

Original Contribution

Economic Incentives in the Socially Optimal Management of Infectious Disease: When R_0 is Not Enough

B. R. Morin,¹ A. P. Kinzig,² S. A. Levin,³ and C. A. Perrings²

¹Mathematics and Statistics Department, Vassar College, 124 Raymond Avenue, Box 257, Poughkeepsie, NY 12604-0257

²School of Life Sciences, Arizona State University, PO Box 4501, Tempe, AZ 85287

³Department of Ecology and Evolutionary Biology, Princeton University, 106A Guyot Hall, Princeton, NJ 08544-1003

Abstract: Does society benefit from encouraging or discouraging private infectious disease-risk mitigation? Private individuals routinely mitigate infectious disease risks through the adoption of a range of precautions, from vaccination to changes in their contact with others. Such precautions have epidemiological consequences. Private disease-risk mitigation generally reduces both peak prevalence of symptomatic infection and the number of people who fall ill. At the same time, however, it can prolong an epidemic. A reduction in prevalence is socially beneficial. Prolongation of an epidemic is not. We find that for a large class of infectious diseases, private risk mitigation is socially suboptimal—either too low or too high. The social optimum requires either more or less private mitigation. Since private mitigation effort depends on the cost of mitigation and the cost of illness, interventions that change either of these costs may be used to alter mitigation decisions. We model the potential for instruments that affect the cost of illness to yield net social benefits. We find that where a disease is not very infectious or the duration of illness is short, it may be socially optimal to promote private mitigation effort by increasing the cost of illness. By contrast, where a disease is highly infectious or long lasting, it may be optimal to discourage private mitigation by reducing the cost of disease. Society would prefer a shorter, more intense, epidemic to a longer, less intense epidemic. There is, however, a region in parameter space where the relationship is more complicated. For moderately infectious diseases with medium infectious periods, the social optimum depends on interactions between prevalence and duration. Basic reproduction numbers are not sufficient to predict the social optimum.

Keywords: Disease risk mitigation, Social optimum, Infectious disease, Economic incentives

INTRODUCTION

During an epidemic, there are several options available to private individuals interested in managing infectious disease risk. These include self-quarantine, vaccination, voluntary travel restrictions, contact reduction, prophylaxis,

and preferential mixing (Philipson 1996; Geoffard and Philipson 1997; Philipson 2000; Taylor and Ampt 2003; Brownstein et al. 2006; Ferguson et al. 2006; Germann et al. 2006; Hunter 2007; Rao et al. 2009; Curtis et al. 2011; Fenichel et al. 2013). The choice of which measures to use frequently depends on costs—the cost of illness and the cost of risk mitigation. The cost of illness is determined by the likelihood that people will become infected if they make

Published online: September 29, 2017

Correspondence to: B. R. Morin, e-mail: bmorin@vassar.edu

contact with an infectious individual, together with the cost of treatment and loss of earnings if they do become infected. The cost of risk mitigation is the resources sacrificed to reduce risk by some amount. Diseases that are perceived to be highly infectious and life-threatening warrant greater private risk mitigation effort than diseases that are perceived to be only moderately infectious and with few symptoms (Sanchez 2001; Leroy et al. 2005; Team 2014). Diseases that are novel, and therefore uncertain in their effects, warrant greater private risk mitigation effort than diseases that are more familiar (Taylor et al. 2001; Presanis et al. 2009; Rubin et al. 2009). It follows that private risk mitigation efforts might be expected to evolve over the course of an epidemic as information improves. A number of studies have, for example, tracked changes in the efforts made by private individuals to reduce disease risk during the 2009 A/H1N1 epidemic (Fenichel et al. 2013; Bayham et al. 2015; Springborn et al. 2015).

The problem addressed in this paper is that disease risk mitigation by one individual affects the well-being of both that individual and others. People vaccinating themselves against seasonal influenza, for example, reduce the likelihood that they will become infected, but also reduce the likelihood that they will become infectious. Vaccination, in such cases, is said to be an impure public good. It confers benefits on the individual, but also on those with whom the individual interacts (Sandler et al. 2002). On the other hand, people avoiding infection now by sequestering themselves may simply delay the moment at which they become infected and infectious and hence may unwittingly prolong an epidemic (Fenichel et al. 2011). Private disease-risk mitigation, in this case, is an impure public bad. In providing benefits to the individual, it imposes costs on society. Because people cannot capture any wider benefits they confer, and are not confronted by any wider costs they impose, they have little incentive to take those benefits and costs into account when making their decisions. The public good (bad) is therefore under-provided (over-provided).

The general response of public health authorities to the under-provision of private disease risk mitigation is the adoption of social distancing measures that include quarantines, travel restrictions, and school closures. These are not, however, the only options available. Since private disease risk mitigation efforts reflect the costs of illness and illness avoidance, interventions that change those costs are also among the options available to public health authorities (Perrings et al. 2014). In this paper, we consider policy instruments that close the gap between the socially and

privately optimal levels of disease risk mitigation. Such policy instruments align the socially and privately efficient solutions by aligning the social and private cost of disease-risk mitigation. We assume private disease-risk mitigation strategies similar to Morin et al. (2015) and consider a control that operates on the cost of illness. This has parallels with the health belief model originally developed in the 1950s, which supposes that private health-risk mitigation depends on beliefs about the likelihood and consequences of contracting disease, the efficacy and costs of the proposed action, and the decision makers' capacity to implement the action (Rosenstock 1974).

We note that risk is used here in the economic sense. It is the product of the probability that an event will occur and the cost of the event if it does occur. The two most common approaches to the management of disease risk are mitigation and adaptation. Mitigation implies an action that reduces the probability that the individual will fall ill. Adaptation implies an action that has no effect on the probability that the individual will fall ill, but reduces the cost of illness. The probability that an infection will occur, $c\beta P_{SI}(\phi(t))$, depends on three things: contact volume c , the probability that contact results in infection β , and the probability that contacts are with infectious individuals $P_{SI}(\phi(t))$, a function of mitigation effort $\phi(t)$. We treat the first two of these parametrically, although we acknowledge that they may be affected by, respectively, strategies to quarantine infectious individuals or to reduce activity levels (Fenichel et al. 2011), and strategies such as prophylaxis or vaccination (Weycker et al. 2005). We focus instead on the probability that contact will be made with infectious individuals. It was shown in Fenichel et al. (2013) that despite travel restrictions during the H1N1 swine flu outbreak, individuals did not cancel travel plans (reduce contact volume) so much as alter the timing of their travel. In previous work, we have shown the equivalence of private strategies targeting c or P_{SI} (Morin et al. 2014) and have identified the social costs and benefits of private strategies that alter disease trajectories (Morin et al. 2015). In this paper, we study the impact of changes in the private cost of illness on mitigation effort, and show how diseases with the same R_0 can trigger very different interventions, depending on which parameter drives the basic reproduction number.

The paper is structured as follows. We first describe a model of preferential mixing based on observable disease states. We then detail two economic decision models, one private and one public, and show how they are coupled. Since the models are not tractable enough to yield analyt-

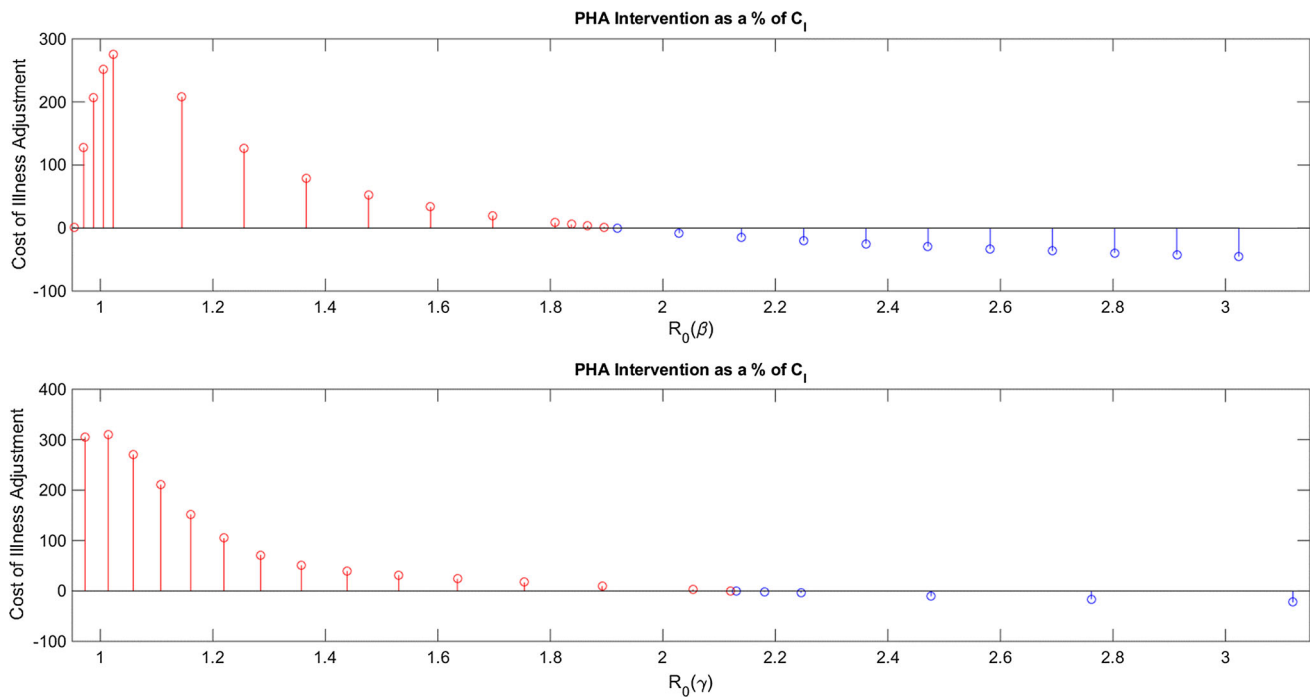


Fig. 1. PHA intervention as a function of R_0 . In the top panel, R_0 is varied as a function of β and on the bottom as a function of γ .

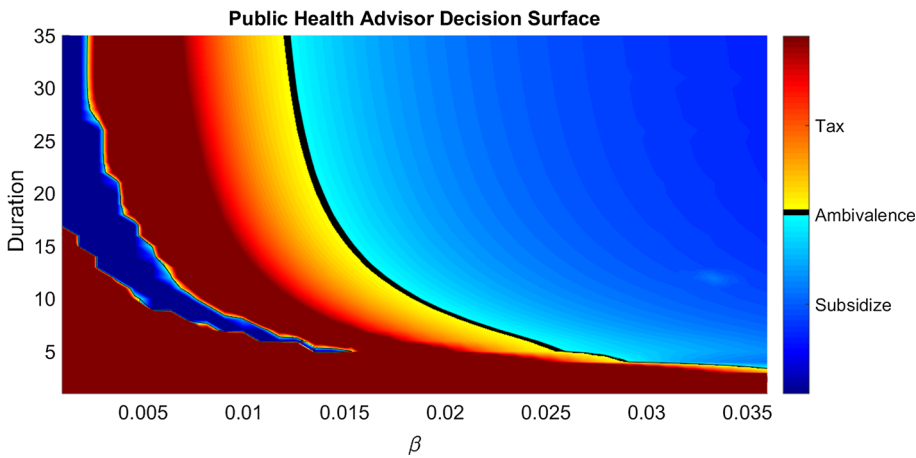


Fig. 2. Socially optimal intervention as a function of the duration and infectiousness of disease.

ical solutions, we solve both numerically in order to show the epidemiological effects of public taxes or subsidies on private cost of illness. This is followed by a discussion.

A MODEL OF PREFERENTIAL MIXING

Our modeling approach builds on existing affinity-based mixing compartment models where compartments represent different disease states (Busenberg and Castillo-Chavez 1989; Blythe et al. 1991; Castillo-Chavez et al. 1991; Fenchel et al. 2011; Morin et al. 2015). We suppose that

individuals mix preferentially, conditional on their own disease state and the (observable) disease states of others. The only information available to individuals is the expression of symptoms in either themselves or others. As in (Morin et al. 2014), the resulting mixing strategy depends on the relative costs of illness and illness avoidance. This framework has been shown to provide the most mathematically general solution to the problem of *who mixes with whom* under the assumptions of symmetric contacts (Blythe et al. 1991).

In what follows, individuals are defined only by their health state, although they could just as easily be grouped

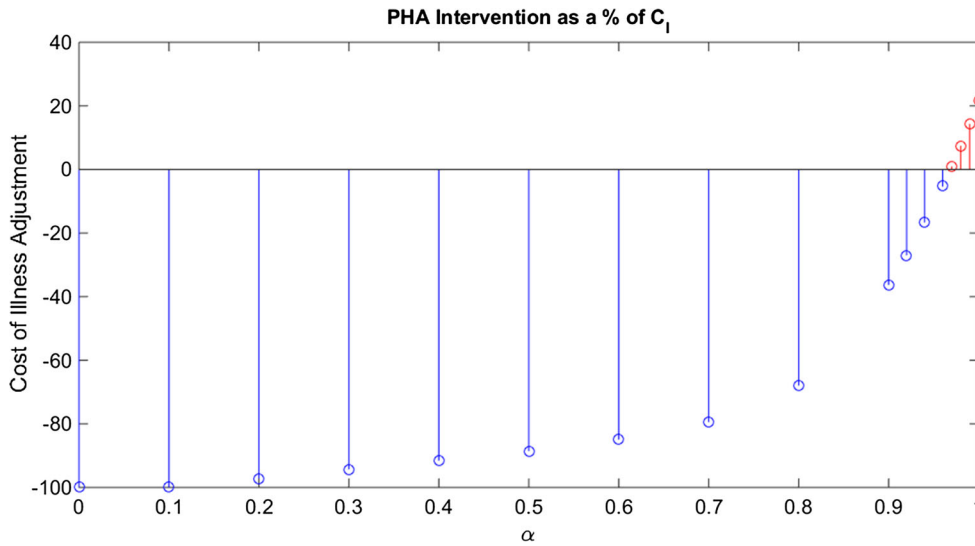


Fig. 3. PHA intervention as a function of the rate of cost recovery.

according to various shared attributes such as economic status, cultural or ethnic identity, geographical location, age, or disease awareness. We suppose that all individuals who do not show symptoms—because they have either not become ill, are asymptotically ill or have recovered from illness—are treated equally by others. The use of the affinity framework allows three factors to determine the volume of contact between groups of like individuals: (1) the size of each group, (2) the nominal activity level or disease-free contact rate of each group and 3) the relative affinity/disaffinity between groups. In what follows, we treat the affinity/disaffinity between groups as endogenous to the epidemiological system.

Susceptible individuals, or at least individuals who believe themselves to be susceptible, choose the people with whom they mix in order to alter the probability that they will encounter infectious individuals and subsequently become ill themselves. In the most general case, people who believe themselves to be susceptible at some time are taken to include all those who have been free of symptoms up to that time. This includes those who are actually susceptible, those who are asymptotically infectious and those who are recovered but have never had symptoms. We hold the nominal level of activity (the contact rate) constant throughout the course of the epidemic and take it to be equal for all individuals. This makes it possible to consider only the effect of changes in mixing preferences. See Fenichel et al. (2011) for a treatment that selects the volume of contacts, and (Morin et al. 2014) for an analysis of the conditions under which choice of contact rates and avoidance effort are equivalent strategies. The main dif-

ference between strategies is that while choice of contact rates allows complete isolation, choice of avoidance effort does not.

To illustrate the approach, we first focus on a susceptible-infectious-recovered (SIR) model—but note that we will be reporting results for a range of other models including susceptible-exposed or latent-infectious-recovered (SEIR), one-path and two-path susceptible-asymptomatic infectious-infectious-recovered (one-path SAIR, two-path SAIR). In the SIR case, only susceptible individuals are free of symptoms. The disease dynamics are summarized in three differential equations:

$$\begin{aligned}\frac{dS(t)}{dt} &= -c\beta S(t)P_{SI}(\phi(t)) \\ \frac{dI(t)}{dt} &= c\beta S(t)P_{SI}(\phi(t)) - \gamma I(t) \\ \frac{dR(t)}{dt} &= \gamma I(t)\end{aligned}\quad (1)$$

As is standard with the SIR model, we let c be the nominal contact volume of all individuals. $P_{SI}(\phi(t))$ is the conditional probability that a contact made by a susceptible individual, committing $\phi(t)$ effort to avoiding infection at time t , is with an infectious individual, and γ is the rate at which an individual recovers and becomes immune. $I(t)$ and $R(t)$ are, respectively, the numbers of infectious and recovered (immune) individuals.

The conditional probability that an individual in the i th disease state encounters an individual in the j th disease state is given by the elements of a time-dependent mixing matrix, $P(t) = (P_{ij}(t))$, that is taken to satisfy three axioms

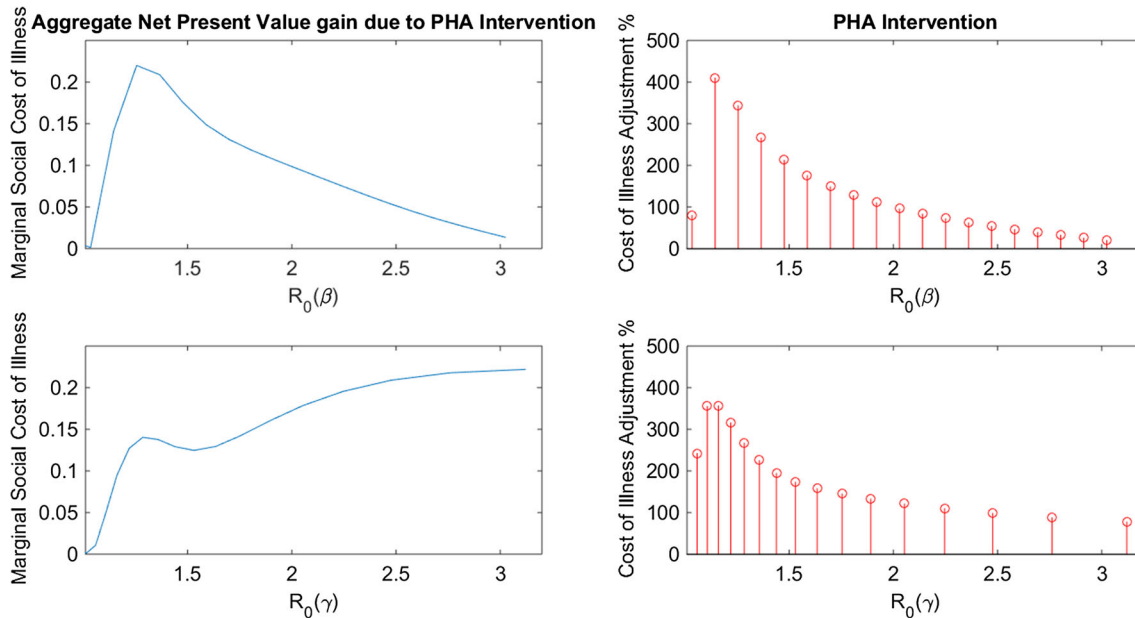


Fig. 4. On the left is the aggregate net present value that PHA intervention induces and on the right is the actual amount of PHA intervention. Each are a function of R_0 which is depicted as a function of β on the top two panels and as a function of γ on the bottom two panels.

(Busenberg and Castillo-Chavez 1989; Blythe et al. 1991; Castillo-Chavez et al. 1991):

1. $0 \leq P_{ij} \leq 1$, for all $i, j \in \{S, E, A, I, R\}$
2. $\sum_{j \in \{S, E, A, I, R\}} P_{ij} = 1$, for all $i \in \{S, E, A, I, R\}$,
3. $i(t)P_{ij} = j(t)P_{ji}$, for all $i, j \in \{S, E, A, I, R\}$.

These three axioms imply that, collectively, susceptible individuals have the same expectation of encountering infected individuals as infected individuals have of encountering susceptible individuals. It has been shown that the unique solution to these mixing axioms is given by

$$P_{ij} = j(t) \left[\frac{M_i M_j}{M} + \phi_{ij}(t) \right],$$

where

$$M_i(t) = 1 - S(t)\phi_{S_i}(t) - I(t)\phi_{I_i}(t) - R(t)\phi_{R_i}(t),$$

$$M(t) = S(t)M_S(t) + I(t)M_I(t) + R(t)M_R(t),$$

and $\Phi(t) = (\phi_{ij}(t))$ is a symmetric affinity matrix, in this case 3×3 .

The main difference between the approach here and previous use of affinity models is that we treat the elements $\phi_{ij}(t)$ of the affinity matrix as private choice variables. They may be interpreted as the effort that the representative individual in disease state i makes to avoid individuals in disease state j , if $\phi_{ij}(t) < 0$, or to associate

with individuals in disease state j , if $\phi_{ij}(t) > 0$. If the representative individual in every disease state i makes no effort to avoid individuals in disease state j , and vice versa, then $\phi_{ij}(t) = 0$. We then have classic proportionate mixing. We take zero elements in the affinity matrix to be evidence of ‘avoidance-neutrality.’ That is, they show the representative individual to be neutral about a pairing event with someone from another disease class. By contrast, negative (positive) elements reflect the desire of the representative individual in one disease state to avoid (seek out) individuals in other disease state. Avoidance, $\phi_{ij}(t) < 0$, can result from individuals in both states wishing to avoid one another; individuals in one state wishing to avoid individuals in other states who may be neutral to the pairing; or individuals in one state wishing to avoid individuals in other states more strongly than those individuals favor the pairing. Similarly, engagement results from individuals in both states favoring the pairing; individuals in one state seeking out individuals in other states who may be neutral to the pairing; or individuals in one state wishing to engage with individuals in other states more strongly than those individuals wish to avoid the pairing. This is a similar measure to that used in models of assortative mating (Karlin 1979) and selective mixing (Hyman and Li 1997) and is a form of a contact kernel (Gurarie and Ovaskainen 2013).

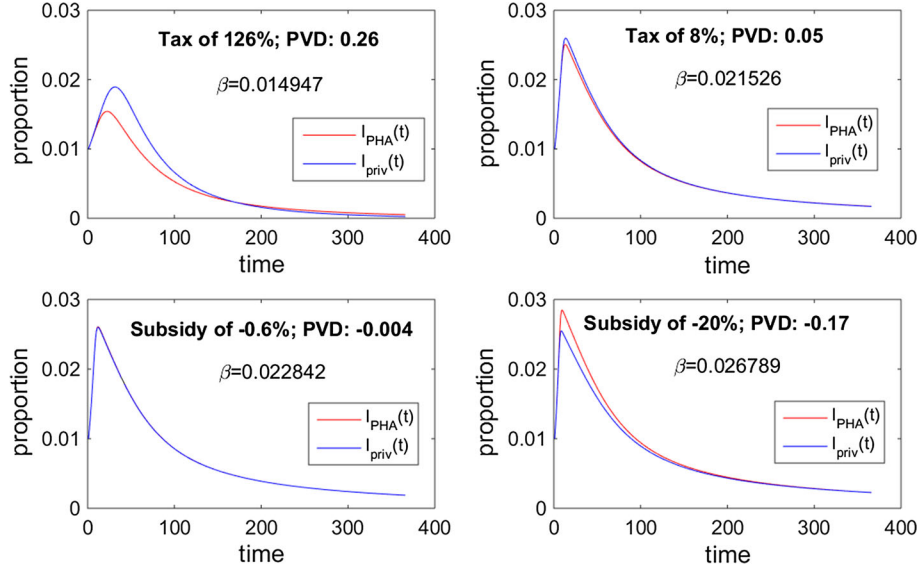


Fig. 5. Each figure shows the prevalence with public health authority intervention, $I_{\text{PHA}}(t)$, and with only the private mitigation effort, $I_{\text{priv}}(t)$, for various β . The present value difference (PVD) in these curves is shown along with the intervention level of the public health authority. Positive PVD, corresponding to an early reduction in incidence that is not offset by the future “fatter” tail, is analogous to a tax on the cost of illness. Negative PVD was always found to include an early increase infection and is analogous to a subsidy on the cost of illness. The magnitude of intervention is directly related to the magnitude of the PVD.

The elements of the affinity matrix, $\phi_{ij}(t)$, describe what the representative individual in each health state wants. What they actually get depends both on the preferences of others in the population and on the relative size of all health classes. More particularly, the elements of the mixing matrix $P = (P_{ij})$ depend both on the proportion of the population in each disease state and on the affinity matrix. They describe the conditional probabilities that an individual of disease state i contacts someone in disease state j .

In what follows, we focus on individuals who believe themselves to be susceptible (who have been symptom free up to that point) and assume that they maximize the net present value of the contacts they make, taking into account the cost of illness and illness avoidance, by choosing the effort to commit to preferential mixing: the elements of $\Phi(t)$. Formally, the decision problem for individuals who believe themselves to be susceptible, collectively labeled X , is to choose the level of mitigation effort, $\phi_{XI}(t)$, in order to maximize the difference between the benefit of not being symptomatic, B , and the cost of mitigation effort, $C(\phi_{XI}(t))$, given the weight they place on future well-being (the discount rate δ) and their planning horizon, T . If susceptible individuals are averse to mixing with symp-

tomatic (infectious or otherwise) individuals in the SIR model, and if all others are neutral, $\Phi(t)$ has the structure:

$$\Phi(t) = \begin{pmatrix} 0 & \phi_{SI}(t) & 0 \\ \phi_{SI}(t) & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix},$$

in which 0 represents neutrality and $\phi_{SI}(t) < 0$ represents the effort susceptible individuals make to avoid mixing with infectious individuals at time t (Morin et al. 2014). This defines

$$M_S = 1 - \phi_{SI}(t)I(t),$$

$$M_I = 1 - \phi_{SI}(t)S(t),$$

$$M_R = 1,$$

$$M = 1 - S(t)M_S - I(t)M_I - R(t)M_R = 1 - 2\phi_{SI}(t)S(t)I(t).$$

We may then write the mixing matrix of conditional probabilities as

$$P = \begin{pmatrix} S(t)\frac{M_S^2}{M} & I(t)\left[\frac{M_S M_I}{M} + \phi_{SI}(t)\right] & R(t)\frac{M_S}{M} \\ S(t)\left[\frac{M_S M_I}{M} + \phi_{SI}(t)\right] & I(t)\frac{M_I^2}{M} & R(t)\frac{M_I}{M} \\ S(t)\frac{M_S}{M} & I(t)\frac{M_I}{M} & \frac{R(t)}{M} \end{pmatrix}.$$

Table 1. List of Epidemiological States and the Instantaneous Net Benefits of Contact in Those States, $V_i(\phi_{ij}(t))$.

	Description	Will mix selectively	Net benefit
S	Susceptible	Yes	$B_S - C_s(\phi_{SI}(t))$
E	Latently infected: asymptomatic and noninfectious	Yes	$B_E - C_s(\phi_{SI}(t))$
A	Asymptomatically infected	Yes	$B_A - C_s(\phi_{SI}(t))$
R_A	Recovered from asymptomatic infection, immune	Yes	$B_{R_A} - C_s(\phi_{SI}(t))$
I	Symptomatically infectious	No	$B_I - C_I(I(t))$
R	Recovered from symptomatic infection, immune	No	B_R

Susceptible, exposed, asymptotically infectious, and recovered (from asymptomatic infection) individuals all choose to mitigate risk as if they were susceptible and thus carry the cost of mitigation.

As shown in Morin et al. (2015), we note that mitigation effort is restricted to the range:

$$\phi_{SI}(t) \in \left[\frac{R(t) - \sqrt{(R(t))^2 + 4S(t)I(t)}}{2S(t)I(t)}, 0 \right],$$

with proportionate mixing resulting from applying 0 effort. $P_{SI}(t) = I(t)$ is the maximum probability of contact and private quarantine of infectious individuals at the left endpoint, $P_{SI}(t) = 0$ is the minimum probability of contact. From an economic perspective, private quarantine implies that the marginal cost of illness is greater than the marginal cost of illness avoidance for all levels of illness avoidance effort.

The Private Economic Problem

We assume that a forward-looking representative individual, who believes himself/herself to be susceptible, seeks to mitigate disease risks by avoiding those who are obviously (symptomatically) infectious. We suppose that individuals may belong to any one of the following epidemiological states at a given time: (S)usceptible to the disease, (E)xposed or latently infected being asymptomatic and noninfectious, (A)symptomatically infectious, (I)nfectious with symptoms, or (R)ecovered and immune to the disease. Individuals who believe themselves to be susceptible may include those in states S , E , and A . We further assume the motivation for selective mixing is the desire to avoid the costs of illness only. We do not allow individuals to behave altruistically. We also assume that individuals who know themselves to be infected have no incentive to avoid others. Only susceptible people (or those who believe themselves to be susceptible) react to disease risk. This includes all individuals in health classes $X = S, E, A$, or R_A (recovered

from asymptomatic infection). Because all individuals who react to disease risk consider themselves to be equally susceptible to the disease, their mixing decisions are both identical and symmetric [i.e., all $\phi_{XI}(t) = \phi_{IX}(t)$ are equal to one another and all other entries in $\Phi(t)$ are 0]. Formally, these individuals screen contacts by choosing the elements $\phi_{ij}(t)$ of the matrix $\Phi(t)$, $i, j \in \{S, E, A, I, R_A\}$ so as to maximize the difference between the expected benefits of contact and the expected cost of illness and illness avoidance, given their current health state.

The benefits of contact range from the satisfaction to be had from purely social engagement to the financial gains to be had from market transactions with others. For simplicity, we assume that the benefits of a contact are financial gains and that they are the same for individuals in all health states, $B_i = B$, $i \in \{S, E, A, I, R\}$. The cost of illness generally includes both forgone earnings, lost wages, and the cost of healthcare. For simplicity, we take the cost of illness to be the cost of treatment. The cost of illness avoidance is simply the cost of the effort made to avoid contact with people who are ill: the cost of choosing $\phi_{iI}(t)$, $i \in \{S, E, A, R\}$. The net benefits of contacts with others by an individual in the i th health state at time t thus comprise the difference between the benefit of contacts made in that health state, B_i , and the cost of disease and disease risk mitigation, $C_i(\phi_{ij}(t), I(t))$ (see Table 1). All individuals within a particular disease class are assumed to behave in the same way.

We define $U_i(\phi_{ij}(t)) = B_i - C_i(\phi_{ij}(t), I(t))$ to be the net benefits of contact to the representative individual in health state i at time t . We define $E\left(V_j\left(t + \tau, \phi_{ij}(t + \tau)\right)\right)$ to be the individual's expected value function from time $t + \tau$ forward, where the probability that the individual will be in health state j in the future depends on the disease

Table 2. Models Studied and Corresponding Compartments.

Model	Respiratory	Nonrespiratory	Compartments
SIR	Influenza		$S \rightarrow I \rightarrow R$
SEIR	SARS	Polio, measles, smallpox, meningitis, West Nile virus	$S \rightarrow E \rightarrow I \rightarrow R$
One-path SAIR		Hepatitis B, rubella	$S \rightarrow A \rightarrow I \rightarrow R$
Two-path SAIR	Influenza	S cholera	$S \rightarrow \begin{cases} A \rightarrow R_A \\ I \rightarrow R_I \end{cases}$

Bolditalic compartments denote individuals who engage in risk-mitigation. Compartment classifications are shown in Table 1. For noncommunicable diseases (such as cholera or West Nile virus), it is assumed that the number of infectious individuals is proportionate to the infectious material, or the disease vector.

dynamics and their mixing strategy while in health state i . The decision problem for the representative susceptible individual in health state i may be expressed via the Hamilton–Jacobi–Bellman equation:

$$V_i(t, H(t)) = \max_{\phi_{ij}(t)} \left\{ \int_t^{t+\tau} U_i(\phi_{ij}(t)) dt + E(V_j(t + \tau, \phi_{ij}(t + \tau))) \right\} \tag{2}$$

where τ is a short interval of time. This is subject to the dynamics of the disease: Eq. (1), Table 3 The H–J–B equation identifies the problem solved by the representative individual in state i : to maximize the expected net value of current and future contacts by choosing the extent to which they mix with individuals in other health classes. Following Fenichel et al. (2011), we assume that individuals form their expectations adaptively. The value function $V_i(t, H(t))$ is defined recursively as the sum of the current net benefits of contact in health state i given the information available at time t plus the discounted stream of expected net benefits over the remaining horizon. This expectation is conditional on the effects of disease risk mitigation decisions on the probability of transitioning between health states. More particularly, we assume individuals observe the state of an epidemic at time t and make a forecast for the epidemiological trajectory over the interval τ . Their mixing strategies are then adapted over time as they make new observations on the state of the epidemic. We assume that individuals make the simplest forecast—that all disease classes are constant over the interval τ , but that they adapt to new data as it emerges.

The representative individual will increase effort to avoid infection up to the point where the marginal cost of illness avoidance is just offset by the marginal benefits it yields—the avoided cost of illness. Efforts to avoid infection will be increasing in the cost of illness and decreasing

in the cost of illness avoidance. In models without risk mitigation, disease dynamics may be completely characterized from initial conditions. With risk mitigation, the evolution of the epidemic reflects feedback between the cost of disease and disease avoidance on the one hand, and averting behavior on the other [see (Fenichel and Horan 2007; Horan et al. 2011) for further discussion].

To solve the problem, we take a discrete time counterpart to Eq. (2) and solve numerically using techniques similar to those in Fenichel et al. (2011). Specifically, we solve the adaptive expectation problem using a method we call Cast-Recast. At each time, the individual solves the H–J–B equation using backwards induction from their time horizon (12 days) to the present, while supposing that there is no change to the state variables over the time horizon. Having determined the optimal mitigation effort, $\phi_{ij}(t)$, the individual commits that effort until the next time step (day). That is, their mitigation effort is held constant for 1 time step. The ordinary differential equations describing the disease dynamics are advanced, and the process is repeated. Note that the private forecast used has little effect on the optimal outcomes so long as the forecast period, the time interval τ , is short relative to the disease dynamics. If the epidemic evolves rapidly relative to the period over which the individual commits to a fixed level of risk mitigation, then the assumption that the disease states are constant may induce errors. In previous work using this method (Fenichel et al. 2011; Morin et al. 2013, 2014, 2015), we found a smooth response—the decisions made by individuals were much the same from day to day. Since the epidemic evolves on a timescale of weeks, this gives us confidence in the Cast-Recast method for the private problem. In the discussion, we outline when the results from such a method may favor more rapid transmission of epidemiological state variables, and when individual

Table 3. Differential Equations Used for Each Epidemic Model, Other than the Previously Stated SIR Model.

SEIR	One-path SAIR	Two-path SAIR
$\frac{dS(t)}{dt} = -c\beta S(t)P_{SI}$	$\frac{dS(t)}{dt} = -cS(t)(\beta_A P_{SA} + \beta_I P_{SI})$	$\frac{dS(t)}{dt} = -cS(t)(\beta_A P_{SA} + \beta_I P_{SI})$
$\frac{dE(t)}{dt} = -c\beta S(t)P_{SI} - \kappa E(t)$	$\frac{dA(t)}{dt} = -cS(t)(\beta_A P_{SA} + \beta_I P_{SI}) - \gamma_A A(t)$	$\frac{dA(t)}{dt} = -cS(t)(\beta_A P_{SA} P_{AA} + \beta_I P_{SI}(1 - p_{II})) - \gamma_A A(t)$
$\frac{dI(t)}{dt} = \kappa E(t) - \gamma I(t)$	$\frac{dI(t)}{dt} = \gamma_A A(t) - \gamma_I I(t)$	$\frac{dI(t)}{dt} = -cS(t)(\beta_A P_{SA}(1 - p_{AA}) + \beta_I P_{SI} p_{II}) - \gamma_A A(t)$
$\frac{dR(t)}{dt} = \gamma I(t)$	$\frac{dR(t)}{dt} = \gamma_I I(t)$	$\frac{dR_A(t)}{dt} = \gamma_A I(t)$
		$\frac{dR_I(t)}{dt} = \gamma_I I(t)$

behavior choices need to be more “agile” in order to match timescales with the disease spread.

The Social Problem

The choice of $\phi_{ij}(t)$ maximizes the private net benefits of contact for the individual over the course of an epidemic, given the private cost of illness and illness avoidance. However, as was shown in Morin et al. (2015) this may well be publically sub-optimal, depending on the social rate of discount or the social cost of illness. The social cost of illness is the sum of the costs borne by all infected and symptomatic individuals, together with the cost of disease avoidance by all others. In a real system, it would also include the infection risk borne by healthcare workers, but we do not address that here. Of the many intervention options open to public health authorities—quarantine, social distancing measures such as school closures, vaccination campaigns and so on—we focus on instruments that change risk mitigation by changing the private cost of illness. In doing this, we follow a literature that integrates epidemiology and economics to explore ways in which economic behavior affects disease spread (Perrings et al. 2014). This literature has concentrated on the economic causes and epidemiological consequences of peoples’ contact decisions (Gersovitz and Hammer 2003, 2004; Barrett and Hoel 2007; Funk et al. 2009; Fenichel et al. 2010; Funk et al. 2010; Springborn et al. 2010). By treating the economic factors behind contact and mixing decisions as central elements in disease transmission, the approach opens up a new set of disease management options.

In what follows, we suppose that the public health authority is able to use an economic policy instrument to alter the private cost of illness. We consider the instrument, r_D , $D \in \{\text{SIR, SEIR, One Path SAIR, Two Path SAIR}\}$ see Tables 2 and 3, which may be interpreted as a disease-specific tax or a subsidy on the cost of illness, $C_I(I, r_D) := C_I(1 + r_D)$. If $r_D > 0$ (illness is ‘taxed’), the private cost of illness is increased. Enforced, uncompensated sick ‘leave’ would be an example of this. We expect this to increase disease-risk

mitigation effort and hence the illness avoidance costs carried by individuals. If $r_D < 0$ (illness is ‘subsidized’), the private cost of illness is reduced. Subsidized health insurance schemes would be an example of this. We expect this to reduce disease-risk mitigation effort, and with it the illness avoidance costs carried by reactive individuals. By changing the privately optimal level of disease-risk mitigation, it is possible to change overall disease dynamics.

There are many ways in which interventions change the private cost of illness in real-world conditions, ranging from direct subsidies or taxes on drugs and treatment, through health insurance costs and coverage, to statutory obligations on sick leave. We suppose that r_D can be applied in a way that proportionately reduces or increases the relative private cost of illness. This would be consistent with, for example, mandatory insurance cover for a specific proportion of potentially forgone earnings. Our baseline case assumes that the policy instrument is revenue neutral. If $r_D < 0$ (illness is subsidized), the cost is met by a levy on all income from contacts. If $r_D > 0$ (illness is taxed), the revenue is returned as a tax benefit on all income from contacts. Since taxes and subsidies both potentially impose an efficiency cost in the form of a deadweight loss of consumer and producer surplus, we include a proxy for this in the optimization problem. More particularly, we include a term that specifies any deadweight loss as a proportion of the cost of taxes or subsidies.

The public health authority’s problem for an SIR disease thus takes the form:

$$\begin{aligned}
 W(t, r_D, H(t)) = & \max_{r_D} \int_0^T \sum_i \sum_j e^{-\delta t} E \left[U_i \left(\phi_{ij}(t, r_D) \right) \right] dt \\
 & - (1 - \alpha) \int_0^T e^{-\delta t} C_I(r_D, I(t)) dt
 \end{aligned}
 \tag{3}$$

subject to the disease dynamics described by the relevant compartmental epidemiological model and to the private decision problem described in Eq. (2). That is, the public

health authority selects r_D so as to maximize the net benefits of risky contacts to society—where society is the sum of all individuals in all health classes. The final term in the public health authority’s problem is our proxy for the deadweight loss associated with taxes or subsidies on the cost of illness. $\alpha \in [0, 1]$ is the proportion of the cost of the intervention that is recovered. To solve these two problems, we maximize the integral in Eq. (3) over the entire epidemic by solving the complete private problem for each “guess” of r_D . This was implemented using MATLAB’s `fminbnd` function.

RESULTS

General Epidemiological Effects of Public Health Authority Intervention on Private Disease-Risk Mitigation

We investigated the effect of a subsidy/tax on the cost of illness on private disease risk mitigation in the key epidemiological classes across four compartmental models: SIR, SEIR, and two SAIR models, a one-path progression and a two-path progression. Within the one-path model, susceptible individuals are first asymptotically infectious and then progress to symptomatically infectious and then immune. In the two-path model, a susceptible individual becomes either asymptotically *or* symptomatically infectious and then recovers. Individuals who recover from asymptomatic infection (R_A) are expected to behave as if they are susceptible. Individuals who recover from symptomatic infection (R_I) are not. Because we have assumed no heterogeneity aside from health status, the use of a disease-specific instrument, r_D , has the potential to be efficient. If the population were heterogeneous, however, we would expect to need more targeted interventions.

We did not consider models with reentry to the susceptible class, e.g., SIS, SIRS, and other cyclical models. This is for two reasons. First, each of these models is capable of endemic levels of infection. This, combined with the fact that people may experience reinfection, would require individuals to form expectations (possess memory) with respect to the impact of different avoidance strategies. Second, numerical simulations of these models reveal very broad oscillations that confound comparison with single outbreak models. There are no entries (births) or removals (deaths) from the system (see Table 2 for example diseases listed for each model). Each model considered here is

therefore a so-called single outbreak model; the population is not only kept at a fixed number, but it is also closed to the introduction of new individuals.

We modeled the dynamics of the epidemic types in Table 2 using ordinary differential equations (see Table 3). This has two main implications for disease dynamics: (a) once nonzero, the state variables will never again be zero in finite time, and (b) in an infinitely small amount of time “mass” will move into each compartment as long as the transition rates are nonzero. These affect the interpretation to be given of the point at which an epidemic is “over.” It is feasible that extreme risk mitigation early in the course of the epidemic could wipe out the infection within a population. However, within the differential equation framework, as soon as mitigation weakens, coupled with the fact that there is a nonzero infectious population with potentially a very large susceptible population, the infection will again spread, potentially causing additional peaks.

The instantaneous transfer of individuals from one compartment to another also serves to induce reaction timing that may not conform to data. In recognition of this, we supposed that the differential equations represent an expectation of outcomes over a population divided between three health classes. In previous work (Morin et al. 2014, 2015), we demonstrated that private disease-risk mitigation reduces peak symptomatic infection levels and the total number of people who experience symptomatic infection while prolonging the epidemic. We seek to understand how public health interventions aimed at minimizing the cost of disease and disease avoidance affect individual risk mitigation decisions.

Our measure of mitigation effort is a relative one. It is the proportion of maximum mitigation effort undertaken, where maximum effort is defined as that which results in $P_{SI}(t) = 0$. This enables us to compare effort across diseases. The background against which public health authority intervention occurs is that private mitigation effort is increasing in infectiousness (β), infectious duration ($1/\gamma$), and the share of the public cost of intervention recovered by the public health authority (α). In other words, private mitigation effort is increasing in the cost of illness. We find that in all cases peak effort occurs early in the epidemic—within the first 30 days for our parameters—and that action is taken sooner, the greater the severity of the disease. Our results on the socially optimal response this induces for different diseases follow. These results are driven by the assumptions we make about the cost of illness (equal to the benefit gained from being

healthy) relative to the cost of mitigation. Specifically, we assume that the cost of mitigation is low relative to the cost of illness.

SIR

To explore the sensitivity of the optimal intervention strategy of the public health authority to the severity of disease, we varied β and γ over intervals so that $\beta/\gamma \in [0.9, 3.29]$, and α in steps of 0.1 from 0 to 1, with baseline values of $\alpha = 1$, $\beta = 0.02$, and $\gamma = 1/7$. While the socially optimal public health authority intervention decision can be viewed as a function of the disease's basic reproduction number $R_0 = \beta/\gamma$, the two components of R_0 (infectiousness and duration) have rather different effects on the public health authority intervention. Nor is the optimal social response monotonic in either case. $R_0(\beta)$ tests the sensitivity of health interventions to the impact of infectiousness on R_0 and tests the sensitivity of health interventions to the impact of disease duration on R_0 . For $R_0(\beta) \in [0.95, 1.91]$, the optimal public health authority intervention involves an increase in the private cost of illness—a 'tax' on illness that will stimulate higher levels of private disease risk mitigation. As infectiousness falls, the optimal public health authority incentive to mitigate first rises and then falls, the turning point being determined by parameters describing both the cost of illness and the cost of illness avoidance. The optimal tax in this case reaches a maximum of 281% for $R_0(\beta) = 1.0345$. For diseases where infectiousness is either very low ($R_0(\beta) < 0.941$) or very high (i.e., $R_0(\beta) > 1.91$), the optimal public health authority intervention involves a reduction in the cost of illness—a 'subsidy' on illness that lowers private disease-risk mitigation effort (Fig. 1).

The duration of illness has a slightly different effect. Note that small $R_0(\gamma)$ values are indicative of diseases of short duration. For $R_0(\gamma) = 0.95$, for instance, the illness is symptomatic for less than 4 days. We found that for diseases of moderate to short duration, $R_0(\gamma) < 2.125$, it would generally be optimal for the public health authority to stimulate an increase in private risk mitigation by raising the private cost of illness—by taxing illness. On the other hand, for diseases of longer duration, $R_0(\gamma) > 2.125$, we found that it would always be optimal for the public health authority to reduce private risk mitigation by lowering the private cost of illness.

Given the assumed cost of illness and mitigation effort, private individuals will overreact to the risk of diseases of

long duration and high infectiousness and will underreact to diseases of short duration and low infectiousness. To see the trade-off between infectiousness and duration, we considered the range of (β, γ) that generates $R_0 \in [1.91, 2.125]$. Figures 2 and 3 show the parameter combinations that leave the public health authority indifferent between intervening or not over this range of values, or that favor intervention to increase or decrease privately optimal disease risk mitigation. While infectiousness and duration trade off against each other, the relation is not linear. Where private individuals take excessive precautions, PHA interventions will generally discourage private risk mitigation for highly infectious diseases of long duration. This changes, however, if the infectiousness of diseases of long duration is very low or the duration of highly infectious diseases is very short. Symmetrically, PHA interventions will generally encourage private risk mitigation for less infectious diseases of short duration. Once again, though, this changes if the infectiousness of diseases of short duration is very high or the duration of less infectious diseases is very long.

The socially optimal intervention will also depend—intuitively—on the degree to which the susceptible and infected individuals carry the cost of public health authority interventions. To capture this, we included a cost recovery parameter, α , and explored the sensitivity of the optimal intervention strategy of the PHA to variation in α (Fig. 3). The proportion of the cost of a subsidy, on the cost of illness, recovered through taxation (or the proportion of tax revenues returned to the wider population) is given by α . Specifically, we varied α in steps of 0.1 from 0 to 1 with a baseline value of $\alpha = 1$. For the baseline values for infectiousness and duration, $\beta = 0.02$ and $\gamma = 1/7$, we found that the greater the value of α , the less one needs to subsidize the cost of illness in order to align private and social optima. The reason is that the cost of illness is affected both by the instrument itself and by cost recovery. As the rate of cost recovery increases, the cost of illness rises along with the privately optimal level of risk mitigation. Indeed, for $\alpha \geq 0.97$ the public health authority switches from discouraging to encouraging private risk mitigation—from a subsidy to a tax.

SEIR

The inclusion of a class of individuals who have unknowingly contracted the illness but are not a danger to others (they are asymptomatic and noninfectious) changes the

economics of the problem in important ways. On the one hand, it unnecessarily increases the aggregate cost of mitigation since exposed individuals continue to mitigate, even though they there is no need. On the other hand, by driving up P_{SE} and reducing P_{SB} , exposed individuals reduce infection rates below those that would occur if only susceptible individuals mitigated risk. That is, the risk mitigation undertaken by exposed individuals confers an external benefit on society. The result is that public health authority intervention in the SEIR case increases private risk mitigation for all but the most severe diseases (i.e., for all $R_0 \leq 3.25$) (Fig. 4). We also note that the basic reproduction number for the SEIR model is identical to that for the SIR model and does not involve the latent period. Over reasonable latent periods from 3 to 14 days, we found an increase in the ‘tax’ rate applied by the public health authority, but no switching or nonlinearities in behavioral response.

The big differences between the SIR and SEIR cases are that the introduction of a latency period slows the spread of disease, while the mitigation by exposed individuals reduces the rate of new infections. Interestingly, when we decomposed the marginal benefit of public health authority intervention by the infectiousness and duration of disease, we found the marginal benefit of intervention to be increasing in the duration of disease, but decreasing in infectiousness (Fig. 4).

As β increases and causes $R_0(\beta)$ to pass 1.45, we found the marginal benefits of public health authority intervention to be monotonically decreasing. This is because $\frac{\partial^2 \phi_{SI}}{\partial r_D \partial \beta} < 0$. However, as duration of disease increases, causing $R_0(\gamma)$ to increase, we found the marginal benefits of PHA intervention to rise monotonically after approximately 1.46. Private disease-risk mitigation in the SIR case reduces peak prevalence but prolongs the epidemic at higher levels than would have occurred in the absence of mitigation. In the SEIR case, by contrast, private disease risk mitigation reduces peak prevalence and prolongs the epidemic, but at lower levels than would have occurred in the absence of mitigation. Moreover, the longer the duration of the disease the greater the social net benefits this offers.

One-Path SAIR

The one-path SAIR model introduces an asymptotically infectious stage between susceptible and symptomatically infectious. These individuals now undertake risk mitigation

because they do not know they’re infected. Unlike the SEIR case, however, this is potentially harmful to the population. While $\frac{\partial P_{SI}}{\partial \phi_{SI}} < 0$, it is also true that $\frac{\partial P_{SA}}{\partial \phi_{SI}} > 0$. We may assume that the two classes have different levels of infectiousness. The net effect of private disease risk mitigation is therefore given by the combined impact of two things: the marginal reduction in the infection rate due to the avoidance of infected people and the marginal increase in infections due to susceptible-asymptomatic contact. The marginal effect of risk mitigation through mixing is $\frac{\partial P_{SI}}{\partial \phi_{SI}} \beta_I$ and the marginal effect of susceptible-asymptomatic contact is $\frac{\partial P_{SA}}{\partial \phi_{SI}} \beta_A$. If $\frac{\partial P_{SI}}{\partial \phi_{SI}} \beta_I + \frac{\partial P_{SA}}{\partial \phi_{SI}} \beta_A < 0$ private risk mitigation reduces prevalence regardless of the actions of asymptotically infectious individuals. If $\frac{\partial P_{SI}}{\partial \phi_{SI}} \beta_I + \frac{\partial P_{SA}}{\partial \phi_{SI}} \beta_A > 0$, it has the opposite effect. Whether the public health authority seeks to increase or decrease private disease risk mitigation effort depends on the net effect. Given our baseline parameters, if $\beta_A < -1.2\beta_I + 0.03$ then the public health authority will seek to increase private mitigation, and if $\beta_A > -1.2\beta_I + 0.03$, it will seek to reduce private mitigation. The strength of intervention increases with the distance of (β_I, β_A) from the line $\beta_A = -1.2\beta_I + 0.03$ (which is roughly analogous with the condition that the basic reproduction number is 1).

Two-Path SAIR

In the two-path SAIR model, asymptotically infectious individuals can spread disease due to their (redundant) risk mitigation actions. As in the SEIR model, those who’ve recovered from asymptomatic infection undertake mitigation that does not benefit them, but may offer a benefit to susceptible individuals by reducing their chance of making infectious contacts. In some sense, the R_A class becomes like a vaccinated class of potentially highly connected individuals. They therefore dilute the contact pool for susceptible individuals. We find the interesting effect that whether to tax or subsidize is sensitive only to the recruitment into the I class. If the process favors the generation of symptomatic infection, then the PHA intervenes with a tax on illness. If the process favors the generation of asymptomatic infection, then the PHA intervenes with a subsidy. A process favorably generates asymptomatic infection when $\beta_A P_{SA} p_{AA} + \beta_I P_{SI} (1 - p_{II}) > \beta_A P_{SA} (1 - p_{AA}) + \beta_I P_{SI} p_{II}$ where p_{XX} is the probability that an infected person of type- X creates another individual of type- X when they spread infection. Therefore, adjustments will be made to the cost of

illness until $\beta_A \frac{\partial P_{SA}}{\partial \phi_{SI}} (1 - 2p_{AA}) < \beta_I \frac{\partial P_{SI}}{\partial \phi_{SI}} (1 - 2p_{II})$, where $\frac{\partial \phi_{SI}}{\partial r_D}$ cancels out from both sides. If $p_{AA}, p_{II} < \frac{1}{2}$, then the marginal condition is never met and no mitigation will occur. Likewise, if both $p_{AA}, p_{II} > \frac{1}{2}$ then mitigation will occur indefinitely until $P_{SI} = 0$. If one is less than 1/2 with the other greater, then the level of mitigation is unclear because it now depends on the relative strengths of infection and the sensitivities of pair probabilities to mitigation.

DISCUSSION

When private disease-risk mitigation has general epidemiological consequences, the optimal level of mitigation should be determined by reference to the costs and benefits to society at large. Because disease-risk mitigation changes the characteristics of an epidemic, it changes the social costs and benefits of disease. If private mitigation reduces prevalence, for example, it confers benefits on society. If it increases the duration of an epidemic, on the other hand, it imposes costs. Since the private and social calculus of the costs and benefits of mitigation are different—individuals calculate only the costs and benefits to themselves, society calculates the costs and benefits to all—we expect the privately and socially optimal level of disease risk mitigation to differ.

We have argued that individuals base their mitigation decisions on the expected net benefits of particular actions. Very similar arguments have long been made in the medical literature. The health belief model, for example, was developed to explain why people undertook (or abstained from) private disease-risk mitigation (Champion and Skinner 2008). The model, originally developed in the 1950s, assumed that individuals engaged in health-risk mitigation to reduce the threat of illness based on their beliefs about four things: the likelihood of contracting disease; the consequences of falling ill (symptoms, loss of work wages, loss of personal interactions); the efficacy and costs of the proposed action; and their capacity to undertake the action adequately. It also assumed that individuals responded to incentives—‘cues to action’ that signaled susceptibility, severity, costs and the like. The perceived benefits and costs of alternative behaviors determined which behaviors would be undertaken (Rosenstock 1974).

In the context of this paper, the perceived susceptibility of an individual at time t , conditional on a mitigation effort $\phi(t)$, is given by $\int_T^t \exp(-\beta P_{SI}(\phi(\tau))) d\tau$, and the expected

cost of illness is the daily forgone income (adjusted by public health authority intervention) multiplied by the expected duration of illness, $C_{II}(1 + r_D)/\gamma$. While we only allow a single risk mitigation choice, affinity-based mitigation, we allow it to vary from “doing nothing” ($\phi(t) = 0$) to “private quarantine” ($\phi(t) = \frac{R(t) - \sqrt{(R(t))^2 + 4X(t)I(t)}}{2X(t)I(t)}$ where $X(t)$ is the total population of susceptible individuals). As in the framework of the health belief model, we expect that mitigation effort applied will be increased up to the point at which its cost is just offset by the benefits, in terms of the reduced probability of illness.

Applications of the health belief model to disease-risk mitigation in the 1974 swine flu outbreak found that the framework adequately segmented the respondent populations into those who vaccinated and those who did not (Janz and Becker 1984). They found that perceived susceptibility, benefits, and barriers were strongly correlated with vaccine-seeking behavior and that these results were strengthened by the fact that, even if it was unsuccessful, the vaccine alleviated the symptoms of illness (Aho 1979; Cummings et al. 1979; Larson et al. 1979; Rundall and Wheeler 1979; Janz and Becker 1984).

The central proposition of this paper is that since disease-risk mitigation is a function of the private cost of disease, it can be managed through changes in the private cost of disease. Whether private disease-risk mitigation is above or below the socially optimal level of mitigation, there exist interventions that have the potential to align private choices with the interests of society. Our numerical results on the difference between the privately and socially optimal mitigation reflect the specific assumptions made about the relative cost of disease and disease avoidance. These were selected to illustrate the range of potential responses and the sensitivity of responses to disease characteristics. Given our cost assumptions, we find that for SIR diseases individuals confronted with highly infectious illnesses of long duration overreact to the associated risks. They mitigate more than is socially optimal. Conversely, individuals confronted with short illnesses of low infectiousness underreact. For SEIR diseases, on the other hand, the privately optimal level of disease-risk mitigation is less than the socially optimal level for almost all parameter values. The gap between privately and socially optimal investments in risk mitigation is partly due to the fact that private individuals base their risk mitigation decisions on the observed health state of others—whether or not others are symptom-free. This is a source of error that is greater in

some cases (SEIR, two-path SAIR) than others (SIR, one-path SAIR). It is also due to the epidemiological effects of actions that influence when a susceptible person becomes infected in the course of an epidemic. One of the main external effects of private disease-risk mitigation is the prolongation of epidemics and with it the cost of disease-risk mitigation.

We therefore considered interventions that ran in both directions: subsidies on the cost of illness (e.g., health insurance) that discourage mitigation and taxes (e.g., uncompensated mandatory days off or other such penalties) that encourage mitigation. It can, by the way, be shown that subsidies or taxes on the cost of mitigation rather than the cost of illness have similar effects. For any adjustment to the cost of illness, there exists an adjustment to the cost of mitigation that produces the same private response. It does not follow, however, that the two instruments are equally efficient. Adjusting the cost of illness alters ‘severity’ in the language of the health belief model while adjusting the cost of mitigation alters ‘barriers.’ Which is the more cost-effective in any real application would depend on the effort needed to achieve the desired epidemiological effect and the relative cost of effort in each case.

Within the SIR model, each individual has perfect information on their infection state and all relevant information on others (infectious or not). The other three models each introduce error due to the existence of asymptomatic classes. They also introduce changes in disease dynamics that affect the socially optimal level of risk mitigation. Whether it is socially optimal to intervene in ways that increase or decrease private disease risk mitigation depends on the effect of mitigation on the time profile of infections. The point has already been made that mitigation reduces peak prevalence but also prolongs epidemics. When the additional cost of the longer (and potentially fatter) tail of an epidemic outweighs the benefits offered by a reduction in peak prevalence, the PHA will intervene to reduce private risk mitigation—and vice versa (as shown in Fig. 5). The advantages of instruments that operate on the cost of illness are that they can be put in place in advance of an outbreak and that they operate automatically. Disease-specific insurance or treatment costs, for example, can be established in advance but only influence private decisions when an outbreak occurs. We have not considered particular instruments in this paper and note only that the effect of each depends on how it would impact the

relative costs and benefits of mitigation actions. There is a range of cost-based instruments that could be applied to diseases whose properties are well understood. The instruments that should be applied to specific diseases are those expected to cost effectively close the gap between socially and privately optimal disease risk mitigation. While we consider only policy instruments determined in advance of an epidemic—i.e., that correspond to the known properties of known diseases—we acknowledge that this would not be reasonable in the case of novel or emerging infectious diseases. There is less scope for the use of incentives of this type in such cases, but it would be worth considering how the choice of r_D might evolve during the course of an epidemic in future work.

Finally, it is worth noting that any intervention of this kind in a real-epidemiological-economic environmental would have distributional consequences. It would make some people better off and some people worse off. We have abstracted from any distributional consequences in this paper by treating all individuals in each disease class as homogeneous, and by assuming that the payoff to contact is the same for everybody. Implementation of a policy instrument of this kind in a real system would, however, need to take account of the distributional goals of the society concerned. Infectious diseases are a comparatively small part of the burden of disease in high-income countries, but still the largest part of the disease burden in low-income countries. This is precisely where distributional issues are of greatest concern and where individuals are least able to bear the cost of disease-risk avoidance.

ACKNOWLEDGEMENTS

This study was made possible by Grant # 1R01GM100471-01 from the National Institute of General Medical Sciences (NIGMS) at the National Institutes of Health and contract HSHQDC-12-C-00058 from the Science and Technology Directorate, Department of Homeland Security. The contents of this manuscript are solely the responsibility of the authors and do not necessarily represent the official views of DHS or NIGMS. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript. It was also funded by NSF Grant 1414374 as part of the joint NSF-NIH-USDA Ecology and Evolution of Infectious Diseases program.

REFERENCES

- Aho WR (1979) Participation of senior citizens in the swine flu inoculation program: An analysis of health belief model variables in preventive health behavior. *Journal of Gerontology* 34:201–208
- Barrett S, Hoel M (2007) Optimal disease eradication. *Environment and Development Economics* 12:627–652
- Bayham J, Kuminoff NV, Gunn Q, Fenichel EP (2015) Measured voluntary avoidance behaviour during the 2009 A/H1N1 epidemic. *Proceedings of the Royal Society B: Biological Sciences* 282(2015):0814
- Blythe S, Castillo-Chavez C, Palmer J, Cheng M (1991) Toward a unified theory of sexual mixing and pair formation. *Mathematical biosciences* 107:379–405
- Brownstein JS, Wolfe CJ, Mandl KD (2006) Empirical evidence for the effect of airline travel on inter-regional influenza spread in the United States. *PLoS Med* 3:e401
- Busenberg, S. & Castillo-Chavez, C. (1989) Interaction, pair formation and force of infection terms in sexually transmitted diseases. *Mathematical and statistical approaches to AIDS epidemiology*, pp. 289–300. Springer.
- Castillo-Chavez C, Busenberg S, Gerow K (1991) Pair formation in structured populations. *Differential equations with applications in biology, physics and engineering* 47:65
- Champion, V.L. & Skinner, C.S. (2008) The health belief model. *Health behavior and health education: Theory, research, and practice* (ed. by K. Glanz, B.K. Rimier and K. Viswanath), pp. 45–65. Jossey-Bass.
- Cummings KM, Jette AM, Brock BM, Haefner DP (1979) Psychosocial determinants of immunization behavior in a swine influenza campaign. *Medical care* 17:639–649
- Curtis V, Barra MD, Aunger R (2011) Disgust as an adaptive system for disease avoidance behaviour. *Philosophical Transactions of the Royal Society B: Biological Sciences* 366:12
- Fenichel E, Horan R (2007) Jointly-determined ecological thresholds and economics trade-offs in wildlife disease management. *Natural Resources Modeling* 20:511–547
- Fenichel E, Castillo-Chavez C, Ceddia M, Chowell G, Gonzales Parra P, Hickling G, Holloway G, Horan R, Morin B, Perrings C, Springborn M, Velazquez L, Villalobos C (2011) Adaptive human behavior in epidemiological models. *Proceedings of the National Academy of Sciences* 108:6306–6311
- Fenichel EP, Horan RD, Hickling GJ (2010) Management of infectious wildlife diseases: bridging conventional and bioeconomic approaches. *Ecological Applications* 20:903–914
- Fenichel EP, Kuminoff NV, Chowell G (2013) Skip the trip: Air travelers' behavioral responses to pandemic influenza. *PLOS One* 8:e58249
- Ferguson NM, Cummings DA, Fraser C, Cajka JC, Cooley PC, Burke DS (2006) Strategies for mitigating an influenza pandemic. *Nature* 442:448–452
- Funk S, Salathe M, Jansen VAA (2010) Modelling the influence of human behaviour on the spread of infectious diseases: a review. *Journal of the Royal Society, Interface* 7:1247–1256
- Funk S, Gilad E, Watkins C, Jansen V (2009) The spread of awareness and its impact on epidemic outbreaks. *Proc National Acad Sci* 106:6872–6877
- Geoffard, P.-Y. & Philipson, T. (1997) Disease eradication: private versus public vaccination. *The American Economic Review*, 222–230.
- Germann TC, Kadau K, Longini IM, Macken CA (2006) Mitigation strategies for pandemic influenza in the United States. *Proceedings of the National Academy of Sciences* 103:5935–5940
- Gersovitz M, Hammer JS (2004) The economical control of infectious diseases. *The Economic Journal* 114:1–27
- Gersovitz M, Hammer JS (2003) Infectious diseases, public policy, and the marriage of economics and epidemiology. *The World Bank Research Observer* 18:129–157
- Gurarie E, Ovaskainen O (2013) Towards a general formalization of encounter rates in ecology. *Theoretical Ecology* 6:189–202
- Horan RD, Fenichel EP, Melstrom RT (2011) Wildlife Disease Bioeconomics. *International Review of Environmental and Resource Economics* 5:23–61
- Hunter, N.D. (2007) 'Public-Private' Health Law: Multiple Directions in Public Health. *Journal of Health Care Law & Policy*, 10.
- Hyman JM, Li J (1997) Behavior changes in SIS STD models with selective mixing. *SIAM Journal on Applied Mathematics* 57:1082–1094
- Janz NK, Becker MH (1984) The health belief model: A decade later. *Health Education & Behavior* 11:1–47
- Karlin S (1979) Models of multifactorial inheritance: I, Multivariate formulations and basic convergence results. *Theoretical Population Biology* 15:308–355
- Larson E, Olsen E, Cole W, Shortell S (1979) The relationship of health beliefs and a postcard reminder to influenza vaccination. *The Journal of family practice* 8:1207–1211
- Leroy EM, Kumulungui B, Pourrut X, Rouquet P, Hassanin A, Yaba P, Délicat A, Paweska JT, Gonzalez J-P, Swanepoel R (2005) Fruit bats as reservoirs of Ebola virus. *Nature* 438:575–576
- Morin BR, Fenichel EP, Castillo-Chavez C (2013) SIR dynamics with economically driven contact rates. *Natural resource modeling* 26:505–525
- Morin BR, Perrings C, Levin S, Kinzig A (2014) Disease risk mitigation: The equivalence of two selective mixing strategies on aggregate contact patterns and resulting epidemic spread. *Journal of Theoretical Biology* 363:262–270
- Morin, B.R., Perrings, C., Kinzig, A. & Levin, S. (2015) The social benefits of private infectious disease-risk mitigation. *Theoretical Ecology*, 1–13.
- Perrings C, Castillo-Chavez C, Chowell G, Daszak P, Fenichel E, Finnoff D, Horan R, Kilpatrick AM, Kinzig A, Kuminoff N, Levin S, Morin B, Smith K, Springborn M (2014) Merging Economics and Epidemiology to Improve the Prediction and Management of Infectious Disease. *EcoHealth* 11:464–475
- Philipson, T. (1996) Private vaccination and public health: an empirical examination for US measles. *Journal of Human Resources*, 611–630.
- Philipson T (2000) Economic epidemiology and infectious diseases. *Handbook of health economics* 1:1761–1799
- Presanis, A.M., De Angelis, D., Flu, T.N.Y.C.S., Hagy, A., Reed, C., Riley, S., Cooper, B.S., Finelli, L., Biedrzycki, P. & Lipsitch, M. (2009) The severity of pandemic H1N1 influenza in the United States, from April to July 2009: a Bayesian analysis. *PLoS medicine*, 6.
- Rao S, de Gier NS, Caram LB, Frederick J, Moorefield M, Woods CW (2009) Adherence to self-quarantine recommendations during an outbreak of norovirus infection. *Infection Control* 30:896–899

- Rosenstock IM (1974) The health belief model and preventive health behavior. *Health education monographs* 2:354–386
- Rubin, G.J., Amlôt, R., Page, L. & Wessely, S. (2009) Public perceptions, anxiety, and behaviour change in relation to the swine flu outbreak: cross sectional telephone survey. *Bmj*, 339.
- Rundall TG, Wheeler JR (1979) Factors associated with utilization of the swine flu vaccination program among senior citizens in Tompkins County. *Medical Care* 17:191–200
- Sanchez, A. (2001) *Ebola viruses*. Wiley Online Library.
- Sandler T, Arce M, Daniel G (2002) A conceptual framework for understanding global and transnational public goods for health. *Fiscal Studies* 23:195–222
- Springborn M, Costello C, Ferrier P (2010) Optimal random exploration for trade-related non-indigenous. In: *Bioinvasions and Globalization: Ecology, Economics, Management, and Policy*, Perrings C, Mooney H, Williamson M (editors), Oxford: Oxford University Press, pp 127–144
- Springborn M, Chowell G, MacLachlan M, Fenichel EP (2015) Accounting for behavioral responses during a flu epidemic using home television viewing. *BMC infectious diseases* 15:21
- Taylor LH, Latham SM, Mark E (2001) Risk factors for human disease emergence. *Philosophical Transactions of the Royal Society of London B: Biological Sciences* 356:983–989
- Taylor MA, Ampt ES (2003) Travelling smarter down under: policies for voluntary travel behaviour change in Australia. *Transport Policy* 10:165–177
- Team WER (2014) Ebola virus disease in West Africa—the first 9 months of the epidemic and forward projections. *N Engl J Med* 371:1481–1495
- Weycker D, Edelsberg J, Elizabeth Halloran M, Longini IM Jr, Nizam A, Ciuryla V, Oster G (2005) Population-wide benefits of routine vaccination of children against influenza. *Vaccine* 23:1284–1293