Cost-Effective Control of Chronic Viral Diseases:

Finding the Optimal Level of Screening and Contact Tracing

Benjamin Armbruster* Margaret L. Brandeau†

January 31, 2010

Abstract

Chronic viral diseases such as human immunodeficiency virus (HIV) and hepatitis B virus (HBV) afflict millions of people worldwide. A key public health challenge in managing such diseases is identifying infected, asymptomatic individuals so that they can receive antiviral treatment. Such treatment can benefit both the treated individual (by improving quality and length of life) and the population as a whole (through reduced transmission). We develop a compartmental model of a chronic, treatable infectious disease and use it to evaluate the cost and effectiveness of different levels of screening and contact tracing. We show that: 1) the optimal strategy is to get infected individuals into treatment at the maximal rate until the incremental health benefits balance the incremental cost of controlling the disease; 2) as one reduces the disease prevalence by moving people into treatment (which decreases the chance that they will infect others), one should increase the level of contact tracing to compensate for the decreased effectiveness of screening; 3) as the disease becomes less prevalent, it is optimal to spend more per case identified; and 4) the relative mix of screening and contact tracing at any level of disease prevalence is such that the marginal efficiency of contact tracing (cost per infected person found) equals that of screening if possible (e.g., when capacity limitations are not binding). We also show how to determine the cost-effective equilibrium level of disease prevalence (among untreated individuals), and we develop an approximation of the path of the optimal prevalence over time. Using this, one can obtain a close approximation of the optimal

---

*armbrusterb@gmail.com Department of Industrial Engineering and Management Sciences, Northwestern University

†brandeau@stanford.edu Department of Management Science and Engineering, Stanford University

solution without having to solve an optimal control problem. We apply our methods to an example of hepatitis B virus.

Keywords: contact tracing – chronic viral disease – compartmental model – cost-effectiveness

1 Introduction

Chronic viral diseases such as hepatitis B virus (HBV), hepatitis C virus (HCV), and human immunodeficiency virus (HIV) infect millions of people worldwide each year (CDC, 2008c; UNAIDS, 2006). Many such infections have no cure and, if untreated, can lead to disease, disability, and death. However, for HBV, HCV, HIV, and other chronic viral infections, antiviral treatments can often extend life, improve quality of life, and in some cases reduce a treated individual's infectivity. Such treatments thus can provide benefit both to the treated individual and to the population as a whole (through reduced transmission). In recent years, the World Health Organization and other health organizations have devoted significant resources to scaling up HIV treatment worldwide (UNAIDS, 2005). Similar efforts are underway, on a smaller scale, to increase treatment of individuals infected with HBV and HCV (AB 158, 2008).

A key public health challenge in managing chronic viral diseases is identifying infected, asymptomatic individuals so that they can receive treatment. Individuals identified before symptoms develop typically derive greater benefit from treatment than individuals who receive treatment only in an advanced stage of disease. For example, highly active antiretroviral therapy for HIV, when given before the development of AIDS, significantly extends length and quality of life (Freedberg et al., 2001). As another example, early management and treatment of chronic HBV infection can provide significant health benefits for treated individuals (Brooks et al., 2001; Hutton et al., 2007; Kanwal et al., 2005; Shepherd et al., 2006).

Two key means of identifying chronically infected, asymptomatic individuals are screening and contact tracing. Screening (via a blood test) may take place during routine care, or may be part of a targeted campaign. In contact tracing (also known as partner notification), the contacts of a newly identified infected individual (a so-called "index case") are located and then screened for the infection. Contact tracing is typically more expensive per case found than screening, but can be an effective means of identifying infected individuals.

In this paper we consider the optimal mix of screening and contact tracing for a chronic viral disease. The goal is to maximize net health benefit in the population, which we define as the value of (quality-adjusted) life years experienced minus the cost of screening and contact tracing.

Analyses of optimal resource allocation for disease control date back to the late 1960s. Taylor (1968) focused on bovine virus diarrhea and Revelle et al. (1969) focused on tuberculosis control measures. Sanders (1971) and Sethi (1974) applied optimal control to curable diseases. Their models bear some similarity to our model except that our model is complicated by the presence of additional states, the lack of a cure, and contact tracing as a second control in addition to screening. Armbruster and Brandeau (2007b) used an optimal control approach to find the best combination of screening and contact tracing to control an infectious disease. They examined a *curable* endemic disease (instead of a chronic one as we do here) and assumed that contact tracing has a fixed efficacy, if it is performed at all. Here we allow the investment in contact tracing (and thus the resulting efficacy) to vary, similar to Armbruster and Brandeau (2007a) who performed a simulation study of contact tracing for a curable endemic disease.

Section 2 presents our model and describes its solution. In section 2.1, we determine the optimal rate of individuals entering treatment and then in section 2.2, we determine the optimal mix of screening and contact tracing. In section 2.3, we describe how the optimal solution, including the optimal equilibrium, can be calculated numerically. In section 2.4, we develop a useful approximation which we use to analytically characterize the optimal trajectory of control and generate insight. We present a numerical example in section 3 and conclude with discussion in section 4.

2   Model of a Chronic Viral Disease

Consider a population in which a chronic viral disease is spreading. Susceptible individuals acquire the infection through contact with infected individuals. Once infected, individuals are initially asymptomatic. Their immune system may resolve the infection, in which case they are immune from further infection. Otherwise, the infection becomes chronic (i.e., not resolved). Chronically infected individuals learn of their infection either through the development of symptoms or via a screening test. Individuals whose infection has been identified enter treatment. All infected individuals are infectious; those in treatment may be less infectious. We allow for the possibility
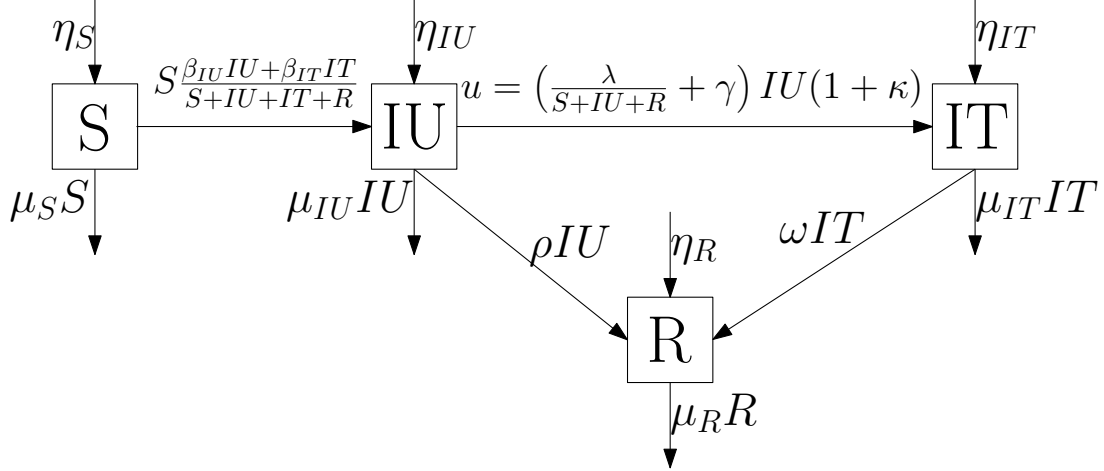
Figure 1: Disease model

that some chronically infected individuals in treatment can recover from the infection.

We model the disease using an epidemic model, figure 1, with four compartments: $S(t)$ susceptible individuals, $IU(t)$ infected but untreated individuals, $IT(t)$ infected and treated individuals, and $R(t)$ recovered individuals who are now immune from infection. We define $\boldsymbol{s}(t) := (S(t), IU(t), IT(t), R(t))$ as the state vector. We define η_S, η_{IU}, η_{IT}, and η_R as the rates of entry into each compartment from outside the population; β_{IU} and β_{IT} as the rates of sufficient contact between susceptibles and infected individuals in compartments IU and IT, respectively, that are sufficient to transmit the infection; ρ as the per person rate at which the infection spontaneously resolves; ω as the per person rate at which treatment effects a cure; γ as the per person rate of symptom development; and $\boldsymbol{\mu} := (\mu_S, \mu_{IU}, \mu_{IT}, \mu_R)$ as the per person rates of leaving the population from each compartment (these rates include deaths and emigration).

We consider two interventions: a screening program that screens individuals not in treatment at a rate $\lambda(t)$, and a contact tracing program implemented at a rate that finds an average of $\kappa(t)$ secondary infections per otherwise reported index case. We let $u(t)$ be the rate at which people enter treatment (i.e., the rate of movement from state IU to state IT). This is the rate at which we identify new index cases (found through screening and symptom development) and their contacts (found through contact tracing): thus, $u(t) = \left(\frac{\lambda(t)}{S(t)+IU(t)+R(t)} + \gamma \right) IU(t)(1 + \kappa(t))$. In solving the problem, we find it convenient to first determine the optimal total level of control $u^*(t)$, and then the optimal values of $\lambda(t)$ and $\kappa(t)$ given $u^*(t)$. For simplicity in notation we drop the dependence

on t in the state and control variables ($\boldsymbol{s}$, κ, λ, and u) unless we want to highlight this dependency.

The dynamics of the model are as follows:

$$\dot{S} = f_1(\boldsymbol{s}, u) := \eta_S \quad -\mu_S S \qquad\qquad -S\frac{\beta_{IU} IU + \beta_{IT} IT}{S + IU + IT + R} \tag{1}$$

$$\dot{IU} = f_2(\boldsymbol{s}, u) := \eta_{IU} -\mu_{IU} IU - \rho IU \qquad +S\frac{\beta_{IU} IU + \beta_{IT} IT}{S + IU + IT + R} - u \tag{2}$$

$$\dot{IT} = f_3(\boldsymbol{s}, u) := \eta_{IT} -\mu_{IT} IT \qquad - \omega IT \qquad\qquad\qquad + u \tag{3}$$

$$\dot{R} = f_4(\boldsymbol{s}, u) := \eta_R \quad -\mu_R R \quad +\rho IU + \omega IT \tag{4}$$

$$u = \left(\frac{\lambda}{S + IU + R} + \gamma\right) IU(1 + \kappa). \tag{5}$$

We impose three constraints on this model: the size of each intervention must be nonnegative, $\lambda \geq 0$ and $\kappa \geq 0$, and the rate at which people can be identified for treatment is bounded from above, $u \leq \bar{u}$. Here $\bar{u}$ is the capacity limit for putting new patients on treatment. From the fact that $\dot{S} \geq 0$ whenever $S = 0$ (and similarly for IU, IT and R), it follows that $\boldsymbol{s} \geq 0$.

We now develop the objective function for the problem. We use a standard cost-effectiveness framework (Gold et al., 1996) and consider both the costs and benefits of disease control. The costs comprise the costs of screening, contact tracing, and health care (including both normal health care costs and the costs of treatment for the disease); the benefits are a function of quality-adjusted life years lived. We define c_s as the cost of screening one person; $c_c(\kappa)$ as the cost of contact tracing per otherwise reported index case as a function of its success, κ; and c_t as the cost of confirmatory tests and initial treatment for anyone newly identified as infected. Using simulations of different levels of contact tracing among a network of individuals, Armbruster and Brandeau (2007a) showed that contact tracing is likely to have diminishing returns to scale. This makes intuitive sense: the larger the number of contacts traced, the less likely it is that the incremental contacts traced will be infected. Thus, we assume that the function $c_c(\kappa)$ is increasing in κ (with $c_c(0) = 0$) and strictly convex. We assume that associated with each health state is an annual per person health care cost (h_S, h_{IU}, h_{IT}, and h_R, respectively) and a quality-of-life multiplier (q_S, q_{IU}, q_{IT}, and q_R, respectively). In general, we would expect $h_S = h_R < h_{IU} < h_{IT}$ and $q_S = q_R > q_{IT} > q_{IU}$. We define m as the monetary value of a healthy year of life. Then, the net value of a year of life in each health state is given by $\boldsymbol{v} := (v_S, v_{IU}, v_{IT}, v_R) := m(q_S, q_{IU}, q_{IT}, q_R) - (h_S, h_{IU}, h_{IT}, h_R)$. Given

these definitions, the total net benefit (value of life years lived minus cost of control) per unit time is

$$J(\boldsymbol{s}, \kappa, \lambda, u) := \boldsymbol{s} \cdot \boldsymbol{v} - \lambda c_s - \left(\frac{\lambda}{S + IU + R} + \gamma \right) IU c_c(\kappa) - u c_t. \tag{6}$$

The first cost term is the cost of screening (at rate λ); the second cost term is the cost of contact tracing (the cost to trace, at intensity κ, all contacts found by screening and by symptom development); and the third cost term is the cost of treating all newly identified index cases and their contacts, who are found at total rate u through screening, symptom development, and contact tracing. Our goal is to determine the mix of screening and contact tracing that will maximize this net benefit over time.

Assuming that the epidemic is initially in state $\boldsymbol{s}_0$, we want to optimize the infinite-horizon net present value of our policy with a discount rate of r. This can be written as an optimal control problem with value function $V(\boldsymbol{s})$:

$$
\begin{aligned}
V(\boldsymbol{s}_0) := \max_{\kappa(t), \lambda(t), u(t)} & \int_0^\infty e^{-rt} J(\boldsymbol{s}(t), \kappa(t), \lambda(t), u(t)) \, \mathrm{d}t \\
& \dot{\boldsymbol{s}}(t) = \boldsymbol{f}(\boldsymbol{s}(t), u(t)), \quad \boldsymbol{s}(0) = \boldsymbol{s}_0, \\
& u(t) = \left(\frac{\lambda(t)}{S(t) + IU(t) + R(t)} + \gamma \right) IU(t)(1 + \kappa(t)), \\
& 0 \leq \kappa(t), \ 0 \leq \lambda(t), \ u(t) \leq \bar{u}.
\end{aligned}
\tag{7}
$$

We solve this problem by first determining the optimal rate $u^*(t)$ at which people enter treatment (section 2.1) and then we determine the cost-minimizing mix of screening and contact tracing, $\lambda^*(t)$ and $\kappa^*(t)$ to achieve $u^*(t)$ (section 2.2).

2.1 Optimal Rate at Which People Enter Treatment

Let $\boldsymbol{\phi} := (\phi_S, \phi_{IU}, \phi_{IT}, \phi_R)$ be our vector of adjoint variables and $H(\boldsymbol{s}, \boldsymbol{\phi}, \kappa, \lambda, u) := J(\boldsymbol{s}, \kappa, \lambda, u) + \boldsymbol{\phi} \cdot \boldsymbol{f}(\boldsymbol{s}, u)$ our Hamiltonian. Along the optimal path, the adjoint variables are the marginal overall benefit of a small change in the current state: $\phi_S^*(t) = \frac{\partial}{\partial S} V(\boldsymbol{s}^*(t))$ where $V(\boldsymbol{s})$ is the value function (and similarly for IU, IT and R). Pontryagin's maximum principle states that the control variables,

$\kappa^*(t)$, $\lambda^*(t)$, and $u^*(t)$, maximize the Hamiltonian along the optimal path:

$$\kappa^*(t), \lambda^*(t), u^*(t) \in \arg\max_{\kappa, \lambda, u} H(\boldsymbol{s}^*(t), \boldsymbol{\phi}^*(t), \kappa, \lambda, u)$$

$$s.t. \ u = \left(\frac{\lambda}{S^*(t) + IU^*(t) + R^*(t)} + \gamma \right) IU^*(t)(1 + \kappa) \tag{8}$$

$$0 \le \kappa, \ 0 \le \lambda, \ u \le \bar{u}.$$

Rewriting the above equality constraint, we have

$$\lambda = (S^*(t) + IU^*(t) + R^*(t)) \left(\frac{u}{IU^*(t)(\kappa + 1)} - \gamma \right). \tag{9}$$

Substituting (9) into (8) we obtain an optimization problem with only κ and u as the decision variables. Then dropping the terms that do not depend on u and κ, and simplifying, we obtain

$$\kappa^*(t), u^*(t) \in \arg\max_{\kappa, u} u \cdot h(\boldsymbol{s}^*(t), \boldsymbol{\phi}^*(t), \kappa)$$

$$s.t. \ 0 \le \kappa, \ IU^*(t)\gamma(\kappa + 1) \le u \le \bar{u}, \tag{10}$$

where

$$h(\boldsymbol{s}, \boldsymbol{\phi}, \kappa) := \phi_{IT} - \phi_{IU} - c_t - \frac{c_c(\kappa) + c_s/p(\boldsymbol{s})}{\kappa + 1} \quad \text{and} \tag{11}$$

$$p(\boldsymbol{s}) := IU/(S + IU + R). \tag{12}$$

The term $h(\boldsymbol{s}, \boldsymbol{\phi}, \kappa)$ represents the marginal benefit of moving an infected individual into treatment, including the costs of finding and treating the individual while holding the contact tracing effort constant. The term $p(\boldsymbol{s})$ represents the disease prevalence among individuals not in treatment.

Optimizing (10) first over u, we determine that $IU^*(t)\gamma(\kappa + 1) \le \bar{u}$ is required for feasibility and that a bang-bang solution,

$$u^*(t) = \begin{cases} \bar{u} & \text{if } h(\boldsymbol{s}^*(t), \boldsymbol{\phi}^*(t), \kappa) > 0, \\ IU^*(t)\gamma(\kappa + 1) & \text{if } h(\boldsymbol{s}^*(t), \boldsymbol{\phi}^*(t), \kappa) < 0, \end{cases} \tag{13}$$

is optimal when $h(\boldsymbol{s}^*(t), \boldsymbol{\phi}^*(t), \kappa) \ne 0$ (when $h(\boldsymbol{s}^*(t), \boldsymbol{\phi}^*(t), \kappa) = 0$ then we are momentarily indif-

ferent about the size of the intervention, $u(t)$). When the marginal benefit of finding and treating a case is positive, then we treat as many people as possible, while if it is negative, then we treat as few people as possible for a given level of κ. Thus the screening level λ will either be at its maximal level or at 0. Substituting the feasibility condition and (13) into (10) we obtain an optimization problem for $\kappa^*(t)$,

$$
\kappa^*(t) \in \arg\max_{\kappa} h(\boldsymbol{s}^*(t), \boldsymbol{\phi}^*(t), \kappa) \cdot
\begin{cases}
\bar{u} & \text{if } h(\boldsymbol{s}^*(t), \boldsymbol{\phi}^*(t), \kappa) > 0, \\
IU^*(t)\gamma(\kappa+1) & \text{if } h(\boldsymbol{s}^*(t), \boldsymbol{\phi}^*(t), \kappa) < 0,
\end{cases}
\tag{14}
$$
$$
s.t. \ 0 \leq \kappa \leq \frac{\bar{u}}{IU^*(t)\gamma} - 1.
$$

The key to solving (14) is to determine the sign of

$$
\max_{\kappa} h(\boldsymbol{s}^*(t), \boldsymbol{\phi}^*(t), \kappa)
$$
$$
s.t. \ 0 \leq \kappa \leq \frac{\bar{u}}{IU^*(t)\gamma} - 1,
\tag{15}
$$

that is, to determine whether a feasible contact tracing level κ exists where the marginal benefit of finding and treating a case is positive. If (15) is negative, then $u^*(t) = IU^*(t)\gamma(\kappa^*(t) + 1)$ and

$$
\kappa^*(t) \in \arg\max_{\kappa} h(\boldsymbol{s}^*(t), \boldsymbol{\phi}^*(t), \kappa) \cdot IU^*(t)\gamma(\kappa+1)
$$
$$
s.t. \ 0 \leq \kappa \leq \frac{\bar{u}}{IU^*(t)\gamma} - 1.
\tag{16}
$$

Substituting (11) into (16) and dropping the terms that do not depend on κ, we obtain

$$
\kappa^*(t) \in \arg\max_{\kappa} \kappa(\phi_{IT}^*(t) - \phi_{IU}^*(t) - c_t)IU^*(t)\gamma - c_c(\kappa)IU^*(t)\gamma
$$
$$
s.t. \ 0 \leq \kappa \leq \frac{\bar{u}}{IU^*(t)\gamma} - 1.
\tag{17}
$$

8

Since $c_c(\kappa)$ is strictly convex,

$$\kappa^*(t) = \begin{cases} \bar{u}/(IU^*(t)\gamma) - 1 & \text{if } c_c'(\bar{u}/(IU^*(t)\gamma) - 1) < \phi_{IT}^*(t) - \phi_{IU}^*(t) - c_t, \\[2mm] c_c'^{-1}(\phi_{IT}^*(t) - \phi_{IU}^*(t) - c_t) & \text{otherwise} \\[2mm] 0 & \text{if } \phi_{IT}^*(t) - \phi_{IU}^*(t) - c_t < c_c'(0). \end{cases} \tag{18}$$

Thus if (15) is negative (i.e., $\lambda^*(t) = 0$), then κ is at a level such that the marginal benefit of getting an additional person into treatment equals the marginal cost of treating the person and finding them using an increase in the contact tracing effort, $\phi_{IT}^*(t) - \phi_{IU}^*(t) = c_t + c_c'(\kappa)$. However, when the number of untreated infections is large enough, then the treatment capacity $\bar{u}$ becomes binding and limits our contact tracing effort.

If on the other hand, (15) is nonnegative, then $\kappa^*(t)$ is the optimizer of (15) (and $\lambda^*(t)$ is given by (9)), and if in addition (15) is strictly positive, then $u^*(t) = \bar{u}$. In the next section we study (15) more closely.

2.2 Optimal Mix of Screening and Contact Tracing

We need to solve the optimization problem (15) to determine $u^*(t)$, and thus whether we need to treat as many or as few people as possible. If it is optimal to treat as many people as possible, $u^*(t) = \bar{u}$, the solution κ^o of (15) will also determine the optimal level of contact tracing ($\kappa^*(t) = \kappa^o$) and thus the optimal mix of screening and contact tracing.

Dropping the terms in (15) that do not depend on κ eliminates the dependence on $\boldsymbol{\phi}^*(t)$. We therefore focus on the equivalent problem,

$$\kappa^o(\boldsymbol{s}) \in \arg\max_{\kappa} \quad -\frac{c_c(\kappa) + c_s/p(\boldsymbol{s})}{\kappa + 1} \tag{19}$$
$$s.t. \ 0 \leq \kappa \leq \frac{\bar{u}}{IU\gamma} - 1,$$

whose solution only depends on the state. Thus the value of (15) is $h(\boldsymbol{s}^*(t), \boldsymbol{\phi}^*(t), \kappa^o(\boldsymbol{s}^*(t)))$. This optimization problem seeks the cheapest combination of screening and contact tracing for identifying a case, over the largest possible range of contact tracing levels (or equivalently assuming that the treatment rate is $\bar{u}$). Its solution, $\kappa^o(\boldsymbol{s})$, depends on the state $\boldsymbol{s}$ only through the disease

prevalence among those not in treatment, $p(\boldsymbol{s})$, and the upper bound on κ. Define the function $g(p)$ as the unique solution of the first-order condition below, for κ in terms of the prevalence p:

$$0 = c_c(\kappa) + c_s/p - (\kappa + 1)c_c'(\kappa). \tag{20}$$

Since $c_c(\kappa) - (\kappa+1)c_c'(\kappa)$ is decreasing in κ (due to the strict convexity of $c_c(\cdot)$), the solution to the first-order condition (20) is unique and corresponds to a maximum: $\kappa = g(p(\boldsymbol{s}))$ is the maximizer of the objective function of (19). Furthermore, the solution to (20) exists and $g(p)$ is defined, as long as $c_s/p \geq c_c'(0)$. When the prevalence p is above the threshold, $c_s/c_c'(0)$, then the right-hand side of (20) is negative for all $\kappa \geq 0$, and thus (20) does not have a solution. In that case, $p(\boldsymbol{s}) > c_s/c_c'(0)$, the objective function of (19) is decreasing in κ, and $\kappa = 0$ is the maximizer of (19). Thus the solution to (19) is

$$\kappa^o(\boldsymbol{s}) = \begin{cases} \bar{u}/(IU\gamma) - 1 & \text{if } \bar{u}/(IU\gamma) - 1 \leq g(p(\boldsymbol{s})), \\ g(p(\boldsymbol{s})) & \text{if } g(p(\boldsymbol{s})) < \bar{u}/(IU\gamma) - 1 \text{ and } p(\boldsymbol{s}) \leq c_s/c_c'(0), \\ 0 & \text{if } c_s/c_c'(0) < p(\boldsymbol{s}). \end{cases} \tag{21}$$

Thus, $\kappa^o(\boldsymbol{s}) = g(p(\boldsymbol{s}))$ except when the right-hand side is undefined or the treatment capacity is a limitation.

Note that when κ^o is not at its upper bound, then κ^o only depends on the incremental effectiveness of contact tracing per dollar spent (via the function $c_c(\cdot)$), the cost of screening, and the fraction of infected people among those not receiving treatment, p: κ^o does not depend on u, γ, or the cost of treatment c_t. (Of course, the control $\kappa^*(t)$ is dynamic and changes over time as $\boldsymbol{s}^*(t)$ changes.)

When the first-order condition (20) holds and the treatment capacity is not a binding limitation, then the optimal level of contact tracing is such that the incremental cost of finding an additional infected person via screening and tracing his or her contacts (yielding a total of $\kappa + 1$ additional cases) equals the cost of finding an additional infected person via contact tracing for each of $\kappa + 1$ otherwise reported index cases. Hence the relative mix of screening and contact tracing at any point in time is such that the marginal efficiency (cost per infected person found) of contact tracing

equals that of screening if possible. This makes intuitive sense: in an optimal resource allocation, the marginal efficiency of competing interventions is balanced (Mas-Collel et al., 1995).

Since the right-hand side of (20) is decreasing in κ (due to the strict convexity of $c_c(\cdot)$), it follows that $g(p)$ is decreasing in p. As the prevalence decreases, the cost of finding an infected person via screening $(c_s/p(\boldsymbol{s}))$ increases, and for prevalences above $c_s/c_c'(0)$, contact tracing is not part of the optimal strategy $(\kappa^o = 0)$. Below this threshold prevalence, the cost of finding an infected person via screening increases further, and the optimal amount of contact tracing $\kappa^o = g(p(\boldsymbol{s}))$ then starts increasing and the relative number of persons screened starts decreasing. The optimal amount of contact tracing increases as prevalence decreases until it is the only intervention $(\lambda^* = 0)$, at which point the optimal contact tracing level $\kappa^o = \bar{u}/(IU\gamma) - 1$ is limited by the treatment rate. The smaller the number of untreated infections, $p(\boldsymbol{s})$, and the greater the screening cost, c_s, the larger the role of contact tracing in identifying people for treatment. This is analogous to a result of Armbruster and Brandeau (2007b) who considered the optimal mix of screening and contact tracing to reduce the prevalence of a curable endemic disease (assuming that contact tracing has fixed efficacy (i.e., that κ is constant)): they found that contact tracing was part of the cost-minimizing solution only when the prevalence of the disease was below a given threshold.

Using (21), we can write the optimum value of (19), the cost of finding an infected individual with the optimal mix of contact tracing and screening:

$$
\frac{c_c(\kappa^o) + c_s/p(\boldsymbol{s})}{\kappa^o + 1}
$$
$$
= \begin{cases} (IU\gamma/\bar{u})(c_c(\bar{u}/(IU\gamma) - 1) + c_s/p(\boldsymbol{s})) & \text{if } \bar{u}/(IU\gamma) - 1 \leq g(p(\boldsymbol{s})), \\ c_c'(g(p(\boldsymbol{s}))) & \text{if } g(p(\boldsymbol{s})) < \bar{u}/(IU\gamma) - 1 \text{ and } p(\boldsymbol{s}) \leq c_s/c_c'(0), \quad (22) \\ c_s/p(\boldsymbol{s}) & \text{if } c_s/c_c'(0) < p(\boldsymbol{s}). \end{cases}
$$

In the first case of (22), the amount of contact tracing is limited by the treatment rate. In the second (interior) case, where κ^o is neither at its minimal nor maximal level, the per person cost of finding an untreated infection equals the marginal cost of additional contact tracing $c_c'(\kappa^o)$. In the third case $(\kappa^o = 0)$, contact tracing is not part of the optimal solution and only screening is performed. In the latter two cases, the per person cost of finding an untreated infection increases

as prevalence decreases: as the disease becomes less prevalent, it is optimal to spend more per case (at least until we hit the treatment rate limitation).

We take two approaches to further characterize the solution to the optimal control problem (7). In section 2.3 we describe how the solution can be calculated numerically, and in section 2.4 we describe an approximation that yields further analytic insight.

2.3   Numerical Solution

To numerically calculate the solution to the optimal control problem (7), we use an indirect method that solves a boundary value problem for the optimal trajectories of the state and adjoint variables $(\boldsymbol{s}(t), \boldsymbol{\phi}(t))$. Our approach is to first calculate the steady state $(\boldsymbol{s}_{eq}, \boldsymbol{\phi}_{eq})$ and then assume that the solution will eventually reach the steady state, thus giving us boundary conditions, $\boldsymbol{s}(0) = \boldsymbol{s}_0$, $\boldsymbol{s}(T) = \boldsymbol{s}_{eq}$, and $\boldsymbol{\phi}(T) = \boldsymbol{\phi}_{eq}$, where T is some large value (we use 300 years in our numerical example). Thus the boundary value problem we solve (the state and adjoint equations together with the boundary conditions) is

$$\dot{\boldsymbol{s}}(t) = \boldsymbol{f}(\boldsymbol{s}(t), u^*(t)) \tag{23}$$

$$\dot{\boldsymbol{\phi}}(t) = r\boldsymbol{\phi}(t) - \frac{\partial}{\partial \boldsymbol{s}} H(\boldsymbol{s}(t), \boldsymbol{\phi}(t), \kappa^*(t), \lambda^*(t), u^*(t)) \tag{24}$$

$$\boldsymbol{s}(0) = \boldsymbol{s}_0, \ \ \boldsymbol{s}(T) = \boldsymbol{s}_{eq}, \ \ \boldsymbol{\phi}(T) = \boldsymbol{\phi}_{eq}, \tag{25}$$

where $u^*(t)$, $\kappa^*(t)$, and $\lambda^*(t)$ can be determined from the optimal trajectories $\boldsymbol{s}^*(t)$ and $\boldsymbol{\phi}^*(t)$ as described in sections 2.1 and 2.2.

The steady state is the combination of state, adjoint, and control variables $(\boldsymbol{s}, \boldsymbol{\phi}, \kappa, \lambda, u)$ such that $\dot{\boldsymbol{s}}(t) = \dot{\boldsymbol{\phi}}(t) = 0$, (that is, the right-hand sides of (23) and (24) equal 0) and the optimality conditions from sections 2.1 and 2.2 are satisfied. The optimality conditions give two additional equations which together with $\dot{\boldsymbol{s}}(t) = 0$ and $\dot{\boldsymbol{\phi}}(t) = 0$ determine the steady state values of $\boldsymbol{s}$, $\boldsymbol{\phi}$, κ, λ, and u. There are three possibilities for these equations, depending on the sign of (15) (or equivalently the sign of $h(\boldsymbol{s}, \boldsymbol{\phi}, \kappa^o(\boldsymbol{s}))$). If the sign is negative, then the remaining two conditions are $u = IU\gamma$ and κ is as given by (18). In this case the steady state where we do as little as possible is optimal. If the sign is positive, then at the steady state we try to get as many people as possible into treatment and the remaining two conditions are $u = \bar{u}$ and $\kappa = \kappa^o(\boldsymbol{s})$, where the function $\kappa^o(\boldsymbol{s})$

12

is given by (21). In this case, the cost-effective steady state is limited by the capacity to get people into treatment and would have fewer untreated infections if $\bar{u}$ were higher. With most parameter combinations though, the sign of (15) is 0 and the cost-effective steady state is between these two extremes; in this case, the remaining two conditions are $\kappa = \kappa^o(\boldsymbol{s})$ and $h(\boldsymbol{s}, \boldsymbol{\phi}, \kappa^o(\boldsymbol{s})) = 0$.

In practice, the bang-bang nature of u (switching between two levels, $\bar{u}$ and $IU\gamma$, depending on the sign of $h(\boldsymbol{s}, \boldsymbol{\phi}, \kappa^o(\boldsymbol{s}))$) causes difficulties for boundary value problem solvers. To mitigate this we use a homotopy approach that solves a sequence of related boundary value problems that converges to our problem: in the first boundary value problem of the sequence, u^* gradually goes from $\bar{u}$ to $IU\gamma$ as $h(\boldsymbol{s}, \boldsymbol{\phi}, \kappa^o(\boldsymbol{s}))$ crosses 0, while in subsequent problems this occurs more and more rapidly until we solve our original problem. We make it easier for the boundary value problem solver to find a solution by taking the solution of one boundary value problem as the initial guess for the next problem in the sequence.

2.4   A Useful Approximation

In this section we approximate the dynamics of the optimal policy (specifically, the path along which the prevalence decreases) in order to gain analytical insight into its dependence on the parameters. We will assume that the initial prevalence is above the level that is cost-effective. On the optimal path, then, we start by getting as many people into treatment as possible, $u^*(0) = \bar{u}$. Hence the system of differential equations, $\dot{\boldsymbol{s}}(t) = \boldsymbol{f}(\boldsymbol{s}(t), \bar{u})$ describes the trajectory of the system until we reach the steady state.

Since this system of differential equations does not have an analytic solution, we make several approximations in order to obtain an analytically solvable problem. We focus on the number of untreated infections, which is given by (2), and substitute in $u = \bar{u}$. This yields

$$\dot{IU} = f_2(\boldsymbol{s}, \bar{u}) = \eta_{IU} - \mu_{IU} IU - \rho IU + \frac{S}{N}(\beta_{IU} IU + \beta_{IT} IT) - \bar{u}, \tag{26}$$

where $N(t) := S(t) + IU(t) + IT(t) + R(t)$ is the total population size. Assuming the demographics are stable, S/N is roughly constant and we obtain a linear approximation of (26):

$$\dot{IU} \approx \eta_{IU} - \mu_{IU} IU - \rho IU + \frac{S(0)}{N(0)}(\beta_{IU} IU + \beta_{IT} IT) - \bar{u}. \tag{27}$$

We assume in addition that people in treatment are not infectious; thus, $\beta_{IT} \approx 0$. Then (27) reduces to

$$\dot{IU} \approx \eta_{IU} - \bar{u} - (\mu_{IU} + \rho - \beta_{IU}S(0)/N(0))IU. \tag{28}$$

We can write (28) as a differential equation of the untreated prevalence $p(\boldsymbol{s}(t)) := IU(t)/(N(t) - IT(t))$, by exploiting the stable-demographics assumption further and assuming that the number of people not in treatment, $N(t) - IT(t)$, is constant:

$$\dot{p} \approx -\frac{\bar{u} - \eta_{IU}}{N(0) - IT(0)} - (\mu_{IU} + \rho - \beta_{IU}S(0)/N(0))p = -\alpha - \sigma p, \tag{29}$$

where we define $\alpha := (\bar{u} - \eta_{IU})/(N(0) - IT(0))$ and $\sigma := \mu_{IU} + \rho - \beta_{IU}S(0)/N(0)$.

We now take a qualitative look at the long-term dynamics of the approximate solution. Let $p_0 := p(\boldsymbol{s}_0)$ be the initial untreated prevalence and $p_{eq} := p(\boldsymbol{s}_{eq})$ be the untreated prevalence of the steady state which we calculated in the previous section. Generally, the cost-effective level is below the current prevalence, $p_0 > p_{eq}$. If $\sigma < 0$ and $p_0 > -\alpha/\sigma$, then the disease is out of control and the number of infections will grow (even when we get as many people into treatment as possible). There are two other cases to consider. In the first case, $\sigma > 0$ and $p_0 > p_{eq} \geq -\alpha/\sigma$ (if $p_{eq} = -\alpha/\sigma$, then the cost-effective steady state is limited by the treatment capacity). In the second case, $\sigma < 0$ and $-\alpha/\sigma > p_0 > p_{eq}$. In both cases the optimal trajectory hits the steady state p_{eq} and remains there. The first case is probably more common. In either case, the time until it reaches some level $p \geq p_{eq}$ is $\sigma^{-1} \log \frac{p_0 + \alpha/\sigma}{p + \alpha/\sigma}$. If the factors in the σ term dominate those in the α term, $\sigma p_0 \gg |\alpha|$, then the dynamics are similar to those of exponential decay and the time until the prevalence reaches p is roughly $\sigma^{-1} \log(p_0/p)$. In this case the time depends greatly on σ and only weakly on the target prevalence p. If, on the other hand, the factors in α dominate those in σ, $|\sigma|p_0 \ll \alpha$, then the dynamics are almost linear and the time until the prevalence reaches p is roughly $(p_0 - p)/\alpha$.

From (29) we can gain insight into the factors that affect the (approximate) change in prevalence among people not in treatment when the optimal control is applied. The decay rate σ is the sum of the death rate from chronic untreated infection plus the spontaneous recovery rate from chronic infection (these represent exits from the infected state) minus the rate of new infections per untreated infection (this represents entry into the infected state). If the death rate is higher

or the recovery rate is faster, then the prevalence associated with the optimal control will decrease more quickly; conversely, if the sufficient contact rate is higher, then the prevalence associated with the optimal control will decrease more slowly. These relationships are only approximate due to the above approximations; for example, the size of the total population and the fraction of susceptible individuals are actually changing over time (and individuals in treatment may infect a few people).

Using the methods in section 2.3, we can calculate the optimal equilibrium prevalence (among people not in treatment). We can then use (29) to estimate the optimal trajectory. This allows us to approximately determine the optimal trajectory without having to solve the full optimal control problem (7). We demonstrate the quality of this approximation in the next section.

3   Numerical Example

In this section we apply our model to a numerical example. We consider the case of hepatitis B virus (HBV) infection among Asian and Pacific Island adults in the US. Parameter values are shown in table 1.

3.1   Parameters

There are approximately 12 million Asian Americans and Native Hawaiian and Pacific Islanders above the age of 18 in the US (U.S. Census Bureau, 2007). Approximately 50% of this population is immune to hepatitis B virus (either from vaccination or previous infection which resolved itself) and 10% is chronically infected (Hutton et al., 2007). Only one-third of chronically infected individuals in the US know that they are infected (Lin et al., 2007); we assume that these individuals are receiving treatment. Thus our initial state is $S = 4.8$ million, $IU = 800,000$, $IT = 400,000$, $R = 6$ million.

We assume that the remaining life expectancy is 40 years for susceptible individuals, immune individuals, and chronically infected individuals in treatment. Approximately 25% of chronically infected individuals not in treatment develop liver disease (Hutton et al., 2007). We assume that these individuals have a remaining life expectancy of 25 years (5 years to develop the disease and 20 years with the disease). Thus $\mu_S = \mu_{IT} = \mu_R = 1/40$ and $\mu_{IU} = 0.75/40 + 0.25/25$.

We assume that entry of infected individuals into the population is due to immigration and that

these individuals are not yet in treatment. Each year approximately 22,000 legal immigrants above the age of 18 with chronic hepatitis B infection emigrate to the US from Oceania and East and South Asia (we take country level HBV prevalence data from table 3 in annex 4 of WHO (2007) and figure 37 of WHO (2001) and immigration statistics from (DHS, 2008)). We assume that the aggregate rate of entry into the population, $\sum_i \eta_i$, matches the exit rate from the initial state, $\boldsymbol{\mu} \cdot \boldsymbol{s}_0$, to keep the population roughly constant, and we assume that half of the people entering the population are immune, $\eta_R = 0.5 \sum_i \eta_i$ (as in the initial state). Thus, $\eta_S = 129,500$, $\eta_{IU} = 22,000$, $\eta_{IT} = 0$, and $\eta_R = 151,500$.

Since our focus is on hepatitis B infections that do not resolve themselves, we set $\rho = 0$. We set $\omega = 0$ because treatment does not effect a cure. In 2006 there were an estimated 46,000 new (acute) hepatitis B infections in the US (CDC, 2008a). Since almost all acute infections in the US occur in adults (see figure 14 of CDC (2008d)) and only 5% of adults develop chronic infection after acute infection (the remaining 95% recover and become immune) (CDC, 2008b), we estimate that approximately 2,300 new cases of adult chronic infection occurred in the US in 2006. Since there were 226 million adults in the US in 2006 (U.S. Census Bureau, 2006) and the incidence of hepatitis B infection in Asian Americans and Pacific Islanders is similar to that of whites (figure 15 of CDC (2008d)), this translates to a 1.0 in 100,000 annual rate of new chronic infection among Asian and Pacific Islander adults. We assume that only chronically infected individuals can infect others and that treatment reduces infectiousness by 50% (i.e., $\beta_{IT} = \beta_{IU}/2$), similar to the reduction for HIV (Quinn et al., 2000). Then we calculate $\beta_{IU} = 3 \cdot 10^{-4}$ and $\beta_{IT} = 1.5 \cdot 10^{-4}$ from the initial state $\boldsymbol{s}_0$. We assume that the rate at which symptoms develop is $\gamma = 0.1$; this equals the fraction of symptomatic infections (CDC, 2008a).

We assume that the average quality multiplier for an untreated infection q_{IU} is 0.985 (we assume that 3/4 of those individuals have no symptoms and the remaining have a quality multiplier of 0.7 in the last 5 years of a 25 year life-span, c.f., the quality multipliers in Hutton et al. (2007)). We assume that the quality multiplier for the remaining states (q_S, q_{IT}, and q_R) is 1 and that the monetary value of a QALY is $m = \$50,000$ (Owens, 1998). We further assume that the additional health care costs h_{IT} for people in the IT state average \$2000 per year (we assume that such individuals incur \$600 in annual monitoring costs for the first 20 years and then 2/3 of the individuals incur a treatment cost of \$5,000 per year during the remaining 20 years, c.f., the costs in Hutton et al.

Table 1: Parameter values for hepatitis B example

Parameter	Value
r	3%
c_s	$27
$c_c(\kappa)$	$400((\kappa + 1)^2 - 1)$
c_t	0
v_S	$50,000
v_{IU}	$49,250
v_{IT}	$48,000
v_R	$50,000
$\bar{u}$	50,000
γ	0.1
β_{IU}	$3 \cdot 10^{-4}$
β_{IT}	$1.5 \cdot 10^{-4}$
ρ	0
ω	0
μ_S	0.025
μ_{IU}	0.02875
μ_{IT}	0.025
μ_R	0.025
η_S	129,500
η_{IU}	22,000
η_{IT}	0
η_R	151,500

(2007)). Thus $v_S = \$50,000$, $v_{IU} = \$49,250$, $v_{IT} = \$48,000$, and $v_R = \$50,000$.

We use a discount factor of 3% (Gold et al., 1996), a screening cost c_s of $27 (the cost of an HBsAg test and the blood test administration in Hutton et al. (2007)), and no cost for the initial treatment, c_t, as we incorporated this into the recurring treatment cost, h_{IT}. We let $\bar{u} = 50,000$ (approximately 5 times larger than the number of clinical hepatitis B cases in the US in 2006 (CDC, 2008a)). We assume that the contact tracing cost is $c_c(\kappa) = \$400((\kappa + 1)^2 - 1)$. This results in a $1,200 cost for an effort level that results in finding one additional infection on average.

3.2 Results

The solution to the optimal control problem (7) is illustrated in figure 2. The maximum level of control $\bar{u}$ is applied until the cost-effective equilibrium is reached. The steady state is $S = 5.2$ million, $IU = 26,000$, $IT = 850,000$, and $R = 6.06$ million. The cost-effective steady-state untreated disease prevalence, $\frac{IU}{S+IU+R}$, is 0.2%, a level comparable to the prevalence of HBV in

the general US population (CDC, 2008a). This suggests that efforts should be made to reduce the untreated prevalence of hepatitis B among Asian and Pacific Islander adults in the US to a level comparable to that for the rest of the population. Using the condition in (20), the threshold prevalence p for contact tracing (this is the prevalence among those not in treatment) is $c_s/c_c'(0) = 3\%$. If this prevalence is above 3%, we perform no contact tracing; if this prevalence is below 3%, we do perform contact tracing. In addition, $\kappa_{eq} = 4.3$: at the cost-effective steady state it is optimal to trace contacts at a rate that finds an average of 4.3 untreated cases per identified index case. Finally, we note that at equilibrium, the marginal overall benefit of moving a person out of the infected untreated state and into the infected treated state exactly equals the cost of finding and treating this person, $\phi_{IT} - \phi_{IU} = \$4,000$.

Figure 2 compares the optimal policy and the approximation of its dynamics that we developed in section 2.4. The results are remarkably similar. This suggests that, if the untreated disease prevalence is higher than the optimal equilibrium level (section 2.3 describes how it can be calculated), and thus the maximum level of control must be applied until the cost-effective equilibrium prevalence is reached, one can estimate the optimal trajectory for prevalence reduction using the simple exponential decay formula given by (29). Instead of solving the full optimal control problem, one can easily calculate the equilibrium prevalence and the rate of decay using the simple formulas in sections 2.3 and 2.4.

Figure 3 shows the optimal equilibrium prevalence among untreated individuals as a function of key model parameters. As contact tracing cost increases (decreases), optimal untreated prevalence increases (decreases), but only by a relatively small amount. Equilibrium prevalence is quite sensitive to the net value of treatment ($v_{IT} = mq_{IT} - h_{IT}$): as this value decreases, optimal prevalence increases rapidly; as this value increases, optimal prevalence decreases less rapidly. A similar sensitivity occurs with the death rate among untreated individuals, μ_{IU}. Finally, optimal prevalence is mildly sensitive to the rate of external infection, η_{IU}: the higher this rate, the higher the optimal equilibrium untreated prevalence.
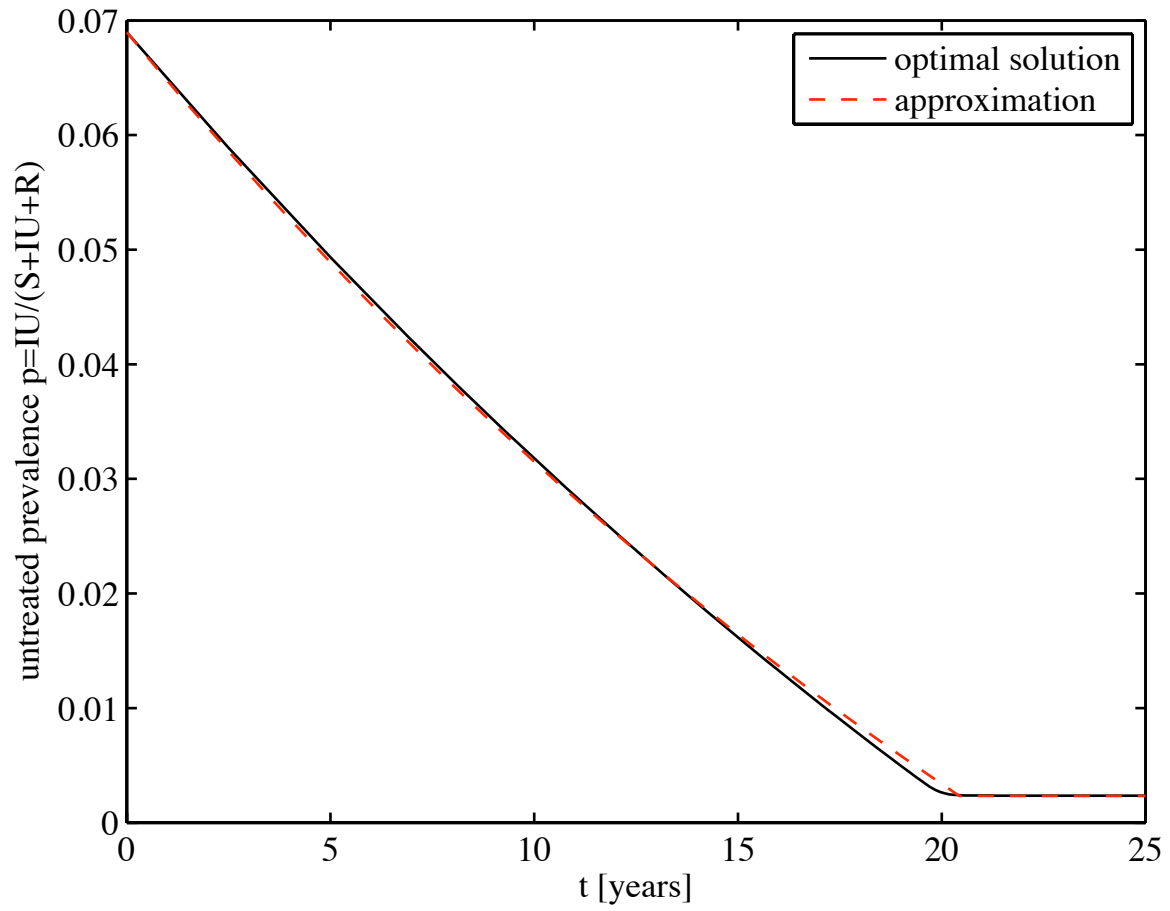
Figure 2: Solutions of the exact and approximate models: Prevalence among untreated individuals over time for the hepatitis B example
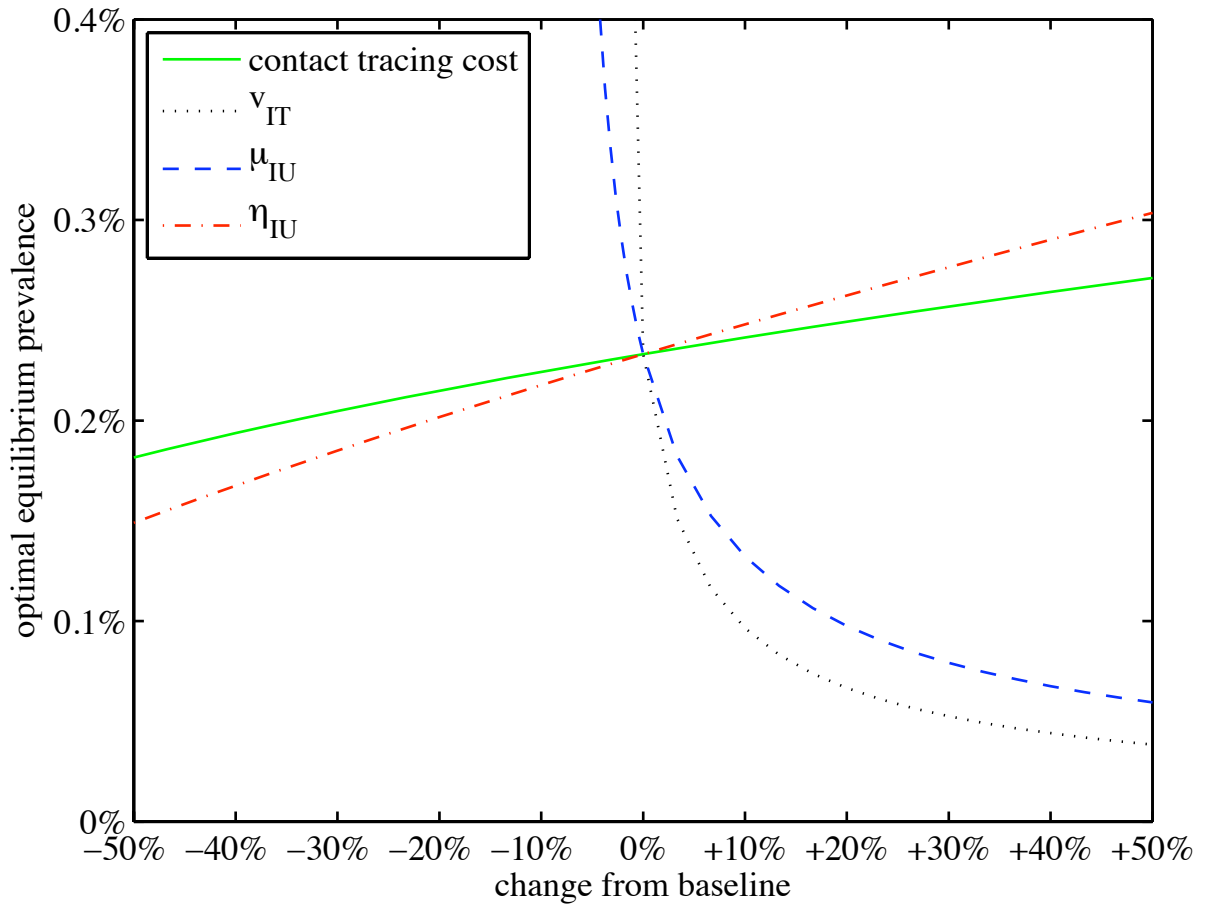
Figure 3: Sensitivity analysis of the optimal steady state: Optimal equilibrium prevalence among untreated individuals as some parameters are varied from their baseline values for the hepatitis B example

4 Discussion

We have considered the optimal mix of screening and contact tracing for a chronic viral disease. We found that when we are above the equilibrium (and moving towards the long-term target steady state) the optimal strategy is to get as many people as possible into treatment per unit time until the incremental health benefits balance the cost of controlling the disease. We further found that the optimal contact tracing effort $\kappa^*(t)$ is a simple function of the current prevalence among untreated individuals, and under the optimal strategy, the marginal effectiveness of contact tracing equals that of screening. Additionally, as the disease becomes less prevalent, it is optimal to spend more per case identified. We showed how to calculate the equilibrium cost-effective state of the system (e.g., number of people in treatment) and developed a simple approximation of the rate at which the system moves toward this equilibrium. Together, the equilibrium calculation and the approximation allow us to obtain a close characterization of the optimal solution without having to solve the full optimal control problem.

Simple compartmental models such as the one we use are standard and analytically tractable but are not completely realistic. Two well-known shortcomings of such models are the assumptions of random mixing and instantaneous transitions. One could ameliorate these shortcomings by constructing a more complicated model with more compartments: for example, one could include compartments for different risk groups (with different mixing rates), different genders in the case of diseases that can be sexually transmitted, and different disease states (such as an infected asymptomatic state). To calculate numerical solutions to optimal control problems arising from more complicated models it may be helpful to use specialized software such as the OCMat toolbox for MATLAB (Grass et al., 2008). Another alternative is to use a stochastic network simulation model to capture more details of mixing and infection transmission (e.g., Armbruster and Brandeau (2007a); Enns et al. (2009); Halloran et al. (2008); Morris and Kretzschmar (1997)). However, in both cases, analytical results may be difficult to develop. Simple models such as ours also ignore co-infection dynamics: in some cases, infection with one disease (e.g., HIV) may facilitate infection with another (e.g., tuberculosis) (Long et al., 2008; Porco et al., 2001). Additionally, our model assumes that the rate of infected individuals developing immunity (and entering compartment R) is proportional to the number of treated and untreated infections, IT and IU. A more accurate model

could capture the possibility that a certain fraction of infections will quickly resolve themselves.

We have considered disease control via treatment but for some chronic infectious diseases, such as hepatitis B, a vaccine exists. A worthwhile extension of this work would consider the optimal allocation of resources between vaccination and treatment. We have modeled contact tracing and screening as two means of finding infected individuals. One could extend our model to also consider health-awareness campaigns as an additional intervention (such campaigns are common for HIV and hepatitis B, for example). The effect of including additional interventions is not additive because any particular infection can only be prevented, detected, or treated once. Thus, one would ideally examine the allocation of resources not only across time but also across a portfolio of possible interventions (e.g., Brandeau et al. (2003); Zaric (2003); Zaric and Brandeau (2001, 2002)).

Our analysis assumed dynamic control: we determined the optimal levels of screening and contact tracing over time to control a chronic viral disease. Dynamic control of this type could be extended to behavior change interventions. Typical models assume static levels of behavior change (e.g., Brandeau et al. (2003); Zaric (2003); Zaric and Brandeau (2001, 2002)). However, studies of risky behavior related to HIV suggest that risky behavior decreases as perceived risk increases (Hethcote et al., 1991; Velasco-Hernández et al., 1996; Zeiler, 2008). It may be possible to extend such models using our ideas to the case in which behavior change depends on prevalence.

Although our model is simple and the mathematics at times complicated, our analysis yields four insights. First, the optimal strategy is to get infected individuals into treatment as fast as possible when the disease prevalence is above the cost-effective equilibrium disease prevalence. Second, as one reduces the disease prevalence by moving people into treatment (which decreases the chance that they will infect others), one should increase the level of contact tracing to compensate for the decreased effectiveness of screening. Third, as disease prevalence decreases, it is optimal to spend more per case identified. Fourth, the relative mix of screening and contact tracing at any level of disease prevalence is such that the marginal efficiency of contact tracing (cost per infected person found) equals that of screening if possible (e.g., when capacity limitations are not binding).

5 Acknowledgments

References

AB 158. Medi-Cal: benefits for nondisabled persons infected with chronic hepatitis B. 2008. URL http://www.aroundthecapitol.com/Bills/AB_158. Proposed California legislation.

B. Armbruster and M. L. Brandeau. Contact tracing to control infectious disease: When enough is enough. *Health Care Management Science*, 10(4):341–355, 2007a. URL http://users.iems.northwestern.edu/~armbruster/HCMS2007.pdf.

B. Armbruster and M. L. Brandeau. Optimal mix of screening and contact tracing for endemic diseases. *Mathematical Biosciences*, 209(2):386–402, 2007b.

M. L. Brandeau, G. S. Zaric, and A. Richter. Resource allocation for control of infectious diseases in multiple independent populations: beyond cost-effectiveness analysis. *J Health Econ*, 22(4):575–598, 2003. URL http://dx.doi.org/10.1016/S0167-6296(03)00043-2.

E. A. Brooks, L. F. Lacey, S. L. Payne, and D. W. Miller. Economic evaluation of lamivudine compared with interferon-alpha in the treatment of chronic hepatitis B in the United States. *Am J Manag Care*, 7(7):677–682, 2001.

CDC. Disease burden from hepatitis A, B, and C in the United States. 2008a. URL http://www.cdc.gov/hepatitis/Statistics.htm.

CDC. Hepatitis B FAQs for health professionals. 2008b. URL http://www.cdc.gov/hepatitis/HBV/HBVfaq.htm.

CDC. Prevention of specific infectious diseases. In *CDC Health Information for International Travel 2008*, chapter 4. 2008c. URL http://wwwn.cdc.gov/travel/contentYellowBook.aspx.

CDC. Surveillance for acute viral hepatitis — United States, 2006. *MMWR*, 57, 2008d. URL http://www.cdc.gov/mmwr/PDF/ss/ss5702.pdf.

DHS. Profiles on legal permanent residents: 2007. 2008. URL http://www.dhs.gov/ximgtn/statistics/data/DSLPR07c.shtm.

E. A. Enns, M. L. Brandeau, T. K. Igeme, and E. Bendavid. Assessing effectiveness and cost-effectiveness of concurrency reduction for HIV prevention. 2009. Working Paper, Stanford University.

K. A. Freedberg, E. Losina, M. C. Weinstein, A. D. Paltiel, C. J. Cohen, G. R. Seage, D. E. Craven, H. Zhang, A. D. Kimmel, and S. J. Goldie. The cost effectiveness of combination antiretroviral therapy for HIV disease. *N Engl J Med*, 344(11):824–831, 2001. URL http://content.nejm.org/cgi/content/abstract/344/11/824.

M. Gold, J. Siegel, L. Russell, and M. Weinstein. *Cost-effectiveness in health and medicine*. Oxford University Press, 1996.

D. Grass, J. P. Caulkins, G. Feichtinger, and D. A. Behrens. *Optimal Control of Nonlinear Processes: With Applications in Drugs, Corruption and Terror*. Springer, 2008.

M. E. Halloran, N. M. Ferguson, S. Eubank, I. M. Longini, D. A. T. Cummings, B. Lewis, S. Xu, C. Fraser, A. Vullikanti, T. C. Germann, D. Wagener, R. Beckman, K. Kadau, C. Barrett, C. A. Macken, D. S. Burke, and P. Cooley. Modeling targeted layered containment of an influenza pandemic in the United States. *Proc Natl Acad Sci USA*, 105(12):4639–4644, 2008. URL http://dx.doi.org/10.1073/pnas.0706849105.

H. W. Hethcote, J. W. V. Ark, and J. M. Karon. A simulation model of AIDS in San Francisco: II. Simulations, therapy, and sensitivity analysis. *Mathematical Biosciences*, 106(2):223 – 247, 1991. URL http://www.sciencedirect.com/science/article/B6VHX-45F52P8-KP/2/4a113da0514d6c361439fdcdaca2594d.

D. W. Hutton, D. Tan, S. K. So, and M. L. Brandeau. Cost-effectiveness of screening and vaccinating Asian and Pacific Islander adults for hepatitis B. *Ann Intern Med*, 147(7):460–469, 2007. URL http://www.annals.org/cgi/content/abstract/147/7/460.

F. Kanwal, I. M. Gralnek, P. Martin, G. S. Dulai, M. Farid, and B. M. R. Spiegel. Treatment

alternatives for chronic hepatitis B virus infection: a cost-effectiveness analysis. *Ann Intern Med*, 142(10):821–831, 2005.

S. Y. Lin, E. T. Chang, and S. K. So. Why we should routinely screen Asian American adults for hepatitis B: a cross-sectional study of Asians in California. *Hepatology*, 46(4):1034–1040, 2007. URL http://dx.doi.org/10.1002/hep.21784.

E. Long, N. Vaidya, and M. Brandeau. Controlling co-epidemics: Analysis of HIV and tuberculosis infection dynamics. *Operations Research*, 56:1366–1381, 2008.

A. Mas-Collel, M. D. Whinston, and J. R. Green. *Microeconomic Theory*, chapter 3. Oxford University Press, 1995.

M. Morris and M. Kretzschmar. Concurrent partnerships and the spread of HIV. *AIDS*, 11(5):641–648, 1997.

D. K. Owens. Interpretation of cost-effectiveness analyses. *J Gen Intern Med*, 13:716–717, 1998.

T. C. Porco, P. M. Small, and S. M. Blower. Amplification dynamics: predicting the effect of HIV on tuberculosis outbreaks. *J Acquir Immune Defic Syndr*, 28(5):437–444, 2001.

T. Quinn, M. Wawer, N. Sewankambo, D. Serwadda, C. Li, F. Wabwire-Mangen, M. Meehan, T. Lutalo, and R. Gray. Viral load and heterosexual transmission of human immunodeficiency virus type 1. *New England Journal of Medicine*, 342(13):921–929, 2000.

C. Revelle, F. Feldmann, and W. Lynn. An optimization model of tuberculosis epidemiology. *Management Science*, 16(4):B190–B211, 1969. URL http://www.jstor.org/stable/2628798.

J. L. Sanders. Quantitative guidelines for communicable disease control programs. *Biometrics*, 27(4):883–893, 1971. URL http://www.jstor.org/stable/2528825.

S. P. Sethi. Quantitative guidelines for communicable disease control program: A complete synthesis. *Biometrics*, 30(4):681–691, 1974. URL http://www.jstor.org/stable/2529232.

J. Shepherd, J. Jones, A. Takeda, P. Davidson, and A. Price. Adefovir dipivoxil and pegylated interferon alfa-2a for the treatment of chronic hepatitis B: a systematic review and economic evaluation. *Health Technol Assess*, 10(28):iii–iv, xi–xiv, 1–183, 2006.

H. M. Taylor. Some models in epidemic control. *Mathematical Biosciences*, 3:383–398, 1968.

UNAIDS. Resource needs for an expanded response to AIDS in low- and middle-income countries. 2005. URL http://www.unaids.org/html/pub/publications/irc-pub06/resourceneedsreport_24jun05_en_pdf.pdf.

UNAIDS. 2006 report on the global AIDS epidemics. 2006.

U.S. Census Bureau. General demographic characteristics. 2006.

U.S. Census Bureau. 2007 American community survey. 2007.

J. X. Velasco-Hernández, F. Brauer, and C. Castillo-Chavez. Effects of treatment and prevalence-dependent recruitment on the dynamics of a fatal disease. *IMA J Math Appl Med Biol*, 13(3):175–192, 1996.

WHO. Health situation and trends assessment: Health situation in the South-East Asia region, 1998-2000. 2001. URL http://www.searo.who.int/en/Section1243/Section1382/Section1386/Section1898.htm.

WHO. Western Pacific regional plan for hepatitis B control through immunization. 2007. URL http://www.wpro.who.int/NR/rdonlyres/CEDD5D4E-71BE-49F4-AEEC-1384751598EE/0/POA_HepB.pdf.

G. S. Zaric. Resource allocation for control of infectious disease epidemics. *Comments on Theoretical Biology*, 8:475–496, 2003.

G. S. Zaric and M. L. Brandeau. Resource allocation for epidemic control over short time horizons. *Math Biosci*, 171:33–58, 2001.

G. S. Zaric and M. L. Brandeau. Dynamic resource allocation for epidemic control in multiple populations. *IMA Journal of Mathematics Applied to Medicine and Biology*, 19:235–255, 2002.

I. Zeiler. *Optimal Dynamic Control with DNSS Curves: Multiple Equilibria in Epidemic Models of HIV/AIDS and Illicit Drug Use.* Ph.D. thesis, Vienna University of Technology, 2008.