3. *There is no optimal path that terminates at either* C *or* C_0. *Depending on the parameter values, at least one and at most two of the steady states* A, A_0, B, B_0 *is the end point of an optimal path.*

be treated with single-dose-therapy as this greatly increases adherence.

Proof: The proof that there is no optimal path that terminates at either C or C_0 is outlined in the online appendix. When feasible, each of these points is a spiral source in (I, λ) space. The remainder of the proposition follows directly from Propositions 1 and 2 ■

When multiple feasible steady states coexist, we can talk of a high prevalence steady state and a low prevalence steady state. In the former steady state, no-one receives treatment, while in the latter steady state, all infected individuals receive treatment. Whereas a feasible steady state may involve keeping prevention at an interior level, the approach to such a steady state always involves maximal or minimal levels of the two policy instruments, as the following result shows:

Proposition 4. *The optimal policy is always of the bang-bang variety. Along the approach path to a steady state, both* $\tau(t), \pi(t) \in \{0, 1\}$ *for all* $t \geq 0$*, except at a finite number of points where there is an instantaneous switch from one control regime to another.*

Proof: This follows directly from the Hamiltonian conditions ■

This result has the implication that from any initial condition, the transition to the steady state is of finite duration.

In the fully interior, but suboptimal steady state C, it follows from the Hamiltonian conditions (each holding with equality) that the marginal cost of prevention relative to that of treatment equals the marginal benefit of prevention, again relative to that of treatment. This kind of equation is the central characterization of optimality in the work of Gersovitz and Hammer (2004), and follows from their twin assumptions that their steady state is unique and that it is fully interior. In contrast, in our setup, we make no such assumptions and find that it is generically true that

$$\frac{\alpha\lambda(t)}{\beta I(t)\lambda(t)} \neq \frac{c_T}{c_P} \tag{16}$$

This result holds everywhere on the transition path, except perhaps at a finite number of switching points, and always holds at the endpoint of an optimal path. In other words, we have proved in our setting that along any optimal path, it is generically true that the relative marginal benefits of the two interventions and their relative marginal costs differ.

5. Description of the Dynamics

The key to understanding the dynamics of the model, is to appreciate the differences between treatment and prevention. In turn, these differences stem from the ways in which the marginal benefits of each instrument depend on disease prevalence. In the case of prevention, the marginal benefits are *increasing* in prevalence: other things equal, higher disease prevalence increases the risk of infection for susceptible individuals and hence increases the return from prevention. Since the value of prevention is increasing in prevalence, a higher level of prevention (which suppresses incidence) reduces the value of additional prevention. Similarly, reducing prevention increases incidence, thereby making prevention more valuable at the margin. The effect of this is that prevention, seen in isolation, tends to force the system towards a unique and interior steady state. That is, prevention creates a *negative stabilizing feedback effect.*

Turning to treatment, the time profile of the benefits is more complex than that for prevention in that the benefits accrue in the future. Treatment increases the proportion of time that a typical individual will spend in the susceptible state. For a given susceptible individual, the probability of infection (or reinfection) is proportional to disease prevalence. The value of treating an individual in the present is therefore a decreasing function of future prevalence. As current treatment is increased, future prevalence decreases, making current and future treatment even more attractive. This virtuous circle (which is formally a complementarity property of the planner's problem) means that with treatment, the marginal benefits are decreasing in prevalence.[11] Thus treatment creates a *positive destabilizing feedback effect*, which is exactly what creates the scope for multiple extremal steady states. In the low infection steady state, the marginal benefits from treatment are high and treatment is thus exerted at the highest possible level, thereby maintaining low infection. In the high infection steady state, the marginal benefits of treatment are low and therefore there is no treatment at all. This keeps the infection at a high level.

Once both instruments are available, the forces described above are essentially superimposed. The presence of treatment creates the potential for multiple steady states, even in the presence of prevention (although the levels are altered accordingly). In the full treatment steady state, disease prevalence is relatively modest. But this means that the marginal benefit of prevention is relatively low, resulting in a low steady state level of prevention. In contrast, in the no treatment steady state, disease prevalence is relatively high, leading to high marginal benefits of prevention. As a consequence, in this steady state the prevention level is relatively high.

5.1. Informal Bifurcation Analysis. Given the complexity of the model, it may be tempting to proceed with the analysis by comparing the welfare levels associated with the many different steady states and then simply steer the system towards the steady state with the highest level of welfare. It turns out that this approach is entirely inappropriate, since the planner is seeking to maximize aggregate discounted welfare rather than steady state welfare. The right way forward is, for a given initial condition, to compare the discounted aggregate welfare along all feasible paths. The superior path then dictates the optimal policy. In this and the following subsection, we outline a systematic approach to such an analysis and give a specific numerical example of how a simple-minded focus on steady state welfare levels can lead to the wrong policy conclusions. The reason that we emphasize this point is that in large parts of the existing literature on infection control, the focus is on steady states rather than equilibrium paths. This focus is unwarranted and may be misleading.

Following Wagener (2003), we can usefully divide the parameter space into three different regimes as follows. In Regime I, there is a unique optimal steady state (i.e. a unique end point of an optimal path) from the set $\{A, B, A_0, B_0\}$. Which of these is feasible depends on the particular parameter constellation in question. In Regime II, there are four potential pairs of saddle-points, namely $\{(A, B), (A_0, B_0), (A, B_0), (A_0, B)\}$, each possibly with an accompanying unstable steady state from the set $\{C, C_0\}$. From each such pair of saddle-points, one or the other equilibrium is always optimal, i.e. is

[11] For a concrete and documented syphilis-related example of current decisions influencing the probability of future reinfection, see Stewart et al. (1951).

the end point of an optimal path for all initial conditions (i.e. the steady state is *globally* optimal). In Regime III, there are again four possible pairs of saddle-points (possibly with corresponding unstable steady states) like in Regime II, but different initial conditions render different equilibria optimal. In this scenario, there is an indifference (or Skiba) point $I_S \in (0,1)$ such that for prevalence levels above this threshold, the high infection steady state is the end point of the optimal path, while for prevalence levels below it, the low infection steady state is the end point of the optimal path. Regime II can be seen as a special case of Regime III, in which the Skiba threshold is outside the unit interval. The three possible regimes can be summarized as follows:

Regime I: Unique eq. point from	$\{A, B, A_0, B_0\}$, path independence
Regime II: Pair of eq. points from	$\{(A,B),(A_0,B_0),(A,B_0),(A_0,B)\}$, path independence
Regime III: Pair of eq. points from	$\{(A,B),(A_0,B_0),(A,B_0),(A_0,B)\}$, path dependence

Even though the interior solutions cannot be end points of optimal paths, it is tempting to think that they demarcate intervals of the state variable from which it is optimal to go to one steady state or the other. For example, it might seem natural that for prevalence levels $I(t) < I_C$, the optimal policy is to go to the low infection steady state I_B, while for prevalence levels $I(t) > I_C$ the optimal policy is to go to the high infection steady state I_A. In fact, this turns out to be wrong. While the optimal policy may indeed have the threshold character just described, the critical prevalence level I_S is generically different from the interior steady state.[12]

For a given set of parameters, it is a routine matter to check the feasibility conditions and determine whether Regime I obtains or not. In order to determine whether the system is in Regime II or III, there is no option but to compute values along all (typically two) paths satisfying the necessary Hamiltonian conditions for optimality. This is because the existence of the indifference point that distinguishes Regimes II and III cannot be formally characterized by a local condition in the same way that local extrema can (see Deisssenberg et al. 2004). This is so since the indifference point is obtained as the point of intersection of two functions for which there are no closed form solutions, namely the value functions evaluated along the different candidate paths.

5.2. Simulated Paths and Steady States. To better illustrate the main features of the analysis in the preceding sections, we now consider some sample simulations of optimal paths and steady states. The simulations were done using a fourth-order Runge-Kutta procedure with the following parameter values:

Parameters	α	β	γ	ω	ρ	c_P	c_T
Values	$\{0.2, 0.4, 0.5\}$	3	0.1	1	0.11	0.5	10

With this choice of the parameters $(\beta, \gamma, \omega, \rho, c_P, c_T)$, the feasible steady states are (A, B, C) and the system is either in Regime II or III, depending on the magnitude of the efficiency of treatment α. This means that both the low and the high infection steady

[12] This property *does* hold when the Hamiltonian is concave, as described in Deissenberg et al. (2004).

states exist. The following table shows the ranges for α where each regime obtains:

Interval	$\alpha \in [0, 0.3]$	$\alpha \in [0.3, 0.41]$	$\alpha \in [0.41, 1]$
Opt. steady state	Point A (Reg. II)	Point A or B (Reg. III)	Point B (Reg. II)

We will consider three specific examples as follows. In Example 1, $\alpha = 0.2$, in Example 2, $\alpha = 0.5$ and in Example 3, $\alpha = 0.4$. The optimal policies corresponding to the paths in the three examples are summarized in the following table:

Example 1 ($\alpha = 0.2$)			**Example 2** ($\alpha = 0.5$)		
Optimal path goes to A	$\tau(t)$	$\pi(t)$	Optimal path goes to B	$\tau(t)$	$\pi(t)$
$I \in [0, 0.0031]$	1	0	$I \in [0, 0.0018]$	1	0
$I \in [0.0031, 0.0370]$	0	0	$I = 0.0018 (= I_B)$	1	0.7996
$I = 0.0370 (= I_A)$	0	0.9654	$I \in [0.0018, 0.0176]$	1	1
$I \in [0.0370, 1]$	0	1	$I \in [0.0176, 1]$	0	1
Example 3 ($\alpha = 0.4$)			**Example 3** ($\alpha = 0.4$)		
Optimal path goes to B	$\tau(t)$	$\pi(t)$	Optimal path goes to A	$\tau(t)$	$\pi(t)$
$I \in [0, 0.0018]$	1	0	$I \in [0.0163, 0.0370]$	0	0
$I = 0.0018 (= I_B)$	1	0.7996	$I = 0.0370 (= I_A)$	0	0.9654
$I \in [0.0018, 0.0115]$	1	1	$I \in [0.0370, 1]$	0	1
$I \in [0.0115, 0.0163]$	0	1			

In Figure 1, we show the simulated candidate paths and associated value functions for the three examples. For completeness, note that the kinks in the optimal paths in the three graphs correspond to switches in the control regimes. In Example 1, it is optimal to pursue the path to steady state A for any initial level of disease prevalence (this case is in Regime II). The paths to the two steady states A and B are illustrated in the first panel of Figure 1. In the second panel, we show the total discounted value of following the paths to steady states A and B respectively, for different initial prevalence levels. It is clear from this figure that the value of going to (and staying at) point A is everywhere higher than the value of going to (and staying at) point B.

In Example 2, it is optimal to follow the path to steady state B for any initial prevalence level (this case is also in Regime II). The paths to A and B are shown in the third panel of Figure 1, which also shows the corresponding values of following the different paths in the fourth panel. It is clear from the figure that going to (and staying at) point B always dominates going to (and staying at) point A.

In Example 3, the system is in Regime III, in which the optimal steady state depends on the initial level of infection. This case is also illustrated in Figure 1, in the fifth and sixth panels. For prevalence levels below $I_S = 0.0163$, the optimal path leads to the low infection steady state B, while for prevalence levels above $I_S = 0.0163$, the optimal path leads to the high infection steady state A. Thus for this parameter constellation, the optimal path is history dependent in the sense that the initial conditions matter for where it is optimal for the system to settle. Note that in the relevant panel of Figure 1, $I_S = 0.0163$ is the prevalence level at which the value functions for the paths to A and

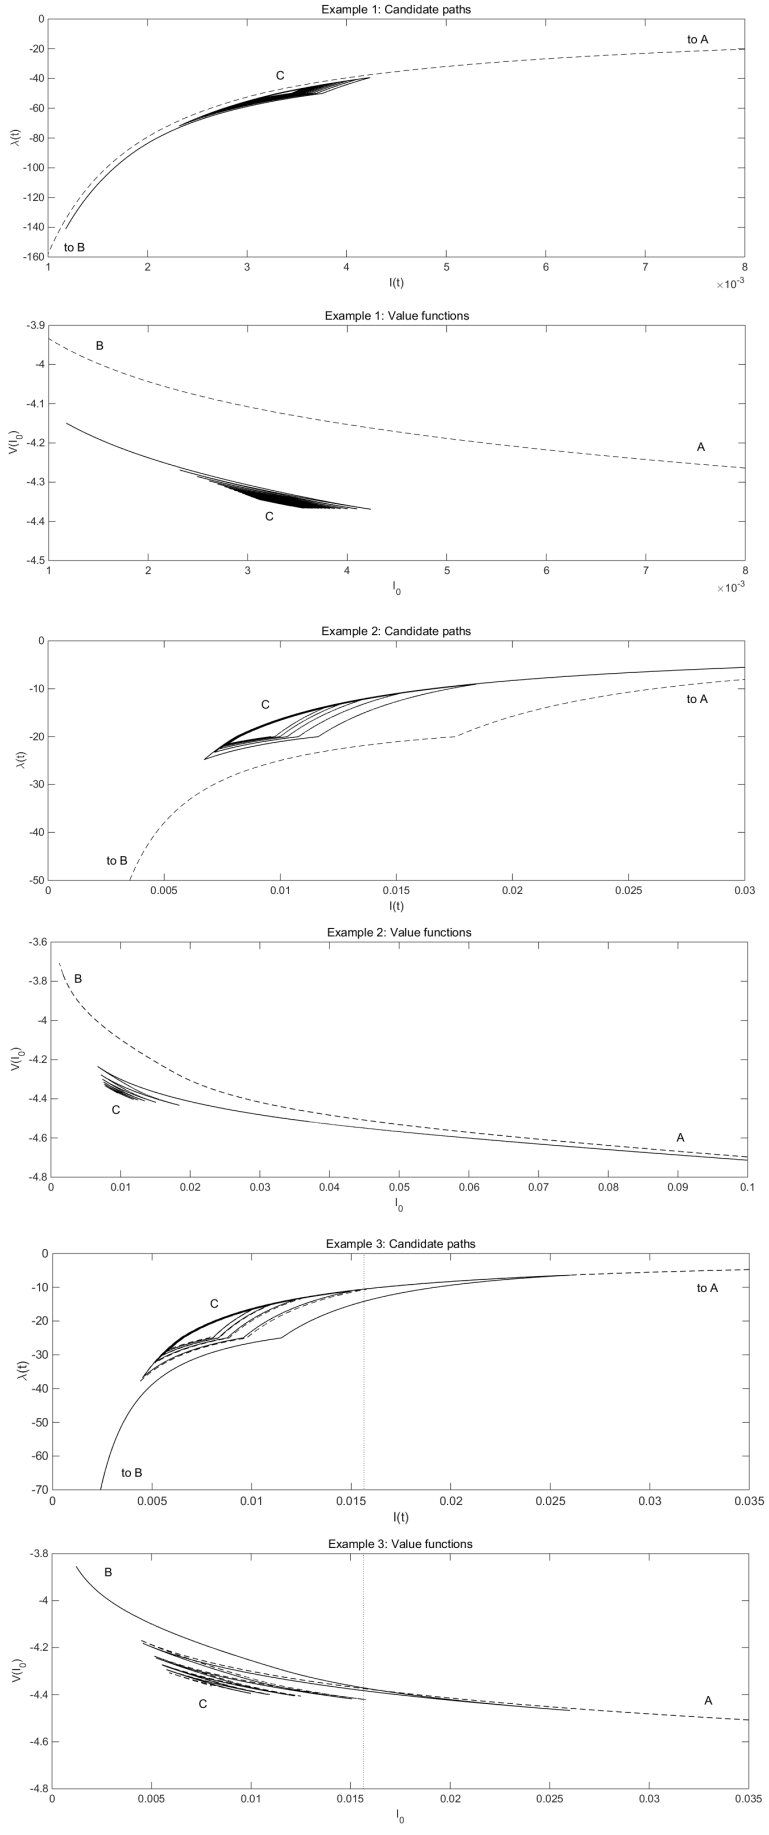

Figure 1: Candidate paths and value functions.

B intersect.

To fully appreciate the pitfalls of focusing on steady state welfare levels, consider the following experiment, based on the parameters of Example 1. We know that in this case, it is always optimal to steer the system to steady state A. Let us compare the discounted steady state welfare levels. Starting at steady state B and staying there in perpetuity yields discounted welfare of $V_B = -4.217$, whereas starting at steady state A and staying there in perpetuity yields $V_A = -4.517$. In other words, it is clearly better to be at B and stay there than it is to be at A and stay there. A simple-minded focus on steady state welfare levels would thus prescribe leaving steady state A and go instead to steady state B, thereby yielding higher welfare in the long run. In fact, this prescription is wrong, because it fails to account for the net loss incurred in the transition from A to B. Indeed, in this example, it is optimal to leave steady state B and going instead to steady state A. This path would yield discounted aggregate welfare equal to -4.023, which is superior to staying at steady state B. In short, properly accounting for welfare along optimal paths may yield the exact opposite policy prescription than pure steady state comparisons would suggest.

6. Equilibria under Decentralized Decision Making

In conducting our analysis so far, we have taken the perspective of a benevolent social planner who can dictate policies and does not have to consider the incentives of the individuals in the population. This raises the important question of the possible decentralization of optimal policy. While the solution to the social planner's problem yields important new insights and is an important benchmark, it is also useful to understand the equilibrium outcomes under decentralized decision making.

In particular, we wish to understand (i) how individual decision makers' choices and aggregate disease dynamics interact, (ii) the extent to which the equilibrium outcomes under decentralized decision making coincide with the solution under centralized decision making, and (iii) how to align private and public incentives through policy intervention. To make progress in answering these three questions, we next analyze the problem faced by an individual $i \in \mathcal{P}$ who is too small to influence aggregate infection dynamics. Such an individual will maximize discounted expected utility, taking the future trajectory of aggregate disease prevalence as *given.*

We will analyze two settings, one with *uncontrolled* decentralized decision making and one with *controlled* decentralized decision making. In the former case, individuals make decisions without any intervention on the planner's part. In the latter case, the planner will seek to modify individuals' equilibrium behavior through two separate incentive schemes, namely *prevention and treatment subsidies* and an *infection tax/health subsidy.*

The individual's problem can be written in the following form (see Appendix C for details), which is amenable to standard optimal control techniques:

$$\max_{\tau_i(t),\pi_i(t)\in[0,1]} \int_0^\infty e^{-\rho t}\left[-q_i(t)\left[\omega+\tau_i(t)c_T\right]-(1-q_i(t))\pi_i(t)c_P\right]dt \tag{17}$$

$$s.t. \quad \dot{q}_i(t) = (1-\pi_i(t))\beta I(t)(1-q_i(t)) - (\gamma+\alpha\tau_i(t))q_i(t), \quad q_i(0)\in\{0,1\} \tag{18}$$

The interpretation of the law of motion (18) is as follows. The first term on the right-hand side is the probability of infection (per unit of time) for a susceptible individual with protection intensity $\pi_i(t)$. The second term is the probability of recovery (per unit

of time), for an infected individual with treatment intensity $\tau_i(t)$. The instruments $\pi_i(t)$ and $\pi_i(t)$ can be interpreted as randomization probabilities across the extreme values 0 and 1. Note that for the individual, the entire path of the aggregate variable $I(t)$ is taken as given.

To derive the optimal policy for individual $i \in \mathcal{P}$ at time $t \geq 0$, we proceed as follows. The individual's current-value Hamiltonian is given by

$$\begin{aligned} H_i^D \equiv & -q_i(t)\left[\omega + \tau_i(t)c_T\right] - (1 - q_i(t))\pi_i(t)c_P \\ & +\mu_i(t)\left[(1 - q_i(t))(1 - \pi_i(t))\beta I(t) - q_i(t)(\gamma + \tau_i(t)\alpha)\right] \end{aligned} \tag{19}$$

where $\mu_i(t)$ is the costate variable in the individual's optimization problem. Note that there is an important difference between the planner's problem and that solved by the individual. The planner's Hamiltonian H^C given in (7) is quadratic in the relevant state variable (I), whereas the Hamiltonian H_i^D is linear in the relevant state variable (q_i).

The evolution of the individual's costate variable is given by the differential equation

$$\dot{\mu}_i(t) = \mu_i(t)\left[\rho + \gamma + \alpha\tau_i(t) + (1 - \pi_i(t))\beta I(t)\right] + \left[\omega + \tau_i(t)c_T - \pi_i(t)c_P\right] \tag{20}$$

For the optimal path, the policy instruments $(\tau_i(t), \pi_i(t))$ must maximize the above current-value Hamiltonian.

For an infected individual, the optimum level of treatment must satisfy the following inequalities:

$$\tau_i(t) = 0 \quad if \quad \alpha\mu_i(t) > -c_T \tag{21}$$

$$\tau_i(t) \in [0,1] \quad if \quad \alpha\mu_i(t) = -c_T \tag{22}$$

$$\tau_i(t) = 1 \quad if \quad \alpha\mu_i(t) < -c_T \tag{23}$$

For a susceptible individual, the optimum level of protection must satisfy the following inequalities:

$$\pi_i(t) = 0 \quad if \quad \beta\mu_i(t)I(t) > -c_P \tag{24}$$

$$\pi_i(t) \in [0,1] \quad if \quad \beta\mu_i(t)I(t) = -c_P \tag{25}$$

$$\pi_i(t) = 1 \quad if \quad \beta\mu_i(t)I(t) < -c_P \tag{26}$$

In addition, it must also be the case that

$$\lim_{t\to\infty} e^{-\rho t} H_i^D(t) = 0 \tag{27}$$

This transversality condition is always satisfied if $\mu_i(t)$ approaches a finite limit.

For any continuous trajectory of aggregate infection $I(t)$, if the equations and conditions (18) to (27) are satisfied, the resulting path satisfies the Arrow conditions for an optimum. This is shown in Appendix F. This is true also in the controlled decentralized setting to be analyzed below.

By symmetry, all individuals who are infected at a given point in time $t \geq 0$, will choose the same level of treatment. Likewise, all individuals who are susceptible at a point in time $t \geq 0$, will choose the same level of prevention. Thus, $\tau_i(t) = \tau(t)$ for all $i \in \mathcal{I}(t)$ and $\pi_i(t) = \pi(t)$ for all $i \in \mathcal{S}(t)$. Both of these variables may be interior,

indicating a mixed strategy. The shadow price of infection (i.e. the costate variable) will also be the same for all individuals and hence $\mu_i(t) = \mu(t)$ for all $i \in \mathcal{P}$. We can therefore omit the subscripts on the costate and control variables when considering the aggregate levels of the instruments in what follows.

Consistency requires that the average infection probability equals aggregate disease prevalence $I(t)$. Hence, averaging equation (18) across all individuals yields

$$\dot{I}(t) = (1 - \pi(t))\beta I(t)(1 - I(t)) - (\gamma + \alpha\tau(t))\, I(t) \tag{28}$$

The above equation is the same as equation (6) in the centralized setting, but with aggregate treatment and protection replaced by the average individual choices.

As in the planner's problem, various parameter restrictions must be satisfied for any particular type of steady state to be feasible. In general, these feasibility restrictions differ between centralized and decentralized decision making. As a consequence, when parameters are altered, there can be both quantitative and qualitative differences between the centralized and decentralized settings. A particular type of steady state that is feasible under the former setting may not be feasible under the latter. This was also pointed out by Goldman and Lightwood (2002).

In Appendix D, we list the steady state values for points A^*, B^*, A_0^*, B_0^*, C_0^* and compare these to the corresponding values from the solution to the planner's problem. A steady state of the form C is never feasible in the decentralized problem.

Before detailing the features of the decentralized equilibria, we formally state the following existence result:

Theorem 5. *Under decentralized decision making, an equilibrium path* $(I^*(t), \tau_i^*(t), \pi_i^*(t))$ *exists if at least one of the fixed points* A^*, B^*, A_0^*, B_0^* *is feasible.*

Proof: Given that aggregate behavior takes the system to a steady state, Arrow's sufficiency theorem establishes that it is individually optimal for any individual to also go to the steady state and to stay there ■

In comparing the possible outcomes under centralized and decentralized decision making, we first note that the set of possible steady states do not coincide. In particular, the equivalent of solution C in the planner's problem, in which $\tau(t) \in (0,1)$ and $\pi(t) \in (0,1)$, does not exist under decentralized decision making. While there are values $\tau_i(t) \in (0,1)$ and $\pi_i(t) \in (0,1)$ for which the individual's problem is at a rest point, i.e. where $\dot{q}_i(t) = 0$, this point does not correspond to a fixed point of the aggregate system, i.e. where $\dot{I}(t) = 0$.

Next, we compare the level of disease prevalence in various steady states. Direct inspection confirms that steady state prevalence levels in the decentralized case are at least as high as the corresponding levels under centralized decision making. In fact, we have that

$$I_A < I_{A^*}, \quad I_B < I_{B^*}, \quad I_{A_0} = I_{A_0^*}, \quad I_{B_0} = I_{B_0^*} \tag{29}$$

An important consequence of decentralized decision making is that, due to the availability of treatment, the path actually chosen in equilibrium may be indeterminate and depend on expectations about future decisions. Specifically, the present model allows for the possibility that there is a range of initial conditions for which there are multiple perfect foresight equilibrium paths. Along each of these, each individual maximizes discounted expected utility given the behavior of others. Furthermore, these paths may go

to qualitatively distinct steady states. This phenomenon is treated in detail in Toxvaerd (2009a) in the context of a treatment-only model.

Although it is convenient to classify the steady states in terms of high and low infection steady states under centralized versus decentralized decision making, it may also be somewhat misleading when comparing what can happen under the two different scenarios. For example, it may be tempting to focus, say, on comparing the high infection steady that the planner chooses to the high infection steady state that can materialize in equilibrium. But this masks the fact that for a given choice of model parameters, one type of steady state under decentralized decision making may be simultaneously feasible with another type of steady state under decentralized decision making. For example, a no-treatment centralized steady state could coexist with a full-treatment decentralized steady state. A complete analysis would therefore compare all feasible steady states across the two scenarios.

We have only taken a few steps in this interesting direction. It is straightforward to show that if the high infection steady state A is feasible, then the low infection steady state B^* is not. But the same is not true for steady state pairs (A, B_0^*), (A_0, B_0^*) or (A_0, B^*). Thus in principle, it could be the case that under decentralized decision making, the system gets stuck in a low infection steady state while the planner would prefer a path that ends in a steady state with a higher level of infection. While this would be a singularly interesting and significant finding, we have carried out extensive simulations of the model and have not been able to identify parameter constellations under which this result is true.[13] In every instance, the socially optimal path terminates at a level of infection which is less than or equal to the minimum steady state infection level achievable under decentralized decision making.

We will now explain what leads to the differences between the outcomes under centralized decision making and the equilibria under decentralized decision making. Under decentralized decision making, the incentives of the individual are different from those of the central planner. As a result, they assign different shadow prices to infection. On an optimum path, the planner's shadow price is state dependent and is equal to $\lambda(I)$. At the same level of aggregate infection, the individual's shadow price for its own infection is equal to $\mu(I)$. The gap between these shadow prices is $z(I) \equiv \lambda(I) - \mu(I) \leq 0$. To elucidate the nature of this gap, we will decompose $z(I)$ as follows. First, there is a pure externality effect $x(I)$. This effect stems from the fact that in deciding what levels of treatment or protection to choose, each individual ignores the impact of its actions on the well-being of other individuals. Second, there is a risk effect $y(I)$. Starting from any given level of aggregate infection I, the future time profile of I is different under centralized and decentralized decision making. As a result, the individual faces a different future time profile of infection risk under the two scenarios. The risk effect encapsulates the influence

[13] In these simulations, we have sought to test whether a feasible steady state B_0^* has a lower prevalence level than would be eventually achieved along the socially optimal path. We assumed that $\beta > \gamma + \alpha$ (to ensure that treatment alone does not eliminate the disease) and normalized by assuming that $\beta = \omega = 1$. We looked at approximately 2.9 million parameter combinations. There were some thousands of combinations for which B_0^* and some high prevalence steady state (either A or A_0) coexisted. However, in none of these cases was it optimal to go to A or A_0. In these cases the optimal path always takes the system to either B or B_0, where the prevalence level is lower than or equal to the prevalence level at steady state B_0^*. We repeated this exercise for B^*. The simulations indicate that it is never optimal to go from B^* to A_0, but optimal to go instead to B, at which prevalence is lower than at B^*.

of this difference on the individual's shadow price.

To quantify these two effects, we will introduce a so-called "maverick" individual. Suppose that under central direction, all individuals, except for one numerically insignificant maverick, behave in the socially optimal fashion. This ensures that aggregate infection $I(t)$ will follow the socially optimal path. The maverick, in contrast, behaves in a purely selfish fashion and maximizes its own personal discounted expected utility while facing the socially optimal time profile of future infection risk. The privately optimal solution for the maverick will then satisfy the system of equations and conditions (18) to (27) that characterize decentralized dynamics, on the assumption that aggregate infection $I(t)$ follows the socially optimal trajectory. Let $\eta(t)$ denote the shadow price of infection for the maverick. This can be written in the state dependent form $\eta(I)$.

Consider the difference between the optimization problem of the maverick and that of the central planner. By assumption, both the planner and the maverick face the same path of aggregate infection (namely the socially optimal one) and hence any differences in valuation stem from the fact that the maverick ignores the impact of its actions on the well-being of others. Thus the difference is a pure externality effect. Similarly, consider the difference between the optimization problem of the maverick and that of an individual on the decentralized equilibrium path of aggregate infection. Both individuals face the same costs of prevention and treatment and both are indifferent to the harm their actions may impose on others, but they face different paths of future aggregate infection, namely the socially optimal one versus the decentralized equilibrium path. This difference in valuation captures the added risk that comes about because the individuals concerned face different future paths of aggregate infection.

More formally, compare the following three paths, expressing the costate variables as functions of the current aggregate infection rate I: (i) the decentralized equilibrium path facing an individual when no subsidies are offered (the shadow price of an individual is then equal to $\mu(I)$); (ii) the centralized optimal path (along which the shadow price is equal to $\lambda(I)$); (iii) the centralized optimal path facing the maverick (along which the shadow price is equal to by $\eta(I)$). We then have the following important result:

Theorem 6. *The shadow price gap $z(I)$ can be decomposed into an externality effect and a risk effect, such that*

$$\underbrace{\lambda(I)-\mu(I)}_{z(t)} = \underbrace{[\lambda(I)-\eta(I)]}_{x(t)} + \underbrace{[\eta(I)-\mu(I)]}_{y(t)} \tag{30}$$

where

$$\dot{\lambda}(t) = \lambda(t)\left[\rho+\gamma+\alpha\tau(t)+(1-\pi(t))\beta(2I(t)-1)\right]+\left[\omega+\tau(t)c_T-\pi(t)c_P\right] \tag{31}$$

$$\dot{\mu}(t) = \mu(t)\left[\rho+\gamma+\alpha\tau(t)+(1-\pi(t))\beta I(t)\right]+\left[\omega+\tau(t)c_T-\pi(t)c_P\right] \tag{32}$$

$$\dot{\eta}(t) = \eta(t)\left[\rho+\gamma+\alpha\tau(t)+(1-\pi(t))\beta I(t)\right]+\left[\omega+\tau(t)c_T-\pi(t)c_P\right] \tag{33}$$

Note that in the equations for $\lambda(t)$ and $\mu(t)$, aggregate infection $I(t)$ follows the socially optimal path. In the equation for $\eta(t)$, aggregate infection follows a decentralized equilibrium trajectory. The values of the instruments $\pi(t)$ and $\tau(t)$ differ across equations.
Proof: See Appendix E ■

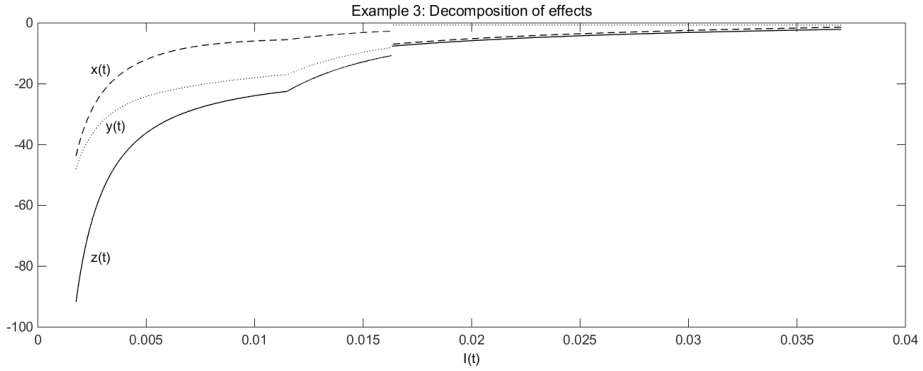

Figure 2: Decomposition of effects

It is worth noting that $\eta(1) = \lambda(1)$, i.e. the shadow values coincide when the entire population is infected. This follows from the straightforward observation that there are no possible externalities when the entire population is already infected.[14] In Figure 2, we plot the two effects $x(I)$ and $y(I)$ for Example 3. As expected, both the pure externality effect and the risk effect are decreasing in absolute value as prevalence increases, tending to zero as the entire population becomes infected.

The decomposition of effects is interesting in its own right, but will also become important when designing optimal corrective interventions, a task we turn to next.

7. Aligning Public and Private Incentives

In this section, we will consider two distinct ways in which the planner can induce the individuals to make socially optimal choices. The first policy intervention relies on a dynamic, state dependent corrective subsidy scheme, in which the planner offers subsidies to treatment and prevention. The second policy intervention relies on dynamic, state dependent corrective subsidies to healthy individuals (or equivalently, taxes on infected individuals). The former policy influences the *flow* between health states because prevention and treatment decisions influence the rates of transition between health states. The latter policy relies on taxes on the *stock* of health. While both schemes induce individuals to make the socially optimal choices, they differ in important ways. Furthermore, there are practical reasons why a planner may wish to use one type of subsidy scheme rather than the other. We will discuss these issues in more detail, once we have formally characterized the two schemes.

In seeking to align public and private incentives, it is natural to resort to Pigouvian-style schemes that rely on the magnitude of non-internalized external effects. There is a large literature in environmental economics that makes use of such state dependent incentive schemes (see e.g. Tahvonen and Kuuluvainen, 1993, Farzin, 1996, Aronsson et

[14] In a similar fashion, note that for the special case $\gamma = \tau(t) = 0$ and $I(t) = 1$, the individuals and the planner have the same shadow prices, i.e. $\mu(t) = \lambda(t)$. This case is relevant when prevention is too costly to be used in steady state and when spontaneous recovery is not possible (i.e. only treatment will induce recovery). In this case, one feasible steady state involves no treatment at all, which leads to the entire population being infected. In this steady state, the planner and the individuals face the same path of infection and since there are no susceptibles on which an infected individual can have external effects, there are no externalities. Therefore, the individuals and the planner value infection equally much.

al., 1998 and Rubio and Escriche, 2001). It turns out that deriving an optimal scheme in our setting is a more subtle endeavour. This is because the incentive scheme needs to be designed to simultaneously correct for both the pure externality effect and the risk effect. As we will see, this means that the optimal corrective incentive schemes are not straightforward Pigouvian taxes.

7.1. Subsidies to Prevention and Treatment. Suppose that subsidies are given to prevention and treatment at rates $s_P(t)$ and $s_T(t)$, respectively. The following proposition specifies how these subsidies should be chosen to induce socially optimal behavior.

Proposition 7. *The following subsidy schedules implement the first-best policy via decentralized decision making:*

$$s_P(t) \equiv \beta I(t)[\phi(t) - \lambda(t)] \geq 0 \tag{34}$$

$$s_T(t) \equiv \alpha[\phi(t) - \lambda(t)] \geq 0 \tag{35}$$

where

$$\dot{\phi}(t) = \phi(t)\left[\rho + \gamma + \beta I(t)\right] + \omega + \tau(t)\left[c_T + \alpha\lambda(t)\right] - \pi(t)\left[c_P + \beta I(t)\lambda(t)\right] \tag{36}$$

Note that $\phi(t)$ is evaluated along the socially optimal path for $\tau(t)$, $\pi(t)$, $\lambda(t)$ and $I(t)$.

Proof: See Appendix E ■

The functional form of the two subsidies have a very nice interpretation, namely that the subsidy rates equal the rates at which prevention and treatment abate the social damage from infection. The preceding results have a surprisingly simple corollary, which we state next:

Corollary 8. *The optimal subsidy ratio $s_P(t)/s_T(t)$ is directly proportional to disease prevalence, i.e.*

$$\frac{s_P(t)}{s_T(t)} = \left(\frac{\beta}{\alpha}\right) I(t)$$

This corollary implies that in a high-infection steady state, protection is subsidized at a higher relative rate than in a low-infection steady state. This result should be seen in the context of the analysis of Gersovitz and Hammer (2004), who arrive at a different result, but under the twin assumption that there is a unique steady state and that it is fully interior, features that generically do not hold in our setting. Gersovitz and Hammer (2004) report that the subsidies should be at equal rates, irrespective of prevalence. They state their results in terms of ad valorem taxes, while we state ours in terms of excise taxes (i.e. as per unit taxes). It is straightforward to verify that in our setting, the optimal ad valorem tax rates for prevention and treatment are given by

$$v_P \equiv \frac{-s_P(t)}{c_P} \tag{37}$$

$$v_T \equiv \frac{-s_T(t)}{c_T} \tag{38}$$

Upon substitution of the optimal subsidies, it follows immediately that $v_P > v_T$ if $I(t) < I_C$ and $v_P < v_T$ if $I(t) > I_C$. In fact, they coincide *only* in the unstable and suboptimal

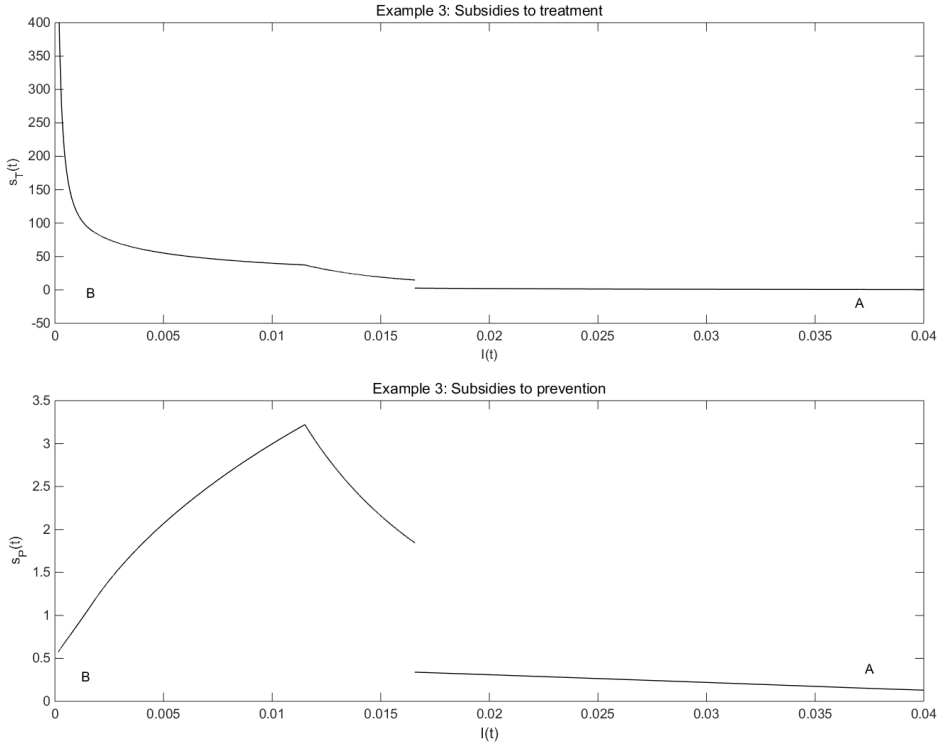

Figure 3: Subsidies to treatment and prevention.

steady state C. In contrast, Gersovitz and Hammer (2004) report that the optimal ad valorem taxes are equal at all times.

In Figure 3, we display the optimal subsidies to prevention and treatment using the parameter values of Example 3. We start by considering the subsidies to treatment, displayed in the first panel. Other than at the Skiba point I_S, the scheme $s_T(t)$ is a continuous function of disease prevalence throughout. Furthermore, it has a kink at the point where optimal treatment switches between zero and one. Note that the schedule is continuous at both steady states. In accordance with the observation that treatment becomes less effective as prevalence increases, the optimal subsidy to treatment decreases as I increases.

Turning to the subsidy for protection, which is displayed in the second panel, the scheme $s_P(t)$ is also continuous throughout, other than at the Skiba point I_S. In particular, it is continuous at the steady states I_A and I_B, even though optimal prevention is discontinuous at these points. However, this subsidy is not monotone in prevalence over its entire domain. We distinguish two cases, depending on whether or not there is treatment on the relevant segment of the optimal path. To the right of the kink in the graphs of $s_P(t)$ and $s_T(t)$ there is no treatment, and in this case, the subsidy to prevention is a decreasing function of disease prevalence. The reason is simply that as prevalence increases, the private incentive to prevent infection increases sufficiently rapidly to induce optimal behavior by individuals, despite the decreasing subsidy. To the left of the kink, full treatment is optimal and induced through a very high subsidy to treatment. In this region, treatment and prevention are substitutes and prevention becomes less and

less useful as a tool for suppressing infection. For that reason, as prevalence decreases, the subsidies to prevention also decrease, tending to zero as the disease approaches the eradication point.

Last, for sufficiently low levels of aggregate infection, the subsidy to treatment is significantly higher than the cost of treatment c_T. In fact, with the present parameterization, this is true for any $I(t) < I_S$. As a consequence, if the planner wants to induce the individuals to reach the low infection steady state I_B from above, the required subsidy is so high that the individual shadow price of infection $\phi(t) > 0$. This is a notable finding, because it means that with optimally chosen subsidies, individuals will privately value becoming infected! This is in contrast to the uncontrolled decentralized equilibrium, in which it is always the case that individuals dislike becoming infected, as $\mu(t) \leq 0$.

At first blush, this might suggest that in the presence of optimal subsidies, individuals might decide not to protect themselves when susceptible. Yet, optimal subsidies ensure that individuals behave in a socially optimal fashion and choose full protection between the Skiba point I_S and the steady state I_B. This may seem strange, since it implies that the planner offers subsidies to treatment that are so high that individuals want to become infected and at the same time offers subsidies in order for the individuals to protect themselves. To make sense of this, it must be kept in mind that if the taxation required to finance the subsidies is non-distortionary, then social welfare is not adversely affected by the fact that subsidies are high. All that matters is that the combined effect of the subsidy schemes $s_P(t)$ and $s_T(t)$ is to ensure that individuals make socially optimal decisions. And the proposed subsidies achieve this.

As this subsidy scheme is not revenue neutral, it must be funded by non-distortionary lump-sum transfers from the general population.

7.2. Taxes and Subsidies on the Individual's Health Status. As an alternative to interventions linked directly to behavior, such as subsidies for prevention and treatment, the social planner could instead seek to influence behavior indirectly by altering the costs and benefits associated with any particular health status. For example, the planner could impose a uniform *infection tax* on all infected individuals or else give a uniform *health bonus* to all uninfected individuals. These two interventions are equivalent if they are offset by a revenue neutral, non-distortionary change in general taxation. We will therefore concentrate on the case of a uniform tax on the infected population.

Theorem 9. *The following infection tax schedule implements the first-best policy via decentralized decision making:*

$$T(t) \equiv -\lambda(t)(1-\pi(t))\beta(1-I(t)) \geq 0 \tag{39}$$

Note that $\lambda(t), \tau(t), \pi(t)$ and $I(t)$ are evaluated along the socially optimal path.

Proof: See Appendix E ■

This result has an intuitive interpretation. The quantity $(1-\pi(t))\beta(1-I(t))$ is the probability that an infected individual will pass on its infection to someone else and $-\lambda(t)$ is the resulting social damage. Thus, $-\lambda(t)(1-\pi(t))\beta(1-I(t))$ is the expected damage per unit of time that an infected person will cause by infecting other members of society. This is the total effect that is internalized by imposing the tax $T(t)$ on infected individuals.

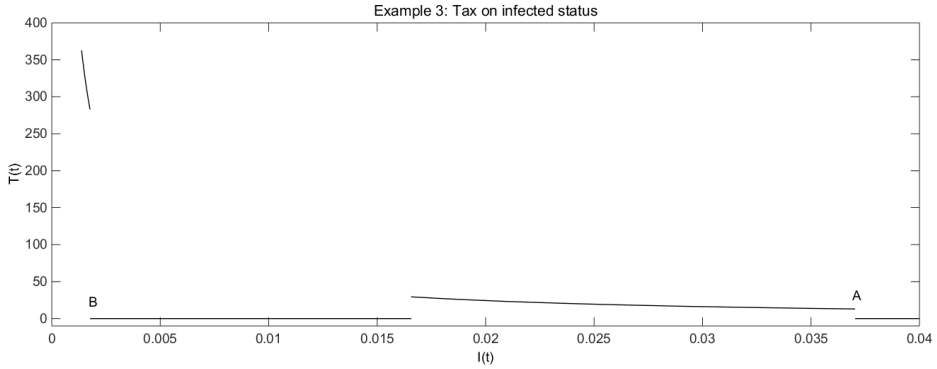

Figure 4: Tax on health status

As noted earlier, when $I(t) = 1$, there are no externalities and thus the optimal tax $T(t) = 0$. More interestingly, the optimal infection tax is positive only when it is optimal to induce less than full prevention (i.e. when $\pi(t) < 1$). In practice, this is the case whenever prevalence is above the steady state that the planner is aiming for. Thus, while approaching the desired steady state from above, no tax is levied on infection. However, once the steady state is reached, a positive tax on infection is imposed. The upshot of this is that the promise of future taxes (at the appropriate level) is sufficient to induce socially optimal behavior on the part of the individuals.

Formally, we have that the optimal tax on infection, as a function of disease prevalence, has the property that

$$T(t) = 0 \quad if \quad I(t) \in [I_B, I_S] \cup [I_A, 1] \tag{40}$$

$$T(t) > 0 \quad if \quad I(t) \notin [I_B, I_S] \cup [I_A, 1] \tag{41}$$

In contrast to the case of subsidies for treatment and prevention, intervention through an infection tax (or a health subsidy) perfectly equalizes the individuals' and the planner's shadow price of infection. Given the same shadow price, and facing the same costs, the individuals and the planner will therefore choose the same levels of prevention and treatment.

In Figure 4, we display the optimal tax for the parameterization of Example 3. First, note that the optimal tax is continuous for all but three points, namely the steady states I_A and I_B and the Skiba point I_S. The discontinuity at the Skiba point stems from the fact that this level of disease prevalence determines the optimality of reaching for steady state I_A versus I_B. Second, along an optimal path, the tax on infection is (weakly) decreasing in disease prevalence. To see this, recall that for $I_0 > I_S$, the optimal path ends at steady state I_A while for $I_0 < I_S$, the optimal path ends at steady state I_B. This feature stems from the fact that the total effects decrease with disease prevalence and thus the optimal tax on infection accordingly decreases.

7.3. Choosing Between Policies. We have shown that the planner can effectively shape individual behavior by appropriately chosen taxes or subsidies. Having said that, the context may dictate that one scheme be chosen rather than the other. For example,

subsidies to prevention and treatment may be difficult to implement in some contexts. While it is straightforward to make condoms freely available (an implicit subsidy), it is not obvious how one would go about subsidizing their actual use (rather than their acquisition). Treatment seems easier to subsidize in practice (as it can be administered through the health practitioner that supplies the medicine), but optimal behavior cannot be induced through treatment subsidies alone. On the other hand, imposing taxes on infected individuals or awarding a health bonus to uninfected individuals may be difficult to implement for political reasons, as they may be seen as unduly harsh treatment of the already vulnerable.

We should emphasize again that the schemes we have proposed are not Pigouvian in the usual sense. Traditionally, Pigouvian taxes are set in order to correct for the fact that individuals do not care about the damage they may cause others. As we have shown above, such external effects are only partly responsible for the fact that individuals make socially undesirable decisions. The fact that individuals are negligible relative to the size of the whole population, gives rise to the additional risk effect. Our proposed incentive schemes are designed to correct for both these effects.[15]

Last, we note that the decomposition results and the method for designing corrective schemes apply to a wider range of models and not just to the SIS setting considered in this paper.

An important issue when considering the practical implementation of policies, is the complexity of the incentive schemes offered to individuals. The corrective incentive schedules we have derived are complicated objects, because the subsidies on offer at any given point in time are relatively complicated functions of aggregate disease prevalence $I(t)$. Furthermore, in the case of subsidies to prevention and treatment, these must be determined jointly (i.e. one cannot simply hold one subsidy constant and then compensate by setting the other subsidy optimally).

In light of the complicated nature of the socially optimal incentive schemes considered here, it may be desirable to look for simplified optimal schemes, i.e. incentive schemes that induce the socially optimal behavior but which are simpler than the ones derived above. For example, the fact that the scheme that we have derived offers positive subsidies to prevention when the planner wants to induce zero prevention, may suggest that a simplified scheme exists in which these subsidies are set to zero. But one must carefully determine the effects of such an alternative scheme, as it may have unintended consequences. To see this, consider the case in which the unsubsidized individual chooses $\pi^*(t) = 0$. Suppose that with subsidy $s_P(t)$, the individual chooses the socially optimal prevention level $\pi(t) = 1$. Suppose also that $\beta\phi(t)I(t) < -(c_P - s_P(t))$. Then for sufficiently small function $\delta(t) > 0$ and with the same value of $\phi(t)$, it is also the case that $\beta\phi(t)I(t) < -(c_P - \widetilde{s}_P(t))$ where $\widetilde{s}_P(t) \equiv s_P(t) - \delta(t)$. This would suggest that we can replace $s_P(t)$ by $\widetilde{s}_P(t)$ without affecting the decision of the individual. However, this argument ignores the fact that by changing the function $s_P(t)$ over a finite interval of time, we alter the trajectory of $\phi(t)$ via the state equation (which includes $s_P(t) > 0$ over

[15] The result that a strictly Pigouvian tax does not implement the first best, is similar in nature to that found by Rubio and Escriche (2001). They find that the optimal Pigouvian tax (in a production setting with market power) is *neutral* in the sense that it only corrects for external effects and not for other inefficiencies in the economy. In this sense, our subsidy scheme takes into account the neutrality of a purely Pigouvian tax by also incentivizing the individuals to correct for the additional risk effect, as explained by our decomposition result.

this interval of time). In turn, this will alter the trajectories of the switching equations and hence the trajectories of the control variables for the individual. Conversely, suppose that $\pi^*(t) = 0$ and $\pi(t) = 0$. Then any subsidy $\widetilde{s}_P(t)$ which satisfies the inequality $\beta\phi(t)I(t) > -(c_P - \widetilde{s}_P(t))$ with our value of $\phi(t)$, will cause the individual to choose zero protection. Moreover, the trajectory of $\phi(t)$ and all the other variables will be unaffected, since the subsidy will not actually be paid. The preceding arguments suggest that our proposed subsidy for protection is not unique over its entire range. But importantly, they also show that the simplified scheme must be designed very carefully indeed, lest one unintentionally induce socially suboptimal behavior.

Short of schemes that induce socially optimal behavior, it is of considerable practical interest to determine second-best corrective policies. In the terminology of Arrow and Kurz (1969), the optimal policy may not be *controllable* with simple schemes, i.e. the planner may be unable to decentralize the optimal policy via schemes that are not complicated functions of the state variable. Having said that, it quickly becomes clear that it is not a simple matter to characterize the second-best subsidy scheme either. A major obstacle is the presence of multiple feasible steady states and the possibility that ill-chosen subsidies may tip the scales in the wrong direction. To see this, consider a fixed (non-state dependent) subsidy that reduces the individual's costs of prevention and/or treatment. As our analysis has made clear, it not always socially optimal to have prevention and/or treatment, and thus the subsidies may lead individuals to over or under demand prevention and/or treatment, depending on the aggregate level of disease prevalence. This is particularly relevant in Regime III, i.e. the case with a Skiba point in which small changes to initial conditions can change the end point of an optimal path. In this case, a carelessly chosen subsidy scheme can conceivably propel the system towards the wrong type of steady state (i.e. a low versus a high infection steady state). This possibility echoes the findings of Cohen et al. (2011), who suggest that subsidies to malaria treatment in Kenya may have lead to significant amounts of over-treatment. This type of finding is significant, because it is generally the case that subsidies are offered in order to correct for a perceived *under* demand of treatment, and the subsidy therefore replaces one distortion with another without necessarily offsetting it. Farzin (1996), Tahvonen and Kuuluvainen (1993), Aronsson et al. (1998) and Rubio and Escriche (2001) consider the welfare effects of non-optimal incentive schemes in different contexts. While the particular application naturally determines the nature of the welfare loss due to non-optimal schemes, such schemes typically modify both transition dynamics and the resulting steady state levels of the state variables, but do so quantitatively. In our setting, the effects can be somewhat more dramatic, as the scheme may drive the system towards the wrong *kind* of steady state and thus have qualitative effects.

8. Conclusion

In this paper, we have analyzed the economic control of an SIS type infectious disease via prevention and treatment. We have conducted our analysis under three different scenarios. First, we analyzed centralized decision making and found that while prevention and treatment can both bring down infection, they work in fundamentally different ways. Prevention is shown to have stabilizing negative feedback effects and treatment is shown to have destabilizing positive feedback effects. These create the potential for multiple steady states and history dependence. Second, we analyzed uncontrolled decentralized decision making and found that equilibrium outcomes generally differ from socially optimal ones.

We showed that the discrepancies are due to two different effects: (i) an externality effect arising from the fact that individuals are indifferent to the well-being of others, and (ii) a risk effect arising from the fact the individual faces a higher risk of future infection in the decentralized environment than in the centralized environment, and hence a different expected payoff from current decisions. Finally, we analyzed controlled decentralized decision making and suggested two incentive schemes that decentralize the command optimum. In particular, we derived a scheme that subsidizes prevention and treatment and a scheme that taxes the infected (or offers a bonus to the healthy). We showed that these schemes are not simple Pigouvian schedules, but complicated state dependent schemes that correct for both the externality effect and the risk effect, reflecting at each point how these effects vary across the stages of the epidemic.

Our analysis is not simply an abstract exercise, but one that has concrete, practical relevance to the formulation of policy. A case in point is the 2009 outbreak of swine flu. While not an SIS type disease, the thinking at the time suggested that there was a "treatment phase" and a "prevention phase" and that the timing of these were functions of disease prevalence.[16] Our analysis would suggest that such a rigid separation is suboptimal.

We would like to mention one omission from the analysis and one avenue for further research. Although policy formulation has taken a prominent role in our analysis, there are some issues that we have omitted due to space considerations. These are policies that work by changing the basic parameters of the model, such as the infectiousness of the disease or the effectiveness of treatment in inducing recovery. We have made some initial analysis along these lines, which is available in the online appendix. The issues involved are non-trivial and at times counter-intuitive, as can be seen in Toxvaerd (2010) in a prevention-only model.

Another issue is the extension of our analysis to the broader class known as susceptible-infected-recovered-susceptible (or SIRS) infections. In this class, recovered individuals become immune to further infection for some period of time before becoming susceptible again. It is clear that when natural immunity is lost at a sufficiently high rate, the SIRS model shares qualitative features with the SIS model and so many if not most of our conclusions remain valid. But when immunity is lost only slowly, then the model is closer to a susceptible-infected-removed model (or SIR). This model is difficult to analyze formally as there are no closed form solutions even in the non-economic version. In such a model, herd immunity is a possibility and so a planner might decide (depending on costs and the biological parameters of the model) to forego any interventions to manage the disease and simply let it run its course to eradication. Such an analysis will rely heavily on simulations, but seems worthwhile.

[16] See *Swine Flu: From Containment to Treatment*, UK Department of Health (2009).

A. Steady State Values for Centralized Setting

The different steady states are given as follows:

Solution A: This case corresponds to $\tau(t) = 0$ and $\pi(t) \in (0,1)$. The steady state solution is then

$$I_A \equiv \frac{\rho c_P}{\beta(\omega - c_P)} \tag{42}$$

$$\lambda_A \equiv \frac{c_P - \omega}{\rho} \tag{43}$$

$$\pi_A \equiv \frac{c_P(\beta - \gamma + \rho) + \omega(\gamma - \beta)}{c_P(\beta + \rho) - \beta\omega} \tag{44}$$

$$\tau_A \equiv 0 \tag{45}$$

Solution B: This case corresponds to $\tau(t) = 1$ and $\pi(t) \in (0,1)$. The steady state solution is then

$$I_B \equiv \frac{\rho c_P}{\beta(c_T + \omega - c_P)} \tag{46}$$

$$\lambda_B \equiv \frac{c_P - \omega - c_T}{\rho} \tag{47}$$

$$\pi_B \equiv \frac{c_P(\beta - \gamma + \rho - \alpha) + (\omega + c_T)(\alpha + \gamma - \beta)}{c_P(\beta + \rho) - \beta(\omega + c_T)} \tag{48}$$

$$\tau_B \equiv 1 \tag{49}$$

Solution C: This case corresponds to $\tau(t) \in (0,1)$ and $\pi(t) \in (0,1)$. The steady state solution is then

$$I_C \equiv \frac{\alpha c_P}{\beta c_T} \tag{50}$$

$$\lambda_C \equiv \frac{-c_T}{\alpha} \tag{51}$$

$$\pi_C \equiv \frac{2\alpha c_P - \alpha\omega + c_T(\gamma + \rho - \beta)}{\alpha c_P - \beta c_T} \tag{52}$$

$$\tau_C \equiv \frac{\alpha c_P - \alpha\omega + \rho c_T}{\alpha c_T} \tag{53}$$

Solution A_0: This case corresponds to $\tau(t) = 0$ and $\pi(t) = 0$. The steady state solution is then

$$I_{A_0} \equiv \frac{\beta - \gamma}{\beta} \tag{54}$$

$$\lambda_{A_0} \equiv \frac{-\omega}{\beta - \gamma + \rho} \tag{55}$$

$$\pi_{A_0} \equiv 0 \tag{56}$$

$$\tau_{A_0} \equiv 0 \tag{57}$$

Solution B_0: This case corresponds to $\tau(t) = 1$ and $\pi(t) = 0$. The steady state solution

is then

$$I_{B_0} \equiv \frac{\beta - \gamma - \alpha}{\beta} \tag{58}$$

$$\lambda_{B_0} \equiv \frac{\omega + c_T}{\alpha - \beta + \gamma - \rho} \tag{59}$$

$$\pi_{B_0} \equiv 0 \tag{60}$$

$$\tau_{B_0} \equiv 1 \tag{61}$$

Solution C_0: This case corresponds to $\tau(t) \in (0,1)$ and $\pi(t) = 0$. The steady state solution is then

$$I_{C_0} \equiv \frac{\alpha\omega + c_T(\beta - \gamma - \rho)}{2\beta c_T} \tag{62}$$

$$\lambda_{C_0} \equiv \frac{-c_T}{\alpha} \tag{63}$$

$$\pi_{C_0} \equiv 0 \tag{64}$$

$$\tau_{C_0} \equiv \frac{-\alpha\omega + c_T(\beta - \gamma + \rho)}{2\alpha c_T} \tag{65}$$

Based on these values, some important observations follow:

Proposition 10. *(i) Steady states with positive treatment have lower disease prevalence than steady states with no treatment, i.e.* $I_A > I_B$ *and* $I_{A_0} > I_{B_0}$. *(ii) Steady states with positive prevention have lower disease prevalence than steady states with no prevention, i.e.* $I_A < I_{A_0}$ *and* $I_B < I_{B_0}$.

Proof: Part (i) follows from direct inspection. Part (ii) follows from the fact that the conditions that ensure that the no prevention steady state prevalence levels are higher than the positive prevention steady state prevalence levels, are exactly the opposite of the conditions that must hold for prevention to be zero in the no-prevention steady states ■

These results are not trivial, since prevention and treatment both work to reduce infection. It is therefore conceivable that the lack of one instrument is compensated for by an increase in the other instrument to the extent that prevalence ends up at a lower level than it otherwise would have been.

The next result follows from direct inspection of the relevant steady state prevention levels:

Proposition 11. *In the steady states with positive prevention, the no treatment steady state involves more prevention than the full treatment steady state, i.e.* $\pi_A > \pi_B$.

We can summarize the ranking of the steady state prevalence levels as follows:

$$I_B \leq \min\{I_A, I_{B_0}\} \leq \max\{I_A, I_{B_0}\} \leq I_{A_0}$$

The prevalence levels I_A and I_{B_0} are not unambiguously ranked.[17] But the Hamiltonian conditions for point B_0 ensure that $I_A \geq I_{B_0}$.

[17] It is easy to check that $I_A \geq I_{B_0}$ if and only if $c_P \geq \omega\left(\frac{\beta-\gamma-\alpha}{\beta-\gamma-\alpha+\rho}\right)$.

B. Non-Optimality of Maximal Prevention

In this appendix, we prove that an optimal path cannot involve eradication. The necessary conditions for an optimum are given by the Hamiltonian conditions, the laws of motion for the state and costate variables and the transversality condition

$$\lim_{t\to\infty}\left[e^{-\rho t}H^C(t)\right]=0 \tag{66}$$

The final condition comes from Michel (1983). Assume that $\beta>\gamma+\alpha$ and suppose that an optimal path exists for which $\lim_{t\to\infty}I(t)=0$. Such a path must satisfy the above conditions. An optimal path is monotone and thus cannot bend back on itself in (I,λ)-space, it can intersect the curve $\beta\lambda(t)I(t)+c_P=0$ at most a finite number of times. Note that an optimal path may not intersect this curve at all. There are three possibilities to consider:

(1) The path terminates at time t_0 at a fixed point $(\hat{I},\hat{\lambda})$ on the curve $\beta\lambda(t)I(t)+c_P=0$. In this case, $I(t)=\hat{I}>0$ for $t\geq t_0$. Thus, $\lim_{t\to\infty}I(t)\neq 0$.

(2) The final segment of the path lies above the curve $\beta\lambda(t)I(t)+c_P=0$. Hence, on the final segment of the path $\pi(t)=0$ and

$$\dot{I}(t)=\left[\beta(1-I(t))-\gamma-\alpha\tau(t)\right]I(t)\geq\left[\beta(1-I(t))-\gamma-\alpha\right]I(t) \tag{67}$$

Since $\beta>\gamma+\alpha$, the right-hand side is strictly positive for $I(t)<\frac{\beta-\gamma-\alpha}{\beta}$. Thus, it cannot be the case that $\lim_{t\to\infty}I(t)=0$.

(3) The final segment of the path lies below the curve $\beta\lambda(t)I(t)+c_P=0$. Since $\beta\lambda(t)I(t)<-c_P<0$ and $\lim_{t\to\infty}I(t)=0$, it must be the case that $\lim_{t\to\infty}\lambda(t)=-\infty$. Thus, on the final segment of the path, there must exist t_1 such that $\alpha\lambda(t)<-c_T$ for all $t\geq t_1$. This implies that $\tau(t)=1$ for $t\geq t_1$. Since $\beta\lambda(t)I(t)<-c_P$ on the final segment, it must also be the case that $\pi(t)=1$ for $t\geq t_1$. Hence over this range we have that

$$\begin{aligned}
\dot{I}(t) &= -I(t)\left[\gamma+\alpha\right] && (68)\\
\dot{\lambda}(t) &= \lambda(t)\left[\rho+\gamma+\alpha\right]+\left[\omega+c_T-c_P\right] && (69)\\
H^C(t) &= -\left[\omega+c_T\right]I(t)-c_P\left[1-I(t)\right]+\lambda(t)\dot{I}(t) && (70)
\end{aligned}$$

Solving, yields

$$\begin{aligned}
I(t) &= e^{-(\gamma+\alpha)(t-t_1)}I(t_1) && (71)\\
\dot{I}(t) &= -\left[\gamma+\alpha\right]e^{-(\gamma+\alpha)(t-t_1)}I(t_1) && (72)\\
\lambda(t) &= -\frac{\left[\omega+c_T-c_P\right]}{\left[\rho+\gamma+\alpha\right]}+e^{(\rho+\gamma+\alpha)(t-t_1)}\left(\lambda(t_1)+\frac{\left[\omega+c_T-c_P\right]}{\left[\rho+\gamma+\alpha\right]}\right) && (73)
\end{aligned}$$

Since $\lim_{t\to\infty}\lambda(t)=-\infty$, it must be the case that

$$\lambda(t_1)+\frac{\left[\omega+c_T-c_P\right]}{\left[\rho+\gamma+\alpha\right]}<0 \tag{74}$$

Therefore, noting that $I(t) \in [0,1]$, it follows that

$$\begin{aligned}
\lim_{t\to\infty}\left[e^{-\rho t}H^C(t)\right] &= \lim_{t\to\infty}\left[e^{-\rho t}\left(-\left[\omega+c_T\right]I(t)+c_P\left[1-I(t)\right]+\lambda(t)\dot{I}(t)\right)\right] \\
&= \lim_{t\to\infty}\left[e^{-\rho t}\lambda(t)\dot{I}(t)\right] \\
&= -\lim_{t\to\infty}\left[e^{-\rho t}e^{(\rho+\gamma+\alpha)(t-t_1)}\left(\lambda(t_1)+\frac{\left[\omega+c_T-c_P\right]}{\left[\rho+\gamma+\alpha\right]}\right)(\gamma+\alpha)\,e^{-(\gamma+\alpha)(t-t_1)}I(t_1)\right] \\
&= -e^{-\rho t_1}\left(\lambda(t_1)+\frac{\left[\omega+c_T-c_P\right]}{\left[\rho+\gamma+\alpha\right]}\right)(\gamma+\alpha)\,I(t_1) > 0 \qquad (75)
\end{aligned}$$

This contradicts the requirement that $\lim_{t\to\infty}\left[e^{-\rho t}H^C(t)\right] = 0$. In conclusion, if $\beta > \gamma+\alpha$ there is no optimal path for which $\lim_{t\to\infty} I(t) = 0$. This concludes the proof ■

C. The Individual's Maximization Problem

In this appendix, we set up the individual's optimization problem in more detail. This problem is most naturally written as

$$\max_{\tau_i(t),\pi_i(t)\in[0,1]} \int_0^\infty e^{-\rho t}v_i(t)^T p_i(t)dt \qquad (76)$$

$$s.t. \quad \dot{p}_i(t) = Q_i(t)p_i(t) \qquad (77)$$

where

$$v_i(t) = \left[\omega_{\mathcal{I}} - \tau_i(t)c_T, \omega_{\mathcal{S}} - \pi_i(t)c_P\right]^T \qquad (78)$$

is the vector of state dependent utilities, $p_i(t)$ is a probability measure on the set of states and $Q_i(t)$ is the transition rate (or *intensity*) matrix, given by

$$Q_i(t) = \begin{pmatrix} -(1-\pi_i(t))\beta I(t) & (1-\pi_i(t))\beta I(t) \\ \gamma+\tau_i(t)\alpha & -\gamma-\tau_i(t)\alpha \end{pmatrix} \qquad (79)$$

Note that the individual's transition rate matrix $Q_i(t)$ is a function of the strategies adopted by the individual and the population as a whole. This formulation of the individual's problem is analogous to that in Reluga (2009). To further analyze the individual's problem, it is useful to rewrite the problem as a standard optimal control problem with a single state variable.[18] First, note that at time $t \geq 0$, the individual's health status is given by the indicator function

$$h_i(t) = \begin{cases} 1 & if \quad i \in \mathcal{I}(t) \\ 0 & if \quad i \in \mathcal{S}(t) \end{cases} \qquad (80)$$

The probability that the individual is infected at instant $t \geq 0$ is given by

$$q_i(t) = E[h_i(t)] \qquad (81)$$

[18] This can be done because there are only two health states and the state probabilities must sum to one at all times.

Thus one may write the vector of state probabilities simply as

$$p_i(t) = [q_i(t), 1 - q_i(t)]^T \tag{82}$$

The probability $q_i(t)$, which we will take as the state variable in the individual's control problem, evolves according to a non-homogeneous continuous-time Markov process[19]

$$\dot{q}_i(t) = (1 - q_i(t))(1 - \pi_i(t))\beta I(t) - (\gamma + \tau_i(t)\alpha)q_i(t) \tag{83}$$

The individual's problem can then be rewritten as the following standard optimal control problem:

$$\max_{\tau_i(t),\pi_i(t)\in[0,1]} \int_0^\infty e^{-\rho t} \left[q_i(t) \left[\omega_{\mathcal{I}} - \tau_i(t)c_T\right] + (1 - q_i(t)) \left[\omega_{\mathcal{S}} - \pi_i(t)c_P\right]\right] dt \tag{84}$$

$$s.t. \quad \dot{q}_i(t) = (1 - q_i(t))(1 - \pi_i(t))\beta I(t) - (\gamma + \tau_i(t)\alpha)q_i(t), \quad q_i(0) \in \{0, 1\} \tag{85}$$

Simplifying this problem further, we obtain the following formulation in the main text.

Note that because each individual is negligible and there is no aggregate uncertainty, each individual's best response can be reduced to a function of time alone. This means that the best responses of the individuals are necessarily of the open-loop variety in the sense that each individual commits to an entire path of the personal choice variables $\pi_i(t)$ and $\tau_i(t)$. Note that we do *not* restrict the individuals to choose open-loop strategies, but because any unilateral deviation by any player has no effect on the aggregate evolution of disease prevalence, the optimal closed-loop (or feedback) strategy happens to be of the open-loop variety (see Fudenberg and Tirole, 1991, chapter 4 for further discussion of this point).

D. Steady State Values in Decentralized Setting

The different steady states are given as follows:

Solution A^*: This case corresponds to $\tau_i(t) = 0$ and $\pi_i(t) \in (0, 1)$. The steady state solution is then

$$I_{A^*} \equiv \frac{(\gamma + \rho)c_P}{\beta(\omega - c_P)} > \frac{\rho c_P}{\beta(\omega - c_P)} = I_A \tag{86}$$

$$\mu_{A^*} \equiv \frac{-(\omega - c_P)}{\gamma + \rho} > \frac{-(\omega - c_P)}{\rho} = \lambda_A \tag{87}$$

$$\pi_{A^*} \equiv \frac{(\beta + \rho)c_P - (\beta - \gamma)\omega}{c_P(\beta + \gamma + \rho) - \beta\omega} < \frac{(\beta - \gamma + \rho)c_P - (\beta - \gamma)\omega}{(\beta + \rho)c_P - \beta\omega} = \pi_A \tag{88}$$

$$\tau_{A^*} \equiv 0 = \tau_A \tag{89}$$

Note that $-\mu_{A^*}I_{A^*} = -I_A\lambda_A = \frac{c_P}{\beta}$. Thus, if $I_{A^*} > I_A$ then $-\mu_{A^*} < -\lambda_A$. Hence $\mu_{A^*} > \lambda_A$. Note that $(1 - \pi_{A^*})\beta(1 - I_{A^*}) - \gamma = 0$ and $(1 - \pi_A)\beta(1 - I_A) - \gamma = 0$. Since $I_{A^*} > I_A$, it follows that $\pi_{A^*} < \pi_A$.

Solution B^*: This case corresponds to $\tau_i(t) = 1$ and $\pi_i(t) \in (0, 1)$. The steady state

[19] It is non-homogeneous because infection prevalence $I(t)$ changes over time.

solution is then

$$I_{B^*} \equiv \frac{(\alpha+\gamma+\rho)c_P}{\beta(\omega+c_T-c_P)} > \frac{\rho c_P}{\beta(\omega+c_T-c_P)} = I_B \tag{90}$$

$$\mu_{B^*} \equiv \frac{-(\omega+c_T-c_P)}{\alpha+\gamma+\rho} > \frac{-(\omega+c_T-c_P)}{\rho} = \lambda_B \tag{91}$$

$$\pi_{B^*} \equiv \frac{(\beta+\rho)\,c_P-(\beta-\gamma-\alpha)\,(\omega+c_T)}{(\beta+\gamma+\alpha+\rho)c_P-\beta(\omega+c_T)} \tag{92}$$

$$< \frac{c_P(\beta-\gamma+\rho-\alpha)+(\omega+c_T)(\alpha+\gamma-\beta)}{c_P(\beta+\rho)-\beta(\omega+c_T)} = \pi_B \tag{93}$$

$$\tau_{B^*} \equiv 1 = \tau_B \tag{94}$$

Note that $-\mu_{B^*}I_{B^*} = -I_B\lambda_B = \frac{c_P}{\beta}$. Thus, if $I_{B^*} > I_B$ then $-\mu_{B^*} < -\lambda_B$. Hence $\mu_{B^*} > \lambda_B$. Note that $(1-\pi_{B^*})\beta(1-I_{B^*})-\gamma-\alpha = 0$ and $(1-\pi_B)\beta(1-I_B)-\gamma = 0$. Since $I_{B^*} > I_B$, it follows that $\pi_{B^*} < \pi_B$.

Solution A_0^*: This case corresponds to $\tau_i(t) = 0$ and $\pi_i(t) = 0$. The steady state solution is then

$$I_{A_0^*} \equiv \frac{\beta-\gamma}{\beta} = I_{A_0} \tag{95}$$

$$\mu_{A_0^*} \equiv \frac{-\omega}{\beta+\rho} > \frac{-\omega}{\beta-\gamma+\rho} = \lambda_{A_0} \tag{96}$$

$$\pi_{A_0^*} \equiv 0 = \pi_{A_0} \tag{97}$$

$$\tau_{A_0^*} \equiv 0 = \tau_{A_0} \tag{98}$$

Solution B_0^*: This case corresponds to $\tau_i(t) = 1$ and $\pi_i(t) = 0$. The steady state solution is then

$$I_{B_0^*} \equiv \frac{\beta-\gamma-\alpha}{\beta} = I_{B_0} \tag{99}$$

$$\mu_{B_0^*} \equiv \frac{-(\omega+c_T)}{\beta+\rho} > \frac{-(\omega+c_T)}{\beta-\gamma-\alpha+\rho} = \lambda_{B_0} \tag{100}$$

$$\pi_{B_0^*} \equiv 1 = \pi_{B_0} \tag{101}$$

$$\tau_{B_0^*} \equiv 0 = \tau_{B_0} \tag{102}$$

Solution C_0^*: This case corresponds to $\tau_i(t) \in (0,1)$ and $\pi_i(t) = 0$. The steady state solution is then

$$I_{C_0^*} \equiv \frac{\alpha\omega-(\gamma+\rho)c_T}{\beta c_T} > \frac{\alpha\omega+(\beta-\gamma-\rho)c_T}{2\beta c_T} = {I_{C_0}}^{20} \tag{103}$$

$$\mu_{C_0^*} \equiv \frac{-c_T}{\alpha} = \lambda_{C_0} \tag{104}$$

$$\pi_{C_0^*} \equiv 0 = \pi_{C_0} \tag{105}$$

$$\tau_{C_0^*} \equiv \frac{(\beta+\rho)c_T-\alpha\omega}{\alpha c_T} < \frac{c_T(\beta-\gamma+\rho)-\alpha\omega}{2\alpha c_T} = {\tau_{C_0}}^{21} \tag{106}$$

Note that $\beta(1-I_{C_0^*})-\gamma-\alpha\tau_{C_0^*}=\beta(1-I_{C_0})-\gamma-\alpha\tau_{C_0}=0$. Hence, if $I_{C_0^*}>I_{C_0}$then $\tau_{C_0^*}<\tau_{C_0}$.

E. Decomposition and Derivation of Incentive Schemes

We start by considering the problem faced by the maverick:

Proof: To formally derive the shadow price of the maverick, consider a situation in which, under central direction, all individuals except for the maverick behave in the socially optimal fashion. Denote the aggregate level of infection on the resulting optimal path by $I(t)$. The maverick individual evades instructions and chooses his or her own levels of treatment $\tau_i(t)$ and protection $\pi_i(t)$. As before, this individual takes the trajectory of aggregate infection $I(t)$ as given. This maverick's objective is to solve

$$\max_{\tau_i(t),\pi_i(t)\in[0,1]}\int_0^\infty e^{-\rho t}\left[-q_i(t)\left[\omega+\tau_i(t)c_T\right]-(1-q_i(t))\pi_i(t)c_P\right]dt \tag{107}$$

subject to the state equation (18). But note that in this problem, we control for optimal aggregate infection, i.e. the path of infection is not allowed to follow its decentralized equilibrium path.

The current-value Hamiltonian for the maverick's problem is given by

$$\begin{aligned} H_m^D \equiv\ & -q_i(t)\left(\omega+\tau_i(t)c_T\right)-(1-q_i(t))\pi_i(t)c_P & (108)\\ & +\eta_i(t)\left[(1-q_i(t))(1-\pi_i(t))\beta I(t)-q_i(t)\left(\gamma+\tau_i(t)\alpha\right)\right] & (109)\end{aligned}$$

Note that in this equation, the path $I(t)$ is the socially optimal one and thus H_m^D differs from H_i^D only because they are evaluated along different paths for aggregate disease prevalence. Next, the costate variable for the maverick's problem evolves according to the differential equation

$$\dot\eta_i(t)=\eta_i(t)\left[\rho+\gamma+\alpha\tau_i(t)+(1-\pi_i(t))\beta I(t)\right]+\omega+\tau_i(t)c_T-\pi_i(t)c_P \tag{110}$$

As in the treatment of the decentralized equilibrium, symmetry allows us to drop the subscript i on the costate and control variables. For completeness, the transversality condition $\lim_{t\to\infty}e^{-\rho t}\eta_i(t)=0$ holds, and the Arrow sufficiency conditions for an individual optimum are also satisfied ■

Finally, we derive the optimal prevention and treatment subsidies:

Proof: Suppose that subsidies $s_P(t)$ and $s_T(t)$ are given for engaging in prevention and treatment, respectively. The individual's objective is then to solve

$$\max_{\tau_i(t),\pi_i(t)\in[0,1]}\int_0^\infty e^{-\rho t}\left[-q_i(t)\left[\omega+\tau_i(t)\left[c_T-s_T(t)\right]\right]-(1-q_i(t))\pi_i(t)\left[c_P-s_P(t)\right]\right]dt \tag{111}$$

subject to the state equation (18). The current-value Hamiltonian now takes the form

$$\begin{aligned}\hat H_i^D \equiv\ & -q_i(t)\left(\omega+\tau_i(t)\left[c_T-s_T(t)\right]\right)-(1-q_i(t))\pi_i(t)\left[c_P-s_P(t)\right]\\ & +\phi_i(t)\left[(1-q_i(t))(1-\pi_i(t))\beta I(t)-q_i(t)\left(\gamma+\tau_i(t)\alpha\right)\right] & (112)\end{aligned}$$

The associated costate variable evolves according to the differential equation

$$\begin{aligned}\dot{\phi}_i(t) &= \phi_i(t)\left[\rho+\gamma+\alpha\tau_i(t)+(1-\pi_i(t))\beta I(t)\right] \\ &\quad +\omega+\tau_i(t)(c_T-s_T(t))-\pi_i(t)(c_P-s_P(t))\end{aligned} \tag{113}$$

Is is straightforward to see that the Hamiltonian conditions for the planner and for the individual coincide if

$$\beta\phi_i(t)I(t)+(c_P-s_P(t))=\beta\lambda(t)I(t)+c_P \tag{114}$$

and

$$\alpha\phi_i(t)+(c_T-s_T(t))=\alpha\lambda(t)+c_T \tag{115}$$

If equations (114) and (115) are satisfied, the individual and the planner will choose the same socially optimal levels of protection and treatment.By symmetry we can drop the subscript i, in which case these equations can be written as follows,

$$\beta\phi(t)I(t)+(c_P-s_P(t))=\beta\lambda(t)I(t)+c_P \tag{116}$$

$$\alpha\phi(t)+(c_T-s_T(t))=\alpha\lambda(t)+c_T \tag{117}$$

and equation (118) can be written

$$\begin{aligned}\dot{\phi}(t) &= \phi(t)\left[\rho+\gamma+\alpha\tau(t)+(1-\pi(t))\beta I(t)\right] \\ &\quad +\omega+\tau(t)(c_T-s_T(t))-\pi(t)(c_P-s_P(t))\end{aligned} \tag{118}$$

where $\pi(t)$ and $\tau(t)$ are socially optimal. Rearranging (116) and (117) yields

$$s_P(t)=\beta I(t)\left[\phi(t)-\lambda(t)\right] \tag{119}$$

$$s_T(t)=\alpha\left[\phi(t)-\lambda(t)\right] \tag{120}$$

Substituting in (118) yields

$$\begin{aligned}\dot{\phi}(t) &= \phi(t)\left[\rho+\gamma+\alpha\tau(t)+(1-\pi(t))\beta I(t)\right] \\ &\quad +\omega+\tau(t)(c_T-\alpha\left[\phi(t)-\lambda(t)\right])-\pi(t)(c_P-\beta I(t)\left[\phi(t)-\lambda(t)\right]) &(121)\\ &= \phi(t)\left[\rho+\gamma+\beta I(t)\right]+\omega+\tau(t)\left[c_T+\alpha\lambda(t)\right]-\pi(t)\left[c_P+\beta I(t)\lambda(t)\right] &(122)\end{aligned}$$

■

Next, we derive the optimal infection tax (or health subsidy):

Proof: Suppose a lump-sum stock tax $T(t)$ is levied on infected individuals. An individual's problem is then given by

$$\max_{\tau_i(t),\pi_i(t)\in[0,1]}\int_0^\infty e^{-\rho t}\left[-q_i(t)\left[\omega+\tau_i(t)c_T+T(t)\right]-(1-q_i(t))\pi_i(t)c_P\right]dt \tag{123}$$

where the maximization is subject to the state equation (18). Note that this problem has the same solution (up to a constant) as a problem with a modified objective function,

in which the tax $T(t)$ on infected individuals is replaced by a subsidy of $T(t)$ given to susceptible individuals. The modified problem is given by

$$\max_{\tau_i(t),\pi_i(t)\in[0,1]} \int_0^\infty e^{-\rho t}\left[-q_i(t)\left[\omega+\tau_i(t)c_T\right]-(1-q_i(t))\pi_i(t)c_P+(1-q_i(t))T(t)\right]dt \quad (124)$$

The current-value Hamiltonian for the former problem is given by

$$\begin{aligned}\bar{H}_i^D \equiv\ & -q_i(t)\left[\omega+\tau_i(t)c_T+T(t)\right]-(1-q_i(t))\pi_i(t)c_P & (125)\\ & +\psi_i(t)\left[(1-q_i(t))(1-\pi_i(t))\beta I(t)-(\gamma+\tau_i(t)\alpha)q_i(t)\right] & (126)\end{aligned}$$

We will now investigate the conditions under which the solution to this problem and the corresponding costate variable coincide with the optimal centralized solution.

The costate equation for the decentralized problem with a stock tax is given by

$$\begin{aligned}\dot{\psi}_i(t) =\ & \psi_i(t)\left[\rho+\gamma+\tau_i(t)\alpha+(1-\pi_i(t))\beta I(t)\right]\\ & +\left[\omega+\tau_i(t)c_T+T(t)-\pi_i(t)c_P\right] \quad (127)\end{aligned}$$

Recall for reference that the costate equation for the centralized problem is given by

$$\begin{aligned}\dot{\lambda}(t) =\ & \lambda(t)\left[\rho+\gamma+\alpha\tau(t)+(1-\pi(t))\beta(2I(t)-1)\right]\\ & +\left[\omega+\tau(t)c_T-\pi(t)c_P\right] \quad (128)\end{aligned}$$

Suppose that $\psi_i(t)=\lambda(t)$ for all t and $i\in\mathcal{P}$. The Hamiltonian conditions will then be identical and thus we can assume that $\tau_i(t)=\tau(t)$ and $\pi_i(t)=\pi(t)$ for all $t\geq 0$ and $i\in\mathcal{P}$. Therefore $I(t)$ is the same in each equation. The decentralized equation (127) can therefore be written as follows

$$\begin{aligned}\dot{\lambda}(t) =\ & \lambda(t)\left[\rho+\gamma+\tau(t)\alpha+(1-\pi(t))\beta I(t)\right]\\ & +\left[\omega+\tau(t)c_T+T(t)-\pi(t)c_P\right] \quad (129)\end{aligned}$$

Subtracting (128) from (129) and solving for the tax $T(t)$ yields

$$T(t)=-\lambda(t)(1-\pi(t))\beta(1-I(t))>0 \quad (130)$$

With this tax on the infected, there is a decentralized equilibrium path which results in socially optimal individual decisions. Again, symmetry has allowed us to drop the subscript i on the costate and control variables. For completeness, the transversality condition $\lim_{t\to\infty} e^{-\rho t}\psi_i(t)=0$ holds, and the Arrow sufficiency conditions for an individual optimum are also satisfied ■

F. Sufficiency of Hamiltonian Conditions in Decentralized Setting

As noted in the main text, the analysis of the centralized problem is complicated by the fact that the Hamiltonian necessary conditions for optimality of paths are not sufficient conditions. In particular, neither Mangasarian's nor Arrow's sufficiency conditions hold. This stems from the convexity of the planner's current-value Hamiltonian in the state variable. In the decentralized setting, each individual's current-value Hamiltonian is linear in the state variable, which raises the hope that in this setting, the Hamiltonian

conditions are both necessary and sufficient for optimality of paths. In this appendix, we show that while the Mangasarian sufficiency condition does not hold for the individual's problem, the condition by Arrow does. This implies that any path that satisfies the Hamiltonian conditions is a perfect foresight equilibrium path.

In general, the current-value Hamiltonian H_i^D is not a jointly concave function of $q_i(t)$, $\tau_i(t)$ and $\pi_i(t)$. Recall that

$$\begin{aligned} H_i^D &\equiv -q_i(t)\left[\omega + \tau_i(t)c_T\right] - (1-q_i(t))\pi_i(t)c_P && (131)\\ &\quad +\mu_i(t)\left[(1-q_i(t))(1-\pi_i(t))\beta I(t) - q_i(t)(\gamma + \tau_i(t)\alpha)\right] && (132)\end{aligned}$$

In simplified notation, we write

$$\begin{aligned} H_i^D &= -q\left[\omega + \tau c_T\right] - (1-q)\pi c_P && (133)\\ &\quad +\mu\left[(1-q)(1-\pi)\beta I(t) - q(\gamma + \tau\alpha)\right] && (134)\end{aligned}$$

The Hessian for this function is given by

$$\begin{aligned} Hessian &= \begin{bmatrix} \frac{\partial^2 H_i^D}{\partial^2 q} & \frac{\partial^2 H_i^D}{\partial q \partial \tau} & \frac{\partial^2 H_i^D}{\partial q \partial \pi} \\ \frac{\partial^2 H_i^D}{\partial q \partial \tau} & \frac{\partial^2 H_i^D}{\partial^2 \tau} & \frac{\partial^2 H_i^D}{\partial \tau \partial \pi} \\ \frac{\partial^2 H_i^D}{\partial q \partial \pi} & \frac{\partial^2 H_i^D}{\partial \tau \partial \pi} & \frac{\partial^2 H_i^D}{\partial^2 \pi} \end{bmatrix} && (135)\\ &= \begin{bmatrix} 0 & -(c_T + \alpha\mu) & c_P + \beta I(t)\mu \\ -(c_T + \alpha\mu) & 0 & 0 \\ c_P + \beta I(t)\mu & 0 & 0 \end{bmatrix} && (136)\end{aligned}$$

The second-order principal minors are 0, $-\left(c_P + \beta I(t)\mu\right)^2$ and $-(c_T + \alpha\mu)^2$ respectively. The conditions for H_i^D to be jointly concave in q, τ and π are for the first-order and third-order principal minors of the Hessian to be non-positive and for the second-order principal minor to be non-negative for all combinations of $q \in (0,1]$ and $\tau, \pi \in [0,1]$. In general, however, the last condition is not satisfied and thus Mangasarian's sufficiency condition is not satisfied. Next, we consider Arrow's sufficiency condition.

F.1. Arrow's Sufficiency Theorem. Suppose that $\hat{\tau}$ and $\hat{\pi}$ maximize the current-value Hamiltonian $H_i^D(q, \tau, \pi, \mu, t)$. Define the maximized current-value Hamiltonian:

$$\hat{H}_i^D(q, \mu, t) \equiv H_i^D(q, \hat{\tau}, \hat{\pi}, \mu, t)$$

The necessary conditions for an optimal path are sufficient if $\hat{H}_i^D(q, \mu, t)$ is a concave function of q taking μ and t as constant. The Hamiltonian conditions for $\hat{\tau}$ and $\hat{\pi}$ are as follows:

$$\begin{aligned} \hat{\tau} &= 0 \quad if \quad \alpha\mu > -c_T && (137)\\ \hat{\tau} &\in [0,1] \quad if \quad \alpha\mu = -c_T && (138)\\ \hat{\tau} &= 1 \quad if \quad \alpha\mu < -c_T && (139)\end{aligned}$$

$$\hat{\pi} = 0 \quad if \quad \beta\mu I(t) > -c_P \tag{140}$$
$$\hat{\pi} \in [0,1] \quad if \quad \beta\mu I(t) = -c_P \tag{141}$$
$$\hat{\pi} = 1 \quad if \quad \beta\mu I(t) < -c_P \tag{142}$$

The maximized current-value Hamiltonian is given by

$$\hat{H}_i^D = -q\left[\omega + \hat{\tau}c_T\right] - (1-q)\hat{\pi}c_P + \mu\left[(1-q)(1-\hat{\pi})\beta I(t) - q(\gamma + \hat{\tau}\alpha)\right] \tag{143}$$

There are nine possible cases, depending on the values of μ and t (via its influence on $I(t)$):

Case 1: $\alpha\mu > -c_T, \beta\mu I(t) > -c_P$. In this case, $\hat{\tau} = 0$, $\hat{\pi} = 0$ and

$$\hat{H}_i^D(q,\mu,t) = -q\omega + \mu\left[(1-q)\beta I(t) - q\gamma\right]$$

Case 2: $\alpha\mu > -c_T, \beta\mu I(t) = -c_P$. In this case, $\hat{\tau} = 0$, $\hat{\pi} \in [0,1]$**.** The coefficient of $\hat{\pi}$ is zero and hence

$$\hat{H}_i^D(q,\mu,t) = -q\omega + \mu\left[(1-q)\beta I(t) - q\gamma\right]$$

Case 3: $\alpha\mu > -c_T, \beta\mu I(t) < -c_P$. In this case, $\hat{\tau} = 0$, $\hat{\pi} = 1$ and

$$\hat{H}_i^D(q,\mu,t) = -q\omega - (1-q)c_P - \mu q\gamma \tag{144}$$

Case 4: $\alpha\mu = -c_T, \beta\mu I(t) > -c_P$. In this case, $\hat{\tau} \in [0,1]$ and $\hat{\pi} = 0$. The coefficient of $\hat{\tau}$ is zero and hence

$$\hat{H}_i^D(q,\mu,t) = -q\omega + \mu\left[(1-q)\beta I(t) - q\gamma\right] \tag{145}$$

Case 5: $\alpha\mu = -c_T, \beta\mu I(t) = -c_P$. In this case, $\hat{\tau}, \hat{\pi} \in [0,1]$. The coefficients of $\hat{\tau}$ and $\hat{\pi}$ are both zero and hence

$$\hat{H}_i^D(q,\mu,t) = -q\omega + \mu\left[(1-q)\beta I(t) - q\gamma\right] \tag{146}$$

Case 6: $\alpha\mu = -c_T, \beta\mu I(t) < -c_P$. In this case, $\hat{\tau} \in [0,1]$ and $\hat{\pi} = 1$. The coefficient of $\hat{\tau}$ is zero and hence

$$\hat{H}_i^D(q,\mu,t) = -q\omega - (1-q)c_P - \mu q\gamma \tag{147}$$

Case 7: $\alpha\mu < -c_T, \beta\mu I(t) > -c_P$. In this case, $\hat{\tau} = 1$, $\hat{\pi} = 0$ and

$$\hat{H}_i^D(q,\mu,t) = -q\left[\omega + c_T\right] + \mu\left[(1-q)\beta I(t) - q(\gamma + \alpha)\right] \tag{148}$$

Case 8: $\alpha\mu < -c_T, \beta\mu I(t) = -c_P$. In this case, $\hat{\tau} = 1$, $\hat{\pi} \in [0,1]$**.** The coefficient of $\hat{\pi}$ is zero and hence

$$\hat{H}_i^D(q,\mu,t) = -q\left[\omega + c_T\right] + \mu\left[(1-q)\beta I(t) - q(\gamma + \alpha)\right] \tag{149}$$

Case 9: $\alpha\mu < -c_T, \beta\mu I(t) < -c_P$. In this case, $\hat{\tau} = 1$, $\hat{\pi} = 1$ and

$$\hat{H}_i^D(q, \mu, t) = -q\left[\omega + c_T\right] - (1-q)c_P - \mu q(\gamma + \alpha) \tag{150}$$

In every case, $\hat{H}_i^D(q, \mu, t)$ is linear in the state q and hence concave in q (holding μ and t constant). Thus Arrow's sufficiency condition is satisfied.

For completeness, it would be noted that analogous arguments hold also for the controlled decentralized setting (i.e. with either subsidies or taxes) and thus the paths described there also constitute decentralized equilibria.

References

[1] Aadland, D., D. Finnoff and K. X. D. Huang (2010): Syphilis Cycles, *mimeo.*

[2] Agusto, F. B., N. Marcus and K. O. Okosun (2012): Application of Optimal Control to the Epidemiology of Malaria, *Electronic Journal of Differential Equations*, 2012(81), 1-22.

[3] Almeder, C., G. Feichtinger, W. C. Sanderson and V. M. Veliov (2007): Prevention and medication of HIV/AIDS: the case of Botswana, *Central European Journal of Operations Research*, 15(1), 47-61.

[4] Anderson, R. M. and R. M. May (1991): Infectious Diseases of Humans: Dynamics and Control, *Oxford University Press.*

[5] Anderson, S., R. Laxminarayan and S. W. Salant (2012): Diversify or Focus? Spending to Combat Infectious Diseases When Budgets Are Tight, *Journal of Health Economics*, 31(4), 658-675.

[6] Aronsson, T., K. Backlund and K.-G. Lofgren (1998): Nuclear Power, Externalities and Non-Standard Pigouvian Taxes, *Environmental and Resource Economics*, 11(2), 177-195.

[7] Arrow, K. J. and M. Kurz (1969): Optimal Public Investment Policy and Controllability with Fixed Private Savings Ratio, *Journal of Economic Theory*, 1(2), 141-177.

[8] Behrens, D. A. , J. P. Caulkins, G. Tragler and G. Feichtinger (2000): Optimal Control of Drug Epidemics: Prevent and Treat – But Not at the Same Time?, *Management Science*, 46(3), 333-347.

[9] Blayneh, K., Y. Cao and H.-D. Kwon (2009): Optimal Control of Vector-Borne Diseases: Treatment and Prevention, *Discrete and Continuous Dynamical Systems - Series B*, 11(3), 587-611.

[10] Brito, D. L., E. Sheshinski and M. D. Intriligator (1991): Externalities and Compulsory Vaccinations, *Journal of Public Economics*, 45(1), 69-90.

[11] Brock, W. A. and D. Starrett (2003): Managing Systems with Non-Convex Positive Feedback, *Environmental and Resource Economics*, 26(4), 575–602.

[12] Cannefax, G. R. (1965): Immunity in Syphilis, *British Journal of Venereal Diseases*, 41(4), 260-274.

[13] Caputo, M. R. (2005): Foundations of Dynamic Economic Analysis: Optimal Control Theory and Applications, *Cambridge University Press.*

[14] Cohen, J., P. Dupas and S. G. Schaner (2011): Price Subsidies, Diagnostic Tests, and Targeting of Malaria Treatment: Evidence from a Randomized Controlled Trial, NBER Working Paper 17943.

[15] Daley, D. J. and J. Gani (2001): Epidemic Modelling: An Introduction, *Cambridge Studies in Mathematical Biology.*

[16] Dasgupta, P. and K.-G. Maler (2003): The Economics of Non-Convex Ecosystems: Introduction, *Environmental and Resource Economics*, 26(4), 499-525.

[17] Deisssenberg, C., G. Feichtinger, W. Semmler and F. Wirl (2004): Multiple Equilibria, History Dependence, and Global Dynamics in Intertemporal Optimization Models, in *Economic Complexity*, W. A. Barnett, C. Deissenberg and G. Feichtinger (eds.), Elsevier.

[18] Dodd, P. J., P. J. White and G. P. Garnett (2010): Notions of Synergy for Combinations of Interventions against Infectious Diseases in Heterogeneously Mixing Populations, *Mathematical Biosciences*, 227(2), 94-104.

[19] Farzin, Y. H. (1996): Optimal Pricing of Environmental and Natural Resource Use with Stock Externalities, *Journal of Public Economics*, 62(1-2), 31-57.

[20] Feichtinger, G. (1984): On the Synergistic Influence of Two Control Variables on the State of Nonlinear Optimal Control Models, *Journal of the Operational Research Society*, 35(10), 907-914.

[21] Fudenberg, D. and J. Tirole (1991): Game Theory, *MIT Press.*

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

[23] Gersovitz, M. (2010): Disinhibition and Immiserization in a Model of Susceptible-Infected-Susceptible (SIS) Diseases, *mimeo.*

[24] Gersovitz, M. and J. S. Hammer (2003): Infectious Diseases, Public Policy and the Marriage of Economics and Epidemiology, *World Bank Research Observer*, 18(2), 129-157.

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

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

[27] 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.

[28] Goyal, S. and A. Vigier (2010): Endogenous Interaction and Vaccination, *mimeo.*

[29] Greenwood, J., P. Kircher and M. Tertilt (2009): An Equilibrium Model of the Malawian HIV/AIDS Epidemic, *mimeo.*

[30] Keeling, M. J. and P. Rohani (2008): Modeling Infectious Diseases in Humans and Animals, *Princeton University Press.*

[31] Hethcote, H. W. and J. A. Yorke (1984): Gonorrhea Transmission Dynamics and Control, Lecture Notes in Biomathematics, vol. 56, *Springer-Verlag.*

[32] Horan, R. D., E. P. Fenichel, K. L. S. Drury and D. M. Lodge (2011): Managing Ecological Thresholds in Coupled Environmental–Human Systems, *Proceedings of the National Academy of Sciences*, 108(18), 7333-7338.

[33] Klein, E., R. Laxminarayan, D. L. Smith and C. A. Gilligan (2007): Economic Incentives and Mathematical Models of Disease, *Environment and Development Economics*, 12(5), 707-732.

[34] Krauthammer, C. (2009): Preventive Care Isn't the Magic Bullet for Health Care Costs, *The Washington Post*, Friday, August 14, 2009.

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

[36] Kremer, M. and C. M. Snyder (2013): When Is Prevention More Profitable than Cure? The Impact of Time-Varying Consumer Heterogeneity, *NBER Working Paper No. 18862.*

[37] Maler, K. G., A. Xepapadeas and A. de Zeeuw (2003): The Economics of Shallow Lakes, *Environmental and Resource Economics*, 26(4), 603-624.

[38] McKinnon, P. S. and S. L. Davis (2004): Pharmacokinetic and Pharmacodynamic Issues in the Treatment of Bacterial Infectious Diseases, *European Journal of Clinical Microbiology and Infectious Diseases*, 23(4), 271-288.

[39] Michel, P. (1982): On the Transversality Condition in Infinite Horizon Optimal Problems, *Econometrica*, 50(4), 1975-1985.

[40] Morin, B. R., C. Perrings, S. Levin and A. Kinzig (2014): Disease Risk Mitigation: The Equivalence of Two Selective Mixing Strategies on Aggregate Contact Patterns and Resulting Epidemic Spread, *Journal of Theoretical Biology, forthcoming.*

[41] Pepin, J. and D. Mabey (2003): Sexually Transmitted Infections in Africa: Single Dose Treatment is Now Affordable, *Sexually Transmitted Infections*, 79(6), 432-434.

[42] Philipson, T. (2000): Economic Epidemiology and Infectious Disease, Handbook of Health Economics, volume 1B, Part 8; Cuyler, A. J. and J. P. Newhouse (eds.), Amsterdam: *North Holland*.

[43] Piot, P., M. Bartos, H. Larson, D. Zewdie and P. Mane (2008): Coming to Terms with Complexity: A Call to Action for HIV Prevention, *The Lancet*, 372(9641), 845-859.

[44] Reluga, T. C. (2009): An SIS Epidemiology Game with Two Subpopulations, *Journal of Biological Dynamics*, 3(5), 515-531.

[45] Reluga, T. C. (2010): Game Theory of Social Distancing in Response to an Epidemic, *PLoS Computational Biology*, 6(5).

[46] Rose, G. (1985): Sick Individuals and Sick Populations, *International Journal of Epidemiology*, 14(1), 32-38.

[47] Rose, G. (1992): The Strategy of Preventive Medicine, *Oxford University Press*.

[48] 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.

[49] Rubio, J. S. and L. Escriche (2001): Strategic Pigouvian Taxation, Stock Externalities and Polluting Non-Renewable Resources, *Journal of Public Economics*, 79(2), 297-313.

[50] Russell, L. B. (1986): Is Prevention Better than Cure?, *Brookings Institution*.

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

[52] Seierstad, A. and K. Sydsaeter (1987): Optimal Control Theory with Economic Applications, *North Holland*.

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

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

[55] 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.

[56] Stewart. J. J., J. G. Thomsen and R. K. Gardner (1951): Reinfection in Early Syphilis That Had Been Treated with Penicillin (Ping-Pong Syphilis), *A.M.A. Archives of Dermatology and Syphilology*, 63(1), 136-137.

[57] Tahvonen, O. and J. Kuuluvainen (1993): Economic Growth, Pollution, and Renewable Resources, *Journal of Environmental Economics and Management*, 24(2), 101-118.

[58] Toxvaerd, F. (2009a): Recurrent Infection and Externalities in Treatment, *mimeo.*

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

[60] Toxvaerd, F. (2010): Recurrent Infection and Externalities in Prevention, *mimeo.*

[61] Wagener, F. O. O. (2003): Skiba Points and Heteroclinic Bifurcations, with Applications to the Shallow Lake System, *Journal of Economic Dynamics and Control*, 27(9), 1533-1561.

[62] World Health Organization (2004): Guidelines for the Management of Sexually Transmitted Infections, http://www.who.int/hiv/pub/sti/pub6/en/.

[63] Zaman, G., Y. H. Kang and I. H. Jung (2007): Optimal Vaccination and Treatment in the SIR Epidemic Model, *mimeo.*